Belief Revision Meets Philosophy of Science

LOGIC, EPISTEMOLOGY, AND THE UNITY OF SCIENCE

VOLUME 21

Logic, Epistemology, and the Unity of Science aims to reconsider the question of the unity of science in light of recent developments in logic. At present, no single logical, semantical or methodological framework dominates the philosophy of science. However, the editors of this series believe that formal techniques like, for example, independence friendly logic, dialogical logics, multimodal logics, game theoretic semantics and linear logics, have the potential to cast new light on basic issues in the discussion of the unity of science.

This series provides a venue where philosophers and logicians can apply specific technical insights to fundamental philosophical problems. While the series is open to a wide variety of perspectives, including the study and analysis of argumentation and the critical discussion of the relationship between logic and the philosophy of science, the aim is to provide an integrated picture of the scientific enterprise in all its diversity.

For further volumes:
http://www.springer.com/series/6936

Erik J. Olsson · Sebastian Enqvist
Editors

Belief Revision Meets Philosophy of Science

Springer

Editors
Prof. Erik J. Olsson
University of Lund
Dept. Philosophy
Kungshuset
222 22 Lund
Sweden
erik_j.olsson@fil.lu.se

Sebastian Enqvist
University of Lund
Dept. Philosophy
Kungshuset
222 22 Lund
Sweden
Sebastian.Enqvist@fil.lu.se

ISBN 978-90-481-9608-1 e-ISBN 978-90-481-9609-8
DOI 10.1007/978-90-481-9609-8
Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2010938250

Printed on acid-free paper

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

Editor's Introduction

Belief revision theory and philosophy of science both aspire to shed light on the dynamics of knowledge – on how our view of the world changes (typically) in the light of new evidence. Yet these two areas of research have long seemed strangely detached from each other, as witnessed by the small number of cross-references and researchers working in both domains. One may speculate as to what has brought about this surprising, and perhaps unfortunate, state of affairs. One factor may be that while belief revision theory has traditionally been pursued in a bottom-up manner, focusing on the endeavors of single inquirers, philosophers of science, inspired by logical empiricism, have tended to be more interested in science as a multi-agent or agent-independent phenomenon.

The aim of this volume is to build bridges between these two areas of study, the basic question being how they can inform each other. The contributors seek their answers by relating the logic of belief revision to such concepts as explanation, coherence, induction, abduction, interrogative logic, conceptual spaces, structuralism, idealization, research agendas, minimal change and informational economy.

Our aim in putting together this volume has been to provide a number of new perspectives that are likely to stir research in new directions, as well as to establish new connections between areas previously assumed unrelated, e.g. between belief revision, conceptual spaces and structuralism. The result is, we believe, a coherent volume of individual papers complementing and shedding light on each other.

We have been very fortunate to be able to attract, as contributors to this volume, some of the best researchers in their respective fields of philosophy, cognitive science and logic, as well as some exceptional scholars in the younger generation. We are extremely proud to have their articles in the volume, and we thank them all for their dedication and fine scholarship.

We hope that this volume will contribute to a greater degree of interaction between the fields of belief revision theory and the philosophy of science. For this reason, we hope that the essays included here will be read by researchers in both fields. However, the fundamental concepts in the philosophy of science are probably known to a significantly wider audience than belief revision theory. For this reason, in order to facilitate for those readers not previously acquainted with belief revision theory, a brief introduction seems in order.

Belief revision theory is a branch of formal epistemology that studies rational changes in states of belief, or *rational theory changes*. The classic framework here, and the best known one, is the socalled AGM theory, named after its creators Carlos Alchourrón, Peter Gärdenfors and David Makinson, which originated in a series of papers in the eighties. Since then, a rather large number of alternative frameworks have emerged, which generalize or deviate from AGM in various ways. But for the purpose of introducing the novice to belief revision theory, a summary of the AGM theory will be enough to give a sense of how belief revision works and what it is about.

The basic way of representing belief states in the AGM theory is to equate an agent's state of belief (at some given time) with a logically closed set of sentences K, i.e. such that K = Cn(K) where Cn denotes the operation of logical closure. Thus, the logical consequences of an agent's beliefs are also counted as beliefs in AGM. Such a set K is sometimes referred to as a "theory", sometimes as a "belief set".

Let α = "Charles is in his office" and let β = "Charles is at home". Let $\mathrm{K}_1 = \mathrm{Cn}(\{\sim\alpha\rightarrow\beta\})$. Then K_1 expresses the state of belief where it is believed that "if Charles is not in his office, then he is at home". Now, say that we learn that Charles is not in his office. Then we need to alter our initial belief state K_1 to include this new belief. How do we do this, rationally? Simply adding the sentence $\sim\alpha$ to K_1 will not do, since the set $K_1 \cup \{\sim\alpha\}$ is not logically closed, i.e. it is not a *bona fide* belief set. What we need to do is to first add $\sim\alpha$ to K_1, and then close the result under logical consequences. This general recipe gives rise to the operation of *expansion*, one of the three basic operations on belief sets in AGM. The expansion of a belief set K by a sentence α, denoted $K + \alpha$, is defined by setting

$$K + \alpha =_{\text{df.}} \mathrm{Cn}\,(K \cup \{\alpha\})$$

In our case, the new belief set $K_1 + \sim\alpha$ after learning that Charles is not in his office is $\mathrm{Cn}\,(\{\sim\alpha \rightarrow \beta, \sim\alpha\})$. This new belief set includes the sentence β, i.e. upon learning that Charles is not in his office, we come to believe that Charles is not at home.

Let us denote this new belief set by K_2. Say that we now learn that Charles is not at home. Then we want to add the sentence $\sim\beta$ to K_2. If we expand again at this point, we run into some trouble: since β is in K_2, the expanded set $K_2 + \sim\beta$ contains a contradiction, and so by logical closure (assuming that the underlying logic contains the classical validities) $K_2 + \sim\beta$ will contain *all* sentences in the language – it is a state where we believe everything. This awkward situation has been referred to as "epistemic hell", and is clearly unattractive. We would like to avoid this consequence, and update the initial state K_2 with $\sim\beta$ while preserving consistency. This enters us into the area where belief revision theory becomes interesting.

The operation used for updating a belief state while maintaing consistency in cases like the one above is called *revision*. The revision of a belief set K by a sentence α is denoted $K^*\alpha$. A useful way of analysing the problem of revision is by introducing a third operator: *contraction*. Take the belief set K_2 again; we would like to add $\sim\beta$ to this belief set without creating an inconsistency. In order to do this, we have to remove something, in particular we have to remove the negation of $\sim\beta$, or equivalently, β. In order to do this, we have to remove some of our earlier beliefs, $\sim\alpha$ and $\sim \alpha \rightarrow \beta$ since they jointly imply β. Now, we could remove *both* these beliefs of course – this would certainly make room for $\sim\beta$ – but that would not be very *economical*, since removing either of $\sim\alpha$ or $\sim \alpha \rightarrow \beta$ while keeping the other would suffice. Beliefs are valuable things, and a rational agent should not be willing to give up more of his beliefs than what is necessary. This intuition is one of the guiding principles in AGM, and usually goes under the name of *the principle of minimal change*. Contraction is intended as an operation that removes a sentence from a given belief set in accordance with this principle.

The contraction of a belief set K by a sentence α is denoted $K \div \alpha$. Given a suitable operator $\div$ of contraction, we can define a revision operator * by setting, for all α:

$$K^{*}\alpha =_{\text{df}} .(K\div \sim \alpha) + \alpha$$

Intuitively: to revise K by α is to first remove the negation of α (to "make room" for α) and then expand with α. The definition is commonly known as the *Levi identity*, after Isaac Levi.

With this definition in place, the problem of revision reduces to the problem of contraction. Thus we are left with the task of devising a satisfactory account of contraction. The way AGM handles this problem can be divided into two distinct approaches: we may call one the *axiomatic* approach, and the other the *constructive* approach. The axiomatic approach consists in narrowing down the class of rationally admissible contraction functions by setting up a list of (intuitively plausible) postulates for contraction. The following six postulates are known as the *basic AGM postulates* for contraction:

(closure) $K \div \alpha = \text{Cn}(K \div \alpha)$
(success) $\alpha \notin \text{Cn}(\emptyset)$ implies $\alpha \notin \text{K} \div \alpha$
(inclusion) $K \div \alpha \subseteq K$
(vacuity) $\alpha \notin K$ implies $K \div \alpha = K$
(extensionality) $\text{Cn}(\alpha) = \text{Cn}(\beta)$ implies $K \div \alpha = K \div \beta$
(recovery) $K \subseteq (K \div \alpha) + \alpha$

This list is usually extended with the following two supplementary postulates:

(conjunctive inclusion) $\alpha \notin K \div (\alpha \wedge \beta)$ implies $K \div (\alpha \wedge \beta) \subseteq K \div \alpha$
(conjunctive overlap) $K \div \alpha \cap K \div \beta \subseteq K \div (\alpha \wedge \beta)$

The postulates all have some intuitive justification. For instance, the *success* postulate says that an admissible contraction operator $\div$ should do its job properly whenever possible, i.e. if α is not a tautology and so can be removed from K while maintaining logical closure, then $\div$ should succesfully remove it. The (highly controversial) *recovery* postulate gives a first formal expresssion to the principle of minimal change: when we contract by α, we should keep so much information that we can regain all our initial beliefs by expanding with α again. A similar set of postulates for revision exists, and these postulates are satisfied by any revision function which is defined from a contraction function satisfying the postulates above (and conversely, if a contraction function gives rise to a revision function that satisfies the AGM postulates for revision, then this function must satisfy the above postulates for contraction).

The *constructive* approach consists in devising explicit constructions of contraction functions. Apart from the socalled *partial meet* contractions, the construction most frequently discussed these days is probably the method of *entrenchment based* contraction. This method uses an auxiliary concept called an *entrenchment relation*, which is a binary relation over the language associated with a given belief set K, satisfying the following postulates:

(transitivity) $\alpha \leq \beta$ and $\beta \leq \chi$ implies $\alpha \leq \chi$
(dominance) $\beta \in \mathrm{Cn}(\alpha)$ implies $\alpha \leq \beta$
(conjunctiveness) either $\alpha \leq (\alpha \wedge \beta)$ or $\beta \leq (\alpha \wedge \beta)$
(minimality) if $\perp \notin K$ then $\alpha \notin K$ iff $\alpha \leq \beta$ for all $\beta \in \mathrm{L}$
(maximality) $\beta \leq \alpha$ for all β only if $\alpha \in \mathrm{Cn}(\emptyset)$

Given a belief set K with an associated entrenchment order $\leq$, we can define a corresponding entrenchment based contraction function $\div$ by setting, for all α:

$$K \div \alpha =_{\mathrm{df.}} \{\beta \in K \mid \alpha < \alpha \vee \beta \text{ or } \alpha \in \mathrm{Cn}(\emptyset)\}$$

The intuition here is that the entrenchment order encodes how entrenched the various beliefs in K are in comparison to each other, or which beliefs the agent would prefer to give up if a choice has to be made. Entrenchment based contractions are then designed so as to remove less entrenched beliefs in favor of the more entrenched ones.

How are the axiomatic approach and the constructive approach related to each other? It turns out that they are very closely related: every entrenchment based contraction satisfies the AGM postulates for contraction (including the supplementary postulates), and vice versa, if a contraction function satisfies the AGM postulates, then there exists an entrenchment order which defines it. In a sense, the AGM postulates are *sound and complete* with respect to entrenchment based contractions. This result is one of the celebrated *representation theorems* of the AGM framework.

Drafts of most of the papers that appear in this volume were originally presented at the first *Science in Flux* conference, organized by Olsson, at Lund University in 2007. (A follow-up conference was organized by Pierre Wagner in 2008 at the CNRS in Paris.) Before they were submitted in their final versions, the papers were revised, often substantially, in order to accommodate various critical points that emerged in the (very lively) discussion at the conference. It has been pointed out to us that there is already a book called "Science in Flux", by J. Agassi, which appears in the Boston Studies in the 1970s. As far as we can see, there is little overlap between the books, and little risk for conflating one with the other, and so we hope we are excused for reusing the title.

We shall briefly present each of the individual contributions:

Raúl Carnota and Ricardo Rodríguez

In their contribution, Raúl Carnota och Ricardo Rodríguez take a closer look at the history behind the influential AGM model of belief revision due to Alchourron, Gärdenfors and Makinson. The AGM theory, as described in the seminal 1985 paper "On The Logic of Theory Change: Partial Meet Contractions and Revision Functions", had a major influence in most subsequent work on belief change. In particular, the constructive approach spelled out in the paper was adopted in AI as a paradigm for how to specify updates of knowledge bases. Throughout the years there has been a steady stream of references to that original AGM paper. Going one step further, Carnota and Rodréguez ask themselves why the AGM theory was so readily accepted within the AI community, their answer being partly that the theory was put forward at a critical time in the history of AI at which the problem of how to update knowledge bases in the face of input possibly inconsistent with the previous corpus was taken to be of utmost importance. The paper also contains a qualitative and quantitative evaluation of the impact of the AGM theory in AI research as well as an account of how the theory has subsequently been developed in different directions.

Sven Ove Hansson

There is a clear connection between belief revision theory and one of the major problems within the philosophy of science, the problem of modelling and understanding the dynamics of empirical theories; both these fields of research deal with the way theories are updated in the light of incoming data. In fact, several of the most influential ideas in 20th century philosophy of science, e.g. Popper's hypothetico-deductive method, Kuhn's theory of paradigm shifts, Lakatos's ideas concerning the "hard core" and the "protective belt" etc. seem to be in essence theories about the dynamics of theories. That is, these theories apparently address the very subject matter of belief revision theory. Given this connection, it is somewhat striking that there has been so little contact between the two fields. Hansson draws the conclusion that belief revision theory as it stands is unsuitable for modelling changes in empirical theories, and sets out to develop a framework which is better suited for this task. He draws some first contours of a model where scientific change is treated as a *partly accumulative* process, through which observational data is added piecewise and theoretical hypotheses are added by a closure operator representing "inference to the best explanation". A set of postulates for this operator is provided, and three versions of a model of theory change are introduced and discussed in the text.

Hans Rott

Scientific change is also the main issue addressed by Hans Rott. In his paper, Rott poses the problem of how, exactly, to explicate the Lakatosian notion of a "progressive problem shift" that plays a crucial role in the understanding of how research programs develop over time, the basic idea being that such a shift is progressive if the transition to a successor theory T′ can somehow explain both the success of its predecessor theory T and the failure of T. That would mean that we would have an account that goes deeper than a plain approximate agreement of the empirical predictions made by the two theories. Rott proposes to accomplish this explication using factual, potential and counterfactual explanations. Thus the successor theory can explain the success of the predecessor theory by implying that the predecessor theory would have been true, had its application conditions been satisfied, but because they are not, the predecessor theory is false. This gives an account of how a single theory can speak, as it were, at the same time in favor of and against another theory. Rott uses the AGM postulates for rational belief changes to spell out these ideas in formal terms, thus connecting a central issue in the philosophy of science with standard theorizing in the logic of belief revision.

Gerhard Schurz

The article by Gerhard Schurz deals with the problem of abduction in the context of belief revision theory. It is noted that belief revision in its usual form lacks an account of the ability of inquiring agents to *learn* in the sense of forming generalized hypotheses on the basis of incoming data, a point which is illustrated with some formal theorems. If belief revision theory is to be applied to problems in the philosophy of science in a fruitful way, then extending the classical models of belief revision to encompass abduction seems to be of great importance: science is not simply the act of collecting and storing data. Arguably, the most important task of scientists is to formulate and test hypotheses on the basis of the empirical data, and any model which claims to capture scientific change must therefore provide an account of this process. After having discussed two attempts in the literature to incorporate such creative elements of belief formation into belief revision theory, Schurz develops an alternative theory based on a theory of abduction developed elsewhere. Various types of abductive revision and expansion are investigated; one of the main observations being that the well known Levi identity breaks down in the context of abductive belief revision.

Sebastian Enqvist

Within the AGM theory, and in belief revision more generally, it is assumed that theories can be represented as sets of sentences or statements. Within the philosophy of science, this seemingly harmless assumption has been challenged by the so-called structuralist theory of science, the seminal exposition of which is Joseph Sneed's 1971 book "The Logical Structure of Mathematical Physics". Structuralism instead reconstructs empirical theories as set theoretical structures, which have no propositional content in themselves, but which can be used to make empirically testable statements about the world. In his paper, Enqvist develops a model of theory change which is founded on the structuralist notion of a "theory net", rather than the classical conception as a set of sentences in a formal language. The notion of theory nets gives an explicit account of the "deep structure" of empirical theories, and Enqvist argues that the fine grained structure of the structuralist's way of representing theories may shed some new light on the problem of theory change. In particular, the specialization relation which forms an essential part of a theory net is investigated in the context of contraction. It is argued that specialization plays an important role in contraction, but that a separate notion of corroboration should also be taken into account. Finally, the possibility of distinguishing novel types of theory changes within the framework is discussed.

Peter Gärdenfors and Frank Zenker

The aim of the contribution by Gärdenfors and Zenker is to apply conceptual spaces, as developed by Gärdenfors in his 2000 book "Conceptual Spaces – The Geometry of Thought", to the dynamics of scientific theories. According to the theory of conceptual spaces, dimensions and their relations provide a topological representation of a concept's constituents and their mode of combination where concepts are seen as n-dimensional geometrical structures and conceptual change, consequently, means the dynamic development of these structures. Gärdenfors and Zenker take the structuralist framework, also addressed by Enqvist, to be a useful contrast to their own thinking on the matter, and in one section they argue that the central notions of structuralism can be expressed in terms of conceptual spaces. As they also observe, however, structuralism is problematic in the context of revolutionary changes. It is here that conceptual spaces may have a distinct worth: as Gärdenfors and Zenker argue, many, or even all, radical changes can be modeled as one of four types of increasingly severe transformations of conceptual spaces.

Bengt Hansson

Bengt Hansson begins his essay by pointing out a certain type of adequacy condition which seems to be present in mature theories like, for example, Newtonian mechanics and which may be demanded of a good theory: a condition which he calls *conceptual closure*. Intuitively, a theory can be said to satisfy conceptual closure if it describes a delineated part of the world which is closed in the sense that everything which influences the factors taken into account in the description is also explicitly part of the description. More formally, Hansson shows that the condition of conceptual closure can be characterized in terms of the notions of homomorphisms and commutative diagrams. While the notion of conceptual closure is derived from examples within the empirical sciences and physics in particular, Hansson argues that the same condition may be required of an adequate theory of belief change or theory change. Some of the well known approaches to modelling epistemic states in belief revision are questioned with regards to whether they satisfy conceptual closure, and the classical principle of minimal change is briefly discussed.

Horacio Arlo-Costa and Arthur Paul Pedersen

Horacio Arlo-Costa and Arthur Paul Pedersen's paper concerns the theory of rational choice and its application to belief revision theory. Correlations between principles of belief change and principles of rational choice have been studied at length by Hans Rott, in particular in his 2001 book "Change, Choice and Inference". This correlation has the nice feature of allowing us to think of principles of belief change in choice theoretic terms, and thus to criticize, assess or justify principles of belief change on choice theoretic grounds. This line of thought is continued in Arlo-Costa and Pedersen's paper, with focus on a particular issue in the theory of rational choice: traditional rational choice theory has been criticized by Amartya Sen for being unable to deal with the role *social norms* play in some cases. After having presented a recent model of rational choice which is intended to take social norms into account, Arlo-Costa and Pedersen develop an alternative model which generalizes the former one. The theory is then applied to belief revision theory. It is argued that the existence of social norms gives rise to counterexamples to some of the classical postulates for belief change, and they attempt to develop a model for "norm-inclusive belief change" which accounts for these examples. Axioms for norm-inclusive belief change are given and related to principles of rational choice by formal correspondence results.

David Westlund

Continuing the choice-theoretic theme, David Westlund's essay concerns the possibility of extending the AGM framework for belief revision to cover the case where collective agents change their beliefs as the result of individual belief changes. He points out the importance of this issue for applications of belief revision to the philosophy of science: scientific theories are not the beliefs of a single person and so changes in the beliefs of the scientific community are not changes in the beliefs of an individual, though what the scientific community may be said to believe is obviously somehow dependent on what individual scientists believe. That is, science is a *social* enterprise, and the dynamics of scientific theories thus contains a social component: to understand how theories change, it is not enough to understand how individual agents change their beliefs. We must also provide an account of how individual belief systems give rise to the beliefs of a scientific community, and how individual belief changes give rise to changes in what the community believes. In order to study this feature of scientific change, Westlund introduces the notion of a *merging function* (borrowed from computer science), which is a function taking a family of belief sets to a "merged" belief set. In terms of these merging functions, some negative results on collective belief change are demonstrated, showing that certain conditions on collective belief change cannot be consistently fulfilled.

Emmanuel Genot

Emmanuel Genot's paper builds on a proposal due to Erik J. Olsson and David Westlund, to extend the representation of epistemic states in the AGM theory to include an account of the research agenda of the inquiring agent. Genot relates this suggestion to the socalled *interrogative model of inquiry* (IMI) due to Jaakko Hintikka. Hintikka's model is a general model of inquiry, which is "Socratic" in the sense that inquiry is treated as a process of asking questions and drawing conclusions from the answers received. Formally, the interrogative model treats inquiry as a game where an agent may make *interrogative moves*, i.e. ask questions in order to retrieve new information, interspersed with *deductive moves*, where information is deduced from given premises together with answers received to previously asked questions. The connection between Hintikka's model and Olsson and Westlund's proposal to extend belief revision with a research agenda is clear: both argue that asking questions somehow plays a central role in inquiry. Apart from this informal similarity, Genot establishes a formal connection between the two approaches, so that results and research problems may be transferred from one to the other. In particular, he applies results from the interrogative model, most notably the so-called "Yes-No theorem", to attack the problem of updating agendas in the case of contraction – a problem which is largely left open in Olsson and Westlund's previous work on the agenda. Borrowing some game theoretical terminology, Genot distinguishes between "stategic" and "extensive" update of questions.

Erik J. Olsson

Olsson's point of departure is a longstanding issue in philosophy of science and belief revision theory concerning what to do when, after thorough investigation, more than one theory or belief set stands out as highly reasonable given all the data. Otto Neurath, the logical empiricist, suggested that it would, in such circumstances, be rationally admissible to decide the matter by coin-flipping. In belief revision theory, this debate has taken the form of a dispute between "functionalists" and "relationalists". Functionalists hold that the result of revising a cognitive state with some new datum is a unique rationally determined belief state. Relationalists, by means of contrast, insist that there may be several rationally admissible results. In an attempt to contribute to conceptual clarity, Olsson distinguishes between three ways of drawing the functionalist-relationalist distinction. This gives rise to six non-contradictory overall positions. He proceeds to consider arguments in the literature for excluding some of these positions on logical, philosophical or other grounds. Finally, Olsson argues that part of what feeds the functionalist relationalist controversy is a false dilemma based on an implausible conception of what it means rationally to suspend judgment. Making this precise requires a formal framework of the kind that includes a representation of the agent's research agenda. Here a connection emerges with the paper by Emmanuel Genot in which the notion of an agenda also plays a prominent role.

Caroline Semmling and Heinrich Wansing

Caroline Semmling and Heinrich Wansing, in their contribution to the present volume, investigate an extension of the socalled Stit Theory in which an operator of "deliberatively seeing to it that" (*dstit*) is investigated. This operator allows representation of deliberate actions of agents in a modal object language. Semmling and Wansing extend this theory by adding operators for beliefs, intentions and desires, forming what they call *bdi-stit* logic. They present a semantics for bdi-stit logic, as well as a complete tableaux-style proof system. With this extension in place, it is possible to express an interesting special class of actions: actions involving seeing-to-it-that you believe something, desire something or intend something. That is, we can express the formation of beliefs, desires or intentions as deliberate actions. In particular, the "belief" part of the logic becomes interesting in connection with belief change: seeing-to-it-that you believe in α looks a lot like *expanding* with α. Semmling and Wansing attempt to make this connection explicit, and apply their bdi-stit logic to AGM theory. It is shown how the language of *bdi-stit* logic can be used to define operators for expansion, revision and contraction. Since the belief fragment of the bdi-stit logic in itself contains only the means to express beliefs, on the one hand, and actions in terms of seeing-to-it-that something is brought about on the other, the result is an analysis of the concepts of expansion, contraction and revision in terms of deliberate actions. It turns out that, when translated with the help of these defined operators, several of the AGM postulates are provable in *bdi-stit* logic.

Isaac Levi

Isaac Levi's seminal work on belief revision theory is widely acclaimed and rightly so: many of the issues that the AGM theory dealt with can be traced back to their origin in Levi's work in the 1970s. This includes the classification of changes of belief in terms of expansions and contractions, and the famous so-called Levi identity according to which a revision can be reconstructed as a contraction followed by an expansion. Levi is also well known for his radical theory of knowledge according to which knowledge is nothing but true full belief, where belief is understood as "standard of serious possibility". This thesis, which runs counter to a long tradition in epistemology according to which knowledge entails that the knower has good reasons for his belief, is intimately related to his endorsement of Peirce's belief-doubt model stating that full belief is a state which is satisfactory in itself and therefore in no need of justification. For Levi, the issue of justification arises only in connection with *changes* of belief. Thus, while an inquirer can be justified or unjustified in expanding his corpus by a given item of information, once this expansion has been implemented there is no issue of justification anymore. In his contribution to this volume, Levi, among other things, usefully contrasts his view with conflicting accounts of knowledge in the recent epistemological literature, focusing on the theory advocated by Edward Craig, and he defends and further clarifies his taxonomy of different types of belief change and their decision-theoretic justification.

Paul Thagard

Paul Thagard has, in a number of papers and books, argued in favor of a coherence based approach to belief revision and reasoning in general. In his contribution to this volume he continues this line of work, showing how the recent debate about climate change in the scientific and political spheres can be modeled in his theory of explanatory coherence. The theory is based on a number of general coherence principles, such as "Similar hypotheses that explain similar pieces of evidence cohere" and "The acceptance of a proposition in a system of propositions depends on its coherence with them". While these principles do not fully determine coherence-based acceptance, Thagard has developed computer-implemented algorithms that can compute acceptance and rejection of propositions on the basis of coherence relations. In his paper, Thagard argues that his model can give a good account of the three main kinds of belief change: expansion, revision and contraction. For instance, "[e]xpansion takes place when a new proposition is introduced into a belief system, becoming accepted if and only if doing so maximizes coherence". Since the driving force behind Thagard's system is to satisfy the goal of maximizing coherence, coherence-based belief revision does not satisfy the AGM postulates, which are rather aimed at capturing minimal change. Thagard argues that this is as it should be, and that the AGM dictum to seek maximally conservative revisions is ill-motivated.

Thagard also maintains that his model has certain computational advantages over competing Bayesian accounts of belief change.

Jonas Nilsson and Sten Lindström

Belief revision theory is concerned with the rationality of changes of belief. The purpose of Jonas Nilsson and Sten Lindström's contribution is rather to inquire into the rationality of changes of methodology, as when a methodological rule such as "prefer simpler theories to less simple ones" is revised in favor of "a successor theory must retain all the corroborated empirical content of its predecessors". While there is a long-standing debate in philosophy of science concerning the proper account of methodological change and its rationality, few, if any of the resulting theories have reached the same level of precision as, for instance, the AGM theory of belief revision. For that reason, Nilsson and Lindström propose to investigate methodological change by means of formal methods. Special problems arise here because they would like to have a bootstrap theory rather than a static theory. According to the former, but not the latter, all standards are open for criticism and correction, meaning that there is no distinguished set of standards that can be used to correct other standards but which is itself immune to objections. While their paper is only a first step in this direction, they propose a number of general principles governing such change. According to their principle of Prospective Acceptability, for instance, a revised set of standards S_2 must be better than the original set S_1, according to the members of S_1, except for those standards in S_1 that are criticized and revised. Finally, Nilsson and Lindström discuss the prospects of modelling standards as a kind of beliefs, namely as beliefs about what rationality requires one to do. This would make methodological change a species of belief change, and raise the expectation that postulates for belief change should also hold for methodological change.

It strikes us that there is another way to connect the two areas of belief and methodological change, viz., to view beliefs, following Isaac Levi, as a certain kind of standards, i.e., as standards of serious possibility. This might be an interesting alternative for the following reason. As Nilsson and Lindström show, the AGM preservation principle is not plausible when viewed as a principle for methodological change, contrary to what we would expect if methodological standards are species of beliefs (and the AGM postulates are taken to apply to all kinds of belief change, rather than merely to, say, our beliefs about the world). However, if we instead view beliefs as species of standards, no such expectation seems to arise. From that perspective, the AGM axioms are about a special class of standards, namely, of standards of serious possibility, and there is little reason to believe that everything that can plausibly be said about the rationality with respect to that special class should be true of standard change in general. An interesting problem for further work would be how to weaken the AGM axioms (or other alternative sets of postulates) for changing standards of serious possibility, so that they hold for standards in general.

We mention this as but one example of how the papers in this volume, individually or in combination, may raise new and potentially fruitful research questions.

We believe there are a great many different angles from which the papers in this volume can be studied, and we invite the reader to be as creative as the contributors in deepening the various connections that emerge between belief revision theory and philosophy of science, or in thinking of new ways entirely for how these areas of philosophical and logical research can be brought into closer contact.

Lund, Sweden

Erik J. Olsson
Sebastian Enqvist

Contents

Contributors

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

Raúl Carnota Universidad Nacional de Tres de Febrero, Av. del Libertador Gral, San Martín 5569 Piso 14 ‘B’, 1426, Buenos Aires, Argentina, carnotaraul@gmail.com

Sebastian Enqvist Department of Philosophy, Lund University, Kungshuset Lundagård, 222 22 Lund, Sebastian.Enqvist@fil.lu.se

Scott Findlay University of Western Ontario, London, Ontario, sfindla@uwo.ca

Peter Gärdenfors Department of Philosophy, University of Lund, Kungshuset, Lundagard, 22 222 Lund, Sweden, peter.gardenfors@lucs.lu.se

Emmanuel Genot Department of Philosophy, University of Lille 3, Lille, France, emmanuel.genot@etu.univ-lille3.fr

Bengt Hansson Lund University, Lund, Sweden, bengt.hansson@fil.lu.se

Sven Ove Hansson Division of Philosophy, Royal Institute of Technology, Tekniktingen 78, 100 44 Stockholm, Sweden, soh@kth.se

Isaac Levi John Dewey Professor Emeritus, Columbia University, New York, NY, USA, levi@columbia.edu

Sten Lindström Department of Historical, Philosophical and Religious Studies, Umeå University, Sweden, sten.lindstrom@philos.umu.se

Jonas Nilsson Department of Historical, Philosophical and Religious Studies, Umeå University, Sweden, jonas.nilsson@philos.umu.se

Erik J. Olsson Lund University, Lund, Sweden, Erik_J.Olsson@fil.lu.se

Arthur Paul Pedersen Department of Philosophy, Carnegie Mellon University, Baker Hall 135, Pittsburgh, PA 15213, USA, apaulpedersen@cmu.edu

Ricardo Rodríguez Departamento de Computación de la Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Intendente Güiraldes 2160 - Ciudad Universitaria - C1428EGA, Buenos Aires, Argentina, ricardo@dc.uba.ar

Hans Rott Department of Philosophy, University of Regensburg, 93040 Regensburg, Germany, hans.rott@psk.uni-r.de

Gerhard Schurz University of Duesseldorf, Geb. 23.21, Universitaetsstrasse 1, D-40225 Duesseldorf, Germany, schurz@phil-fak.uni-duesseldorf.de

Caroline Semmling Institute of Philosophy, Dresden University of Technology, Dresden, Germany, Caroline.Semmling@gmx.de

Paul Thagard Department of Philosophy, University of Waterloo, Waterloo, ON, Canada, N2L 3G1, pthagard@uwaterloo.ca

Heinrich Wansing Institute of Philosophy, Dresden University of Technology, Dresden, Germany, Heinrich.Wansing@tu-dresden.de

David Westlund Fjärdingsvägen 29, 241 36 Eslöv, Sweden, david@westlund.fm

Frank Zenker Department of Philosophy, University of Lund, Kungshuset, Lundagard, 22 222 Lund, Sweden, frank.zenker@fil.lu.se