T·L·G
Texts in Logic and Games
Volume 3

Logic and the Foundations of Game and Decision Theory (LOFT 7)

EDITED BY
GIACOMO BONANNO, WIEBE VAN DER HOEK AND
MICHAEL WOOLDRIDGE

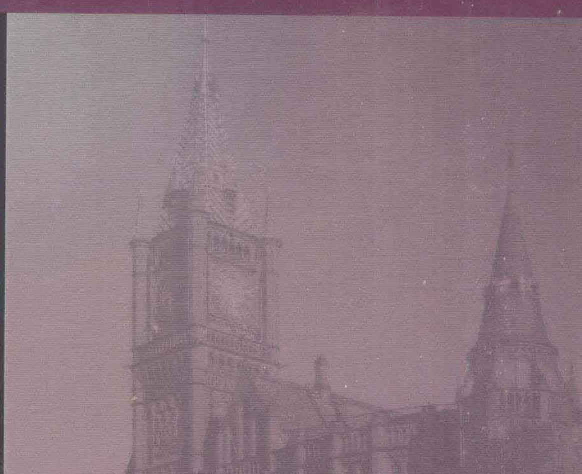

AMSTERDAM UNIVERSITY PRESS

LOGIC AND THE FOUNDATIONS OF GAME AND DECISION THEORY (LOFT 7)

T·L·G
Texts in Logic and Games
Volume 3

Logic and the Foundations of Game and Decision Theory (LOFT 7)

EDITED BY
GIACOMO BONANNO
WIEBE VAN DER HOEK
MICHAEL WOOLDRIDGE

Texts in Logic and Games
Volume 3

AMSTERDAM UNIVERSITY PRESS

Cover design: Maedium, Utrecht

ISBN 978 90 8964 026 0
NUR 918

Table of Contents

Preface

This volume in the *Texts in Logic and Games* series was conceived as a ramification of the seventh conference on *Logic and the Foundations of the Theory of Games and Decisions* (LOFT7), which took place in Liverpool, in July 2006.[1]

The LOFT conferences have been a regular biannual event since 1994. The first conference was hosted by the Centre International de Recherches Mathematiques in Marseille (France), the next four took place at the International Centre for Economic Research in Torino (Italy), the sixth conference was hosted by the Graduate School of Management in Leipzig (Germany) and the most recent one took place at the University of Liverpool (United Kingdom).[2]

The LOFT conferences are interdisciplinary events that bring together researchers from a variety of fields: computer science, economics, game theory, linguistics, logic, multi-agent systems, psychology, philosophy, social choice and statistics. In its original conception, LOFT had as its central theme the application of logic, in particular modal epistemic logic, to foundational issues in the theory of games and individual decision-making. Epistemic considerations have been central to game theory for a long time. The

[1] The conference was organized by the editors of this volume with the assistance of a program committee consisting of Thomas Ågotnes, Johan van Benthem, Adam Brandenburger, Hans van Ditmarsch, Jelle Gerbrandy, Wojtek Jamroga, Hannes Leitgeb, Benedikt Löwe, Marc Pauly, Andrés Perea, Gabriella Pigozzi, Wlodek Rabinowicz, Hans Rott, and Krister Segerberg.

[2] Collections of papers from previous LOFT conferences can be found in a special issue of *Theory and Decision* (Vol. 37, 1994, edited by M. Bacharach and P. Mongin), the volume *Epistemic logic and the theory of games and decisions* (edited by M. Bacharach, L.-A. Gérard-Varet, P. Mongin and H. Shin and published by Kluwer Academic, 1997), two special issues of *Mathematical Social Sciences* (Vols. 36 and 38, 1998, edited by G. Bonanno, M. Kaneko and P. Mongin), two special issues of *Bulletin of Economic Research* (Vol. 53, 2001 and Vol. 54, 2002, edited by G. Bonanno and W. van der Hoek), a special issue of *Research in Economics*, (Vol. 57, 2003, edited by G. Bonanno and W. van der Hoek), a special issue of *Knowledge, Rationality and Action* (part of *Synthese*, Vol. 147, 2005, edited by G. Bonanno) and the volume *Proceedings of the 7th Conference on Logic and the Foundations of Game and Decision Theory* (edited by G. Bonanno, W. van der Hoek and M. Wooldridge, University of Liverpool, 2006).

expression "interactive epistemology" has been used in the game-theory literature to refer to the analysis of decision making by agents involved in a strategic interaction, when these agents recognize each other's intelligence and rationality. What is relatively new is the realization that the tools and methodologies that were used in game theory are closely related to those already used in other fields, notably computer science and philosophy. Modal logic turned out to be the common language that made it possible to bring together different professional communities. It became apparent that the insights gained and the methodologies employed in one field could benefit researchers in other fields. Indeed, new and active areas of research have sprung from the interdisciplinary exposure provided by the LOFT conferences.[3]

Over time the scope of the LOFT conference has broadened to encompass a wider range of topics, while maintaining its focus on the general issue of rationality and agency. Topics that have fallen within the LOFT umbrella include epistemic and temporal logic, theories of information processing and belief revision, models of bounded rationality, non-monotonic reasoning, theories of learning and evolution, mental models, etc.

The papers collected in this volume reflect the general interests and interdisciplinary scope of the LOFT conferences.

The paper by Alexandru Baltag and Sonja Smets falls within the recent literature that deals with belief revision and update within the Dynamic Epistemic Logic paradigm. The authors develop a notion of doxastic action general enough to cover many examples of multi-agent communication actions encountered in the literature, but also flexible enough to deal with both static and dynamic belief revision. They discuss several epistemic notions: knowledge, belief and conditional belief. For the latter they distinguish between the statement 'if informed that P, the agent would believe that Q *was* the case (before the learning)' and the statement 'if informed that P, the agent would come to believe that Q *is* the case (in the world after the learning)'. They also study a "safe belief" operator meant to express a weak notion of "defeasible knowledge": it is belief that is persistent under revision with any true information. Baltag and Smets provide a complete axiomatization of the logic of conditional belief, knowledge and safe belief. In the second part of the paper the authors discuss dynamic belief revision in the context of action models.

The paper by Giacomo Bonanno deals with the question of what choices are compatible with rationality of the players and common belief of rationality. He takes a syntactic approach and defines rationality axiomatically.

[3] There is substantial overlap between the LOFT community and the community of researchers who are active in another regular, biannual event, namely the conferences on Theoretical Aspects of Rationality and Knowledge (TARK).

Furthermore, he does not assume von Neumann-Morgenstern payoffs but merely ordinal payoffs, thus aiming for a more general theory of rationality in games. The author considers two axioms. The first says that a player is irrational if she chooses a particular strategy while believing that another strategy of hers is better. He shows that common *belief* of this weak notion of rationality characterizes the iterated deletion of pure strategies that are strictly dominated by another pure strategy. The second axiom says that a player is irrational if she chooses a particular strategy while believing that a different strategy is at least as good and she considers it possible that this alternative strategy is actually better than the chosen one. The author shows that common *knowledge* of this stronger notion of rationality characterizes the iterated deletion procedure introduced by Stalnaker (1994), restricted—once again—to pure strategies.

The paper by Hans van Ditmarsch and Barteld Kooi investigates a dynamic logic describing "epistemic events" that may change both the agents' information (or beliefs) and what the authors call "the ontic facts" of the world (that is, objective, non-epistemic statements about the world). A sound and complete axiomatization is provided. Some original and interesting semantic results are also proved, in particular the fact that any model change can be simulated by "epistemic events", and thus any consistent goal can be achieved by performing some such event. The authors illustrate their results in several examples, including card games and logical puzzles.

The paper by Wiebe van der Hoek, Mark Roberts and Michael Wooldridge extends the authors' previous work on Alternating-time Temporal Logic and its ramifications. They extend it by introducing the notion of a legally possible strategy, that they oppose to a physically possible strategy, and define social belief as truth in all states that are (1) possible for the agent, and (2) are obtained from the initial state by a legally possible strategy. They use this framework to reason about social laws. In a system with social laws, every agent is supposed to refrain from performing certain forbidden actions. Rather than assuming that all agents abide by the law, the authors consider what happens if certain agents act socially, while others do not. In particular, they focus on the agents' strategic abilities under such mixed conditions.

The paper by Alexander Nittka and Richard Booth deals with the traditional "static" belief revision setting, but with a different twist: rather than answering the question of how an agent should rationally change his beliefs in the light of new information, they address the question of what one can say about an agent who is observed in a belief change process. That is, the authors study the problem of how to make inferences about an agent's beliefs based on observation of how that agent responded to a sequence of revision inputs over time. They start by reviewing some earlier results for the case where the observation is complete in the sense that (1) the logical content of all formulas appearing in the observation is known, and (2) all revision inputs received by the agent during the observed period are recorded in the observation. They then provide new results for the more general case where information in the observation might be distorted due to noise or because some revision inputs are missing altogether. Their analysis is based on the assumption that the agent employs a specific, but plausible, belief revision framework when incorporating new information.

The paper by R. Ramanujam and Sunil Simon deals with the most important notion of non-cooperative game, namely extensive game. Extensive games provide a richer description of interactive situations than strategic-form games in that they make the order of moves and the information available to a player when it is his turn to move explicit. A strategy for a player in an extensive game associates with every information set of that player a choice at that information set. The authors observe that the game position (or information set) may be only partially known, in terms of properties that the player can test for. Thus—they argue—strategies can be thought of as programs, built up systematically from atomic decisions like *if* b *then* a where b is a condition checked by the player to hold (at some game position) and a is a move available to the player at that position. This leads them to propose a logical structure for strategies, where one can reason with assertions of the form "(partial) strategy σ ensures the (intermediate) condition α". They present an axiomatization for the logic and prove its completeness.

The paper by Giacomo Sillari contributes to the very recent and fast growing literature on the notion of (un)awareness. An open problem in this literature has been how to model the state of mind of an individual who realizes that he may be unaware of something, that is, the problem of formalizing the notion of "awareness of unawareness". Sillari offers a solution to this problem using a new system of first-order epistemic logic with awareness. He also offers a philosophical analysis of awareness structures and proves that a certain fragment of the first-order epistemic language with awareness operators is decidable.

The papers went through a thorough refereeing and editorial process. The editors would like to thank the many referees who provided invaluable help and the authors for their cooperation during the revision stage.

Davis, CA & Liverpool G.B. W.v.d.H. M.W.

A Qualitative Theory of Dynamic Interactive Belief Revision

Alexandru Baltag[1]
Sonja Smets[2,3]

[1] Computing Laboratory
Oxford University
Oxford OX1 3QD, United Kingdom

[2] Center for Logic and Philosophy of Science
Vrije Universiteit Brussel
Brussels, B1050, Belgium

[3] IEG Research Group on the Philosophy of Information
Oxford University
Oxford OX1 3QD, United Kingdom

baltag@comlab.ox.ac.uk, sonsmets@vub.ac.be

Abstract

We present a logical setting that incorporates a belief-revision mechanism within Dynamic-Epistemic logic. As the "static" basis for belief revision, we use *epistemic plausibility models*, together with a modal language based on *two epistemic operators*: a *"knowledge" modality* K (the standard S5, fully introspective, notion), and a *"safe belief" modality* $\Box$ ("weak", non-negatively-introspective, notion, capturing a version of Lehrer's "indefeasible knowledge"). To deal with "dynamic" belief revision, we introduce *action plausibility models*, representing various types of "doxastic events". Action models "act" on state models via a modified update product operation: the *"Action-Priority" Update*. This is the natural dynamic generalization of AGM revision, giving priority to the incoming information (i.e., to "actions") over prior beliefs. We completely axiomatize this logic, and show how our update mechanism can "simulate", in a uniform manner, *many different belief-revision policies*.

1 Introduction

This paper contributes to the recent and on-going work in the logical community [2, 14, 24, 8, 10, 9, 7] on dealing with mechanisms for belief revision and update within the Dynamic-Epistemic Logic (DEL) paradigm. DEL originates in the work of Gerbrandy and Groeneveld [30, 29], anticipated by Plaza in [44], and further developed by numerous authors

Giacomo Bonanno, Wiebe van der Hoek, Michael Wooldridge (eds.). *Logic and the Foundations of Game and Decision Theory (LOFT 7).* Texts in Logic and Games 3, Amsterdam University Press 2008, pp. 11–58.

[6, 31, 22, 4, 23, 39, 5, 15, 16] etc. In its standard incarnation, as presented e.g., in the recent textbook [25], the DEL approach is particularly well fit to deal with *complex multi-agent learning actions* by which groups of interactive agents update their beliefs (including *higher-level beliefs* about the others' beliefs), *as long as the newly received information is consistent with the agents' prior beliefs.* On the other hand, the classical AGM theory and its more recent extensions have been very successful in dealing with the problem of *revising one-agent, first-level (factual) beliefs when they are contradicted by new information.* So it is natural to look for a way to combine these approaches.

We develop here a notion of *doxastic actions*[1], general enough to cover most examples of multi-agent communication actions encountered in the literature, but also flexible enough to deal with *(both static and dynamic) belief revision*, and in particular to *implement various "belief-revision policies" in a unified setting.* Our approach can be seen as a natural extension of the work in [5, 6] on "epistemic actions", incorporating ideas from the AGM theory along the lines pioneered in [2] and [24], but using a *qualitative* approach based on *conditional beliefs*, in the line of [50, 20, 19, 14].

Our paper assumes the general distinction, made in [24, 8, 14], between *"dynamic" and "static" belief revision.* It is usually acknowledged that the classical AGM theory in [1, 28] (and embodied in our setting by the *conditional belief* operators $B_a^P Q$) is indeed "static", in the sense that it captures *the agent's changing beliefs about an unchanging world.* But in fact, when we take into account all the higher-level beliefs, the "world" (that these higher-level beliefs are about) includes all agent's (real) beliefs.[2] Thus, such a world is *always changed by our changes of beliefs*! So we can better understand a belief conditional on P as capturing the agent's beliefs *after revising with* P about the state of the world *before the revision*: the statement $B_a^P Q$ says that, *if agent a would learn* P, *then she would come to believe that* Q *was the case (before the learning).* In contrast, "dynamic" belief revision uses dynamic modalities to capture the agent's revised beliefs about the world *as it is after revision*: $[!P]B_a Q$ says that *after learning* P, *agent a would come to believe that* Q *is the case (in the world after the learning).* The standard alternative [37] to the AGM theory calls this *belief update*, but like the AGM approach, it only deals with "first-level" (factual) beliefs from a non-modal perspective, neglecting any higher-order "beliefs about beliefs". As a result, *it completely misses the changes induced* (in our own or the other agents' epistemic-doxastic states) *by the learning actions themselves* (e.g., the learning of a Moore sentence, see Section 3). This

[1] Or "doxastic events", in the terminology of [14].

[2] To verify that a higher-level belief about another belief is "true" we need to check the content of that higher-level belief (i.e., the existence of the second, lower-level belief) against the "real world". So the real world has to include the agent's beliefs.

is reflected in the acceptance in [37] of the AGM "Success Axiom": in dynamic notation, this is the axiom $[!P]B_aP$ (which cannot accommodate Moore sentences). Instead, the authors of [37] exclusively concentrate on the possible changes of (ontic) facts that may have occurred during our learning (but *not due to our learning*). In contrast, our approach to belief update (following the DEL tradition) may be thought of as "dual" to the one in [37]: we completely neglect here the ontic changes[3], considering only the changes induced by *"purely doxastic" actions* (learning by observation, communication, etc.).

Our formalism for "static" revision can best be understood as a modal-logic implementation of the well-known view of belief revision in terms of *conditional reasoning* [50, 52]. In [8] and [10], we introduced two equivalent semantic settings for conditional beliefs in a multi-agent epistemic context (*conditional doxastic models* and *epistemic plausibility models*), taking the first setting as the basic one. Here, we adopt the second setting, which is closer to the standard semantic structures used in the literature on modeling belief revision [34, 49, 52, 27, 19, 14, 17]. We use this setting to define notions of *knowledge* K_aP, *belief* B_aP and *conditional belief* B_a^QP. Our concept of "knowledge" is the standard S5-notion, partition-based and fully introspective, that is commonly used in Computer Science and Economics, and is sometimes known as "Aumann knowledge", as a reference to [3]. The conditional belief operator is a way to "internalize", in a sense, the "static" (AGM) belief revision within a modal framework: saying that, at state s, agent a believes P conditional on Q is a way of saying that Q belongs to a's revised "theory" (capturing her revised beliefs) after revision with P (of a's current theory/beliefs) at state s. Our conditional formulation of "static" belief revision is close to the one in [50, 47, 19, 20, 45]. As in [19], the preference relation is assumed to be well-preordered; as a result, the logic CDL of conditional beliefs is equivalent to the strongest system in [19].

We also consider other modalities, capturing other "doxastic attitudes" than just knowledge and conditional belief. The most important such notion expresses a form of "weak (non-introspective) knowledge" $\Box_aP$, first introduced by Stalnaker in his modal formalization [50, 52] of Lehrer's *defeasibility analysis of knowledge* [40, 41]. We call this notion *safe belief*, to distinguish it from our (Aumann-type) concept of knowledge. Safe belief can be understood as belief that is *persistent under revision with any true information.* We use this notion to give a new solution to the so-called "Paradox of the Perfect Believer". We also solve the open problem posed in [19], by providing a *complete axiomatization of the "static" logic $K\Box$ of conditional belief, knowledge and safe belief.* In a forthcoming paper, we

[3] But our approach can be easily modified to incorporate ontic changes, along the lines of [15].

apply the concept of safe belief to Game Theory, improving on Aumann's epistemic analysis of backwards induction in games of perfect information.

Moving thus on to *dynamic belief revision*, the first thing to note is that (unlike the case of "static" revision), *the doxastic features of the actual "triggering event"* that induced the belief change *are essential* for understanding this change (as a "dynamic revision", i.e., in terms of the revised beliefs about the state of the world after revision). For instance, our beliefs about *the current situation after* hearing a *public* announcement (say, of some *factual* information, denoted by an atomic sentence p) are different from our beliefs after receiving a *fully private* announcement with the same content p. Indeed, in the public case, we come to believe that p is now *common knowledge* (or at least *common belief*). While, in the private case, we come to believe that the content of the announcement forms now our *secret knowledge*. So the agent's *beliefs about the learning actions* in which she is currently engaged affect the way she updates her previous beliefs.

This distinction is irrelevant for "static" revision, since e.g., in both cases above (public as well as private announcement) we learn the same thing about the situation that existed *before the learning*: our beliefs about that past situation will change in the same way in both cases. More generally, our beliefs about the "triggering action" are irrelevant, as far as our "static" revision is concerned. This explains a fact observed in [14], namely that by and large, the standard literature on belief revision (or belief update) *does not usually make explicit the doxastic events that "trigger" the belief change* (dealing instead only with types of abstract operations on beliefs, such as update, revision and contraction etc). The reason for this lies in the "static" character of AGM revision, as well as its restriction (shared with the "updates" of [37]) to one-agent, first-level, factual beliefs.

A "truly dynamic" logic of belief revision has to be able to capture the *doxastic-epistemic features (e.g., publicity, complete privacy etc.) of specific "learning events"*. We need to be able to model the agents' "dynamic beliefs", i.e., their *beliefs about the learning action itself*: the *appearance* of this action (while it is happening) to each of the agents. In [5], it was argued that a natural way to do this is to use *the same type of formalism that was used to model "static" beliefs*: *epistemic actions should be modeled in essentially the same way as epistemic states*; and this common setting was taken there to be given by *epistemic Kripke models.*

A similar move is made here in the context of our richer doxastic-plausibility structures, by introducing *plausibility pre-orders on actions* and developing a notion of "action plausibility models", that extends the "epistemic action models" from [5], along similar lines to (but without the quantitative features of) the work in [2, 24].

Extending to (pre)ordered models the corresponding notion from [5], we introduce an operation of *product update* of such models, based on the *anti-lexicographic order* on the product of the state model with the action model. The simplest and most natural way to define a connected pre-order on a Cartesian product from connected pre-orders on each of the components is to use either the *lexicographic* or the *anti-lexicographic* order. Our choice is the second, which we regard as the *natural generalization of the AGM theory*, giving *priority to incoming information* (i.e., to "actions" in our sense). This can also be thought of as a generalization of the so-called *"maximal-Spohn" revision*. We call this type of update rule the *"Action-Priority" Update*. The intuition is that the beliefs encoded in the action model express *the "incoming" changes of belief*, while the state model only captures that *past beliefs*. One could say that the new "beliefs about actions" are *acting* on the prior "beliefs about states", producing the updated (posterior) beliefs. This is embedded in the Motto of Section 3.1: *"beliefs about changes encode (and induce) changes of beliefs"*.

By abstracting away from the quantitative details of the plausibility maps when considering the associated *dynamic logic*, our approach to dynamic belief revision is in the spirit of the one in [14]: instead of using "graded belief" operators as in e.g., [2, 24], or probabilistic modal logic as in [39], both our account and the one in [14] concentrate on the simple, qualitative language of *conditional beliefs, knowledge and action modalities* (to which we add here the *safe belief* operator). As a consequence, we obtain *simple, elegant, general logical laws of dynamic belief revision*, as natural generalizations of the ones in [14]. These "reduction laws" give a *complete axiomatization of the logic of doxastic actions*, "reducing" it to the "static" logic $K\Box$. Compared both to our older axiomatization in [10] and to the system in [2], one can easily see that the introduction of the safe belief operator leads to a major simplification of the reduction laws.

Our qualitative logical setting (in this paper and in [8, 10, 9]), as well as van Benthem's closely related setting in [14], are conceptually very different from the more "quantitative" approaches to dynamic belief revision taken by Aucher, van Ditmarsch and Labuschagne [2, 24, 26], approaches based on "degrees of belief" given by ordinal plausibility functions. This is not just a matter of interpretation, but it makes a difference for the choice of dynamic revision operators. Indeed, the update mechanisms proposed in [49, 2, 24] are essentially quantitative, using various binary functions in transfinite ordinal arithmetic, in order to compute the degree of belief of the output-states in terms of the degrees of the input-states and the degrees of the actions. This leads to an increase in complexity, both in the computation of updates and in the corresponding logical systems. Moreover, there seems to be no canonical choice for the arithmetical formula for updates, various authors proposing various formulas. No clear intuitive justification is provided to any of these formulas, and we see no transparent reason to prefer one to the others. In contrast, classical (AGM) belief revision theory is a qualitative theory, based on natural, intuitive postulates, of great generality and simplicity.

Our approach retains this qualitative flavor of the AGM theory, and aims to build a theory of "dynamic" belief revision of equal simplicity and naturality as the classical "static" account. Moreover (unlike the AGM theory), it aims to provide a *"canonical"* choice for a dynamic revision operator, given by our "Action Priority" update. This notion is a *purely qualitative one*[4], based on a *simple, natural relational definition.* From a *formal point of view*, one might see our choice of the anti-lexicographic order as *just one of the many possible options* for developing a belief-revision-friendly notion of update. As already mentioned, it is a generalization of the "maximal-Spohn" revision, already explored in [24] and [2], among many other possible formulas for combining the "degrees of belief" of actions and states. But here we justify our option, arguing that our *qualitative interpretation of the plausibility order makes this the only reasonable choice.*

It may seem that by making this choice, we have confined ourselves to *only one of the bewildering multitude of "belief revision policies"* proposed in the literature [49, 45, 48, 2, 24, 17, 14]. But, as argued below, *this apparent limitation is not so limiting after all*, but can instead be regarded as an *advantage*: the power of the "action model" approach is reflected in the fact that *many different belief revision policies* can be recovered as *instances of the same type of update operation.* In this sense, our approach can be seen as a *change of perspective*: the diversity of possible revision policies is replaced by the diversity of possible action models; the differences are now viewed as *differences in input, rather than having different "programs"*. For a computer scientist, this resembles "Currying" in lambda-calculus: if every "operation" is encoded as an input-term, then *one operation* (functional application) *can simulate all operations.*[5] In a sense, this is nothing but the idea of Turing's universal machine, which underlies universal computation.

The title of our paper is a paraphrase of Oliver Board's "Dynamic Interactive Epistemology" [19], itself a paraphrase of the title ("Interactive Epistemology") of a famous paper by Aumann [3]. We interpret the word "interactive" as referring to the *multiplicity of agents* and the *possibility*

---

[4] One could argue that our plausibility pre-order relation is equivalent to a quantitative notion (of ordinal degrees of plausibility, such as [49]), but unlike in [2, 24] the way belief update is defined in our account does not make any use of the ordinal "arithmetic" of these degrees.

[5] Note that, as in untyped lambda-calculus, the input-term encoding the operation (i.e., our "action model") and the "static" input-term to be operated upon (i.e., the "state model") are essentially *of the same type*: epistemic plausibility models for the same language (and for the same set of agents).