Introduction and Overview

Theo Kuipers · Gerhard Schurz

Received: 22 October 2010 / Accepted: 15 December 2010 / Published online: 28 June 2011
© Springer Science+Business Media B.V. 2011

1 Truth Approximation

It is a widespread view among more-or-less realistic philosophers of science that scientific progress consists in progress towards truth. This position has been elaborated within the fallibilistic program of Karl Popper, who emphasized that scientific theories are always conjectural and corrigible. Nevertheless scientific progress is possible insofar scientific theories, even if they have some false consequences, may be more or less close to the truth.

In (1963) Popper launched a definition of his idea that one theory may be closer to the truth than another which became 10 years later the start of a research program under the heading of *verisimilitude*, *truthlikeness* or *truth approximation*. Crucial for this research program was the discovery by David Miller and Pavel Tichý that Popper's definition could not be correct because it left no room for *false* theories being closer to the truth than other ones. A number of approaches were developed that attempted to solve or to circumvent this problem (an early overview is found in Kuipers 1987). Although these approaches were in agreement or at least compatible in several main respects, they were deviating from each other in some other essential respects, concerning questions of logical reconstruction (qualitative and quantitative, syntactic and semantic, disjunction- vs. conjunction-based, content- vs. likeness-based) or concerning adequacy conditions for verisimilitude. In the course of the years a large number of articles and a few monographs appeared, and some useful overviews were published that represented the state of the art at that moment (cf. Niiniluoto 1998, Oddie 2008).

T. Kuipers
Faculty of Philosophy, University of Groningen, Groningen, The Netherlands

G. Schurz (✉)
Department of Philosophy, University of Duesseldorf, Universitaetsstrasse 1,
Geb. 23.21, 40225 Duesseldorf, Germany
e-mail: schurz@phil.uni-duesseldorf.de

Whereas the primary problem of the research on verisimilitude was the *logical* problem of finding an optimal definition, several authors also engaged in the *epistemic* problems of verisimilitude. One of these epistemic problems consists in the question how the verisimilitude of theories changes when theories are expanded or revised in the face of growing evidence. It is this epistemic question where the connection of verisimilitude with belief revision comes into play.

2 Belief Revision

In the course of the eighties Carlos Alchourrón, Peter Gärdenfors, and David Makinson (1985) developed a research program, now known as the Belief Revision (BR) program, and more specifically, the AGM program. In spite of their common interest in scientific theory change, the philosophical motivations of the BR program and the TA (truth approximation) program are quite different. While the epistemological orientation of the TA program is truth as a relation of theories to an *external* world, the orientation of the BR program consisted mainly in *internal* investigations of change operations among sets of sentences or propositions without considerations of language-world relations. Instead of truth, the aims of belief revision, at least in the AGM style, are consistency and informative content. These aims guide the three main operations that have been introduced to model the change of a belief system in the face of *new information* or evidential *input*, and that are called *expansion*, *contraction* and *revision*. Expansion is an operation which simply adds a new piece of information to the old belief corpus, via the logical operations of conjunction or its set-theoretic counterpart, in the case when both are compatible with each other. Contraction and revision are more complicated and apply when the new information contradicts old beliefs: in that case, a revision consists of a contraction step and an expansion step (this composition of a revision is called the ‘Levi-identity’). The contraction step withdraws old beliefs are incompatible with the new information, thereby producing a contracted belief set that is no longer incompatible with the new information, which is then added to the contracted belief set.

While the correct definition of the operation of expansion is unproblematic, it has turned out to be very difficult to get consensus on the right definitions of the operations of contraction and revision. According to Hansson (2008) one of the most debated topics in belief revision theory is the recovery postulate, according to which all original beliefs are regained if one of them is first removed and then reinserted. Another much discussed topic is how iterated belief revision can be adequately represented. However this may be, there is a productive family of BR systems, each member having its own pros and cons. Many of the systems are purely syntactic, but in the course of time there has come a growing interest in taking semantic considerations into account, starting with the purely semantic approach of Adam Grove (1988).

A more substantial deviation from the AGM-model of belief revision is belief base (BB) revision introduced by Hansson (1999). In AGM-revision the belief state is represented as a *deductively closed* set of sentences or beliefs. Revision is performed over *all* beliefs of this set, independently from the question whether they are basic or derived. In this respect the AGM-model is *coherentistic*. In BB-revision, revision is

performed only over sets of basic beliefs, so-called *belief bases*. Formally, these are sets of beliefs which are no longer deductively closed; informally, they contain only beliefs that have some independent justification. So BB-revision is more in line with *foundationalist* epistemology. Although AGM-revision and BB-revision have many logical operations in common, their epistemological difference is rather deep (for detailed comparisons cf. Rott 2001, ch. 3.4.1).

Concerning the relation of BR to TA, one may argue that at least the BR-assumption that the input information is always accepted—the so-called success postulate—signifies an implicit orientation towards truth, insofar as this assumption means that the input information is assumed to be true. But neither is this assumption shared by all BR-accounts, nor is it substantial for the BR-account in its present development. This latter account studies belief revision in abstraction from the notion of truth. A telling quote can be drawn from Gärdenfors' gripping introduction (1988, p. 20):

[...] the concepts of truth and falsity are irrelevant for the analysis of belief systems. These concepts deal with the relation between belief systems and the external world, which I claim is not essential for an analysis of epistemic dynamics. [...] My negligence of truth may strike traditional epistemologists as heretical. However, one of my aims is to show that many epistemological problems can be attacked without using the notions of truth and falsity.

3 Belief Revision Aiming at Truth Approximation

Also Sven Ove Hansson concludes in his survey article on belief revision (2004, pp. 275–276) that “the relations between states of belief and the objects that these states refer to remain an essentially unexplored issue”. However, as pointed out by Niiniluoto (this volume), a BR theorist need not deny that beliefs have truth values and that the exploration of the relation between BR and TA would be an important issue. Indeed, the BR-theorist Hansson (ibid.) seems to agree with this standpoint. Likewise, a TA theorist who is interested in the epistemic question of how the verisimilitude of theories changes in the light of new evidence need not ignore that there exist elaborated models of belief revision in abstraction from the question of truth. This consideration brings us to the central topic of this volume: to the question how belief revision and truth approximation are related.

The most gripping question is the following one: can it be shown that the expansion or revision of a theory with a new piece of true information always increases its verisimilitude? Of course, the question is only interesting if the new information is assumed to be true, because that a theory's verisimilitude may decrease if we revise it with false information is trivial. Moreover, the really interesting case of progress in truthlikeness is given when the theory to be revised is false, i.e., has at least some false consequences. For this most interesting case, Niiniluoto (1999) was the first who proved two remarkable negative results. He showed that the expansion as well as the revision of a false theory *T* by a new true evidential input *e* may decrease *T*'s verisimilitude, for a very broad range of

truthlikeness accounts which have only to share a few adequacy conditions. Let us illustrate Niiniluoto's discovery at hand of a simple example in a propositional language. Assume there are two propositional variables p_1 and p_2 ; the full truth is $p_1 \wedge p_2$, and the false theory T asserts $\neg p_2$, so $T = \neg p_2$ (or $T = \text{Cn}(\{\neg p_2\})$) if T is assumed to be deductively closed; "Cn" for "consequences"). The new evidence asserts $p_1 \leftrightarrow p_2$. Then the expansion of T by e , which is written as $T + e$, is defined as $T \wedge e$, which is logically equivalent with $\neg p_1 \wedge \neg p_2$; hence $T + e = \neg p_1 \wedge \neg p_2$. The majority of truthlikeness accounts (though not all) agree that $T + e = \neg p_1 \wedge \neg p_2$ is less truthlike than $T = \neg p_2$, because $T + e$ asserts one basic falsity ($\neg p_1$) more than T without entailing more basic truths than T . This shows that the expansion of a false theory by true evidence may decrease T 's verisimilitude. A similar counterexample applies to the case of revision.

It cannot be argued that this negative result is a consequence of a wrong definition of belief revision, because the result obtains also in the case of expansion whose definition is unproblematic. Also, the negative result can hardly be attributed to a wrong definition of verisimilitude, because the only presupposition of this result is the adequacy condition that if p_1 and p_2 are true, then $\neg p_1 \wedge \neg p_2$ is less truthlike than $\neg p_2$ —a condition which is shared by a majority of—though not by all—accounts of verisimilitude. So a deeper investigation of this negative result is required. The major (though not the only) research question of the following papers is this: under which additional conditions, either on the belief revision procedure, on the logical form of theories, or on the verisimilitude measure, can it be shown that expansion or revision by true inputs increases the verisimilitude of theories?

4 An Overview of Papers in this Volume

We arranged the papers in order to facilitate readability in view of relations of content: papers dealing with more basic topics have been put first; papers which built upon topics of other papers have been put behind the latter ones.

The volume starts with a paper by Ilkka Niiniluoto, whose preferred framework is the representation of theories as disjunctions (or disjunctively interpreted sets) of *constituents*. Constituents are maximally complete descriptions of possible worlds in the given language. For a propositional language with n variables $p_1, \dots, p_n$ the 2^n constituents c_i have the form $(\pm)p_1 \wedge \dots \wedge (\pm)p_n$, with "±" for "unnegated" or "negated" (Niiniluoto's preferred linguistic framework, however, are the constituents of a monadic first order language). Fundamental to measures of verisimilitude within this account is a distance (or similarity) function between constituents. Most natural is the Hamming-distance $\Delta_{\text{Ham}}(c_1, c_2)$ between two constituents c_1 and c_2 , which is in the propositional case defined as the number of propositional variables on whose truth value c_1 and c_2 disagree (for normalization purposes this number is divided by 2^n). Let $\|T\|$ denote the set of constituents entailing (or verifying) theory T . The Hamming-distance function between constituents can be extended to a distance function $\Delta(T, c)$ between a theory T and a given constituent c by defining $\Delta(T, c)$ either as the minimum of the distances of constituents in $\|T\|$ to c , or as their average, or as a weighted average of the minimum distance and the sum of the

distances (the latter choice, known as *minsum*, is preferred by Niiniluoto). The quantitative verisimilitude of T to the true constituent τ , $V(T)$, is then simply defined as $1 - \Delta(T, \tau)$.

Also Grove's semantic modelling of BR is defined in terms of a distance (or similarity) relation between possible worlds as follows: if the evidence e is incompatible with T , then the revised proposition $\|T^*e\|$ is the set (or disjunction) of those possible worlds in $\|e\|$ with minimal distance to T . Niiniluoto starts his discussion of the relation between BR and TA with the positive message that while in standard BR theory the similarity relation between worlds is left undefined (which leads to a kind of "anything goes effect" as shown in section 6.2 of the Schurz's paper), the Hamming-distance account provides a natural definition and thus operationalizes BR theory. But Niiniluoto continues with the negative message that BR of false theories with true evidences, even if based on natural Hamming distance functions, does not always increase verisimilitude, neither in the expansion nor in the revision case. He illustrates this negative result with an example concerning the number of planets, which makes clear that the negative result is independent from the special choice of the extended distance function $\Delta(T, e)$, be it the min-, average- or min-sum-distance. Niiniluoto goes so far to conclude that belief revision in the AGM-tradition (including BB-revision) is generally *not effective* for truth approximation.

Although Niiniluoto himself mentions two cases to which his negative result doesn't apply—and more results of this sort are provided in the other papers—he continues in section 6 by presenting an alternative to theory revision, namely theory choice based on maximizing *estimated* verisimilitude, $\text{ver}(T|e) = \sum_{1 \leq i \leq n} P(c_i|e) \cdot V(T, c_i)$. Here new evidence doesn't revise our theories, but merely changes our subjective *probabilities* $P(c_i|e)$ of the constituents c_i given evidence e (either by Bayesian conditionalization or by Lewisian imaging) and hence the estimated verisimilitudes of given theories. In this alternative account theories are not revised in the light of new evidence; one rather works with a fixed partition of competing theories $PT = \{T_j; 1 \leq j \leq m\}$ and selects that theory in PT with maximal expected verisimilitude. The initial question how the revision of theories affects their verisimilitude can not directly be asked in this account, because in Niiniluoto's alternative account a change of the preferred theory in the light of new evidence is not produced by a process of theory revision but rather by a process of theory selection. The question whether Niiniluoto's alternative method of theory selection based on estimated verisimilitude always converges to the truth may still be asked and deserves further investigation.

In the next paper, Gustavo Cevolani, Vincenzo Crupi and Roberto Festa present a positive result on the relation between belief revision and verisimilitude that is based on the restriction of the compared theories to *conjunctive* theories, or in short *c-theories*. Generally speaking, *c-theories* are represented as conjunctions of unnegated or negated elementary statements which are mutually logically independent from each other. In the propositional case the elementary statements are the propositional variables $p_1, \dots, p_n$, and the theories consist of conjunctions (or conjunctively interpreted sets) of unnegated or negated propositional variables $(\pm)p_i$ which are also called *b(asic)-claims* or *literals*. Every *c-theory* T can be viewed as a

‘partial’ constituent that decides only the truth value of some variables and leaves that of the others undecided. Letting $t(T)$ stand for the set of true and $f(T)$ for the set of false b-claims of a c-theory T , the comparative relation of verisimilitude between c-theories T_1 and T_2 is then defined as $T_1 \geq_v T_2$ iff $t(T_2) \subseteq t(T_1)$ and $f(T_1) \subseteq f(T_2)$.

Cevolani et al. continue with their definition of belief revision of a c-theory T with a conjunctive input e (i.e., also the evidence e is a conjunction of b-claims). Their revision operation is a variant of BB (belief base) revision that is performed over b-claims as follows. The expansion of c-theory T by (T -compatible) c-evidence e is simply given as the conjunction of T and e (or the union of their b-claims). The revision $T * e$ for T -incompatible e is the conjunction of e with the contraction of T by $\neg e$, abbreviated as $T \div \neg e$. The contraction is in turn defined as a maximal subconjunction of T ’s b-claims that is consistent with e ; it can be proved that for c-theories and c-evidences this maximal subset is *uniquely* given by the removal of all b-claims of T which are incompatible with e . Based on these natural definitions of BR performed over c-statements Cevolani et al. are able to prove the following robust positive results concerning the relation of BR and TA: both the expansion and the revision of a c-theory by true c-evidence e increase T ’s verisimilitude (i.e., $T + e \geq_v T$ and $T * e \geq_v T$), where this increase is *strict* ($>$) if the evidence is partially new (i.e., e contains b-claims not contained in T).

Cevolani et al. extend their basic result in two ways. The first extension is about *completely false* input information (where a c-statement is completely false if all of its b-claims are false): it can be shown that revision with completely false inputs always decreases verisimilitude. The second extension concerns the quantitative extension V of the comparative verisimilitude concept $\geq_v$. Cevolani et al. demonstrate that under certain conditions on $V(T)$ and $V(e)$ the revision of a theory by evidence which is not true but more-or-less verisimilar will increase T ’s verisimilitude. One may object against Cevolani et al. that their positive results are restricted to conjunctive theories, which are a very special subcase of theories. Indeed, their results don’t hold for theories containing disjunctive or implicative combinations of b-claims. However, Cevolani et al. have a good reply to this objection: they demonstrate that the conjunctive approach can be extended from simple propositional examples to other areas where one is interested only in conjunctions of b-claims, and for which there exists a language with a suitably defined notion of a *constituent*, i.e., a maximally complete conjunctive theory. C-theories correspond to *subconjunctions* of such constituents—Oddie (1986) calls them ‘quasi-constituents’. One example are the constituents of a monadic 1st order language that assert which possible states are realized in the domain and which are not. Another example are nomic constituents in the sense of Kuipers (see below) that assert which states are nomologically possible and which aren’t.

In the third paper, Gerhard Schurz attempts to generalize the explained positive results to theories of arbitrary logical form, which in particular may also contain *implications* between b-claims (which is typical for scientific theories). He starts from a general distinction between *disjunction-of-possibilities* versus *conjunction-of-parts* accounts of verisimilitude. Next he introduces his favoured conjunction-of-part account to verisimilitude, namely the *relevant element* account. The set of

relevant elements of a theory T is written as T_r and consists of all minimal but still relevant ‘conjunctive parts’ of T . In propositional languages, these relevant elements are given as *clauses*, i.e., *disjunctions* of b-claims (or literals) satisfying the additional ‘relevance’ condition that no proper subdisjunction of them is entailed by T . It can be proved that the set T_r *preserves* the logical content of T . This is the difference with the view of conjunctive parts as b-claims, which preserve a theory’s content only if it is a c-theory in the sense of Cevolani et al. Based on earlier work, Schurz goes on to introduce his comparative definition of verisimilitude, according to which $T_1 \geq_v T_2$ iff $(T_1)_{tr} \parallel\!\!-\! (T_2)_{tr}$ and $(T_2)_{fr} \parallel\!\!-\! (T_1)_{fr}$, where $(T)_{tr}$ and $(T)_{fr}$ denote the set of T ’s true and T ’s false relevant elements, respectively (“ $\parallel\!\!-\!$ ” for “logical consequence”). He extends this comparative notion to a quantitative concept of verisimilitude defined over relevant elements that has been introduced in Schurz and Weingartner (2010).

Next Schurz turns to the relation between verisimilitude and theory *expansion*. Because his theories may have arbitrary form, Schurz is able to represent Niiniluoto’s counterexample to V-increasing expansions in full generality; he shows that counterexamples of this sort are even possible when the evidence is purely conjunctive, i.e., consists of conjunctions of true b-claims (though untypical for theories, this is typical for evidence). The reason for these counterexamples to expansions are what Schurz calls “true-to-false-implications” of the form $t \rightarrow (f_1 \wedge f_2)$ (where t is a true b-claim and f_i are false b-claims): if the true b-claim t is added to the false theory $t \rightarrow (f_1 \wedge f_2)$ the expanded theory entails the new false b-claims f_1, f_2 which may produce a V-decrease in spite of the new true b-claim t . However, Schurz provides conditions under which expansion by true conjunctive evidences is guaranteed to increase a theory’s verisimilitude. The simplest of these conditions requires that each b-claim of the evidence resolves with at most one false relevant element of the theory. The corresponding theorem covers the positive result for conjunctive theories as a special case.

Turning to the revision of theories T by T -incompatible evidence e , Schurz first argues that the AGM-account to theory-revision is too liberal insofar as it allows too many unreasonable revision operations. Instead of AGM-revision Schurz suggests to use belief base (BB) revision which is directly performed over the set T_r of relevant elements of a theory, based either on maxichoice contractions of T_r (i.e., maximal T_r -subsets not entailing $\neg e$) or on suitably defined “disjunctions” of them. In Schurz’s framework, the reason behind Niiniluoto’s negative result on revision is what Schurz calls the “Duhem-problem of revision”: if several (instead of just one) relevant elements of T entail $\neg e$, then the contraction of T_r by $\neg e$ may remove true instead of false relevant elements of T_r and, thus, decrease T ’s verisimilitude. In his final theorem Schurz proves that theory revision by true conjunctive evidence is guaranteed to increase verisimilitude if two conditions are satisfied: (a) the contraction-operation is “truth-preserving”, i.e., removes only false elements, and (b) the above-mentioned condition for expansion holds.

The fourth paper of Theo A.F. Kuipers is based on his preferred framework of *nomio* verisimilitude. Kuipers starts from a given partition of logically *possible states* $PS = \{S_i; i \in I\}$ of an individual or a system of individuals. He represents theories T as sets $T \subseteq PS$ of possible states. Although this is formally similar to the

disjunction-of-possibilities representation of theories, Kuipers' sets have a *conjunctive* interpretation, insofar as they assert that, among all possible states in PS, exactly the states in T (and no other states) are *nomologically* possible (cf. Niiniluoto 1987, 381, and the remarks in section 5 of the Cevolani et al. paper). With N = the *nomic truth* = the *true* set of nomological possibilities, Kuipers defines his *basic* and comparative notion of “having as least as much verisimilitude as”, $T_1 \geq_v T_2$, by the two conditions (a) $T_2 \cap N \subseteq T_1 \cap N$ and (b) $T_1 - N \subseteq T_2 - N$. Informally speaking, (a) T_2 's true possibility-claims are contained in those of T_1 , and (b) T_1 's false possibility-claims are contained in T_2 .

Based on a “success theorem” which is based on earlier work (Kuipers 2000) Kuipers then introduces a natural method of theory revision over possibilities that is, however, *prima facie* unrelated to BR in the AGM tradition. He assumes that new evidence consists in (1) experimental results R which provide new true possibilities i.e., $R \subseteq N$, and (2) inductive generalizations which provide new empirical laws that are assumed as true and thus constrain conjectures about nomological possibilities, captured by a set $S \supseteq N$. Kuipers defines the basic revision of a given theory T by an evidence-pair R/S as the transition from T to $(T \cap S) \cup R$ (or equivalently to $(T \cup R) \cap S$). He is able to prove that the so defined basic revision of T by R/S will never decrease the theory's verisimilitude, and in the long run it will increase it.

In the light of the general negative results explained above, Kuipers' positive result must be due to the logical form of his theories and his definition of revision. Kuipers asks how his method of revision relates to AGM-revision, and he finds out that his BR method is equivalent to a sequence of an AGM-expansion step and a so-called full meet AGM-contraction step. While Kuipers' basic revision method is adequate only in the case when T and S are compatible, he also introduces a refined revision method and represents it in terms of Grove's semantic BR approach. Based on his definition of *refined* verisimilitude, Kuipers finally shows that the refined method of theory revision never decreases refined verisimilitude. Kuipers adds, however, that his two positive results have the disadvantage that they apply only to the “weak necessity claim” of a theory T , which asserts that all true nomic possibilities are in T , as opposed to the “sufficiency claim” which asserts that all possibilities in T are true nomic possibilities.

The fifth paper of Gerard R. Renardel de Lavalette and Sjoerd D. Zwart has something in common with the paper of Niiniluoto: it works within the framework of disjunctions of possibilities. Theories are represented as disjunctions of constituents $c_i = (\pm)p_1 \wedge \dots \wedge (\pm)p_n$ of a propositional language (which Renardel and Zwart call “atoms” because they figure as atoms in the corresponding Boolean algebra). The authors assume two kinds of pre-orderings and corresponding ordinal similarity (or preference) functions $t/p: C \rightarrow \{0, 1, \dots\}$ which assign ordinals to constituents in C . Both similarity functions are extended to a similarity function between a constituent and a set of constituents via the explained ‘minimum distance’ (or ‘maximal similarity’) criterion. The first function $t(c_i)$ ranks constituents according to their distance to the truth, i.e., $t(c_i) = 0$ for constituents which are maximally truthlike. The second function p_T is used for BR; it ranks constituents according to their similarity with the given theory T ($p_T(c) = 0$ for

constituents which are maximally T-similar). Renardel and Zwart define belief revision in the line of Grove's semantical approach, i.e., $\|T^*e\|$ = the set of constituents in $\|e\|$ which have minimal rank according to p_T . They extend Grove's BR account by formulating adequacy criteria for the revision of the T-preference relation p_T into the revised preference relation $p_{(T^*e)}$ —this is needed for laying down how T^*e is to be revised in the light of further evidence.

Concerning the relation of BR and TA, the basic idea of Renardel and Zwart is to demonstrate that theory-revision increases verisimilitude if the truthlikeness similarity function t and the revision preference function p_T are “sufficiently similar” to each other. There are, however, three important differences of Renardel and Zwart's account with most of the other accounts. First, Renardel and Zwart allow for the (in their view degenerated) case that the full truth τ is incomplete, being a proper disjunction of constituents that leaves the truth value of some propositional variables undecided. They share this deviation from standard accounts with Kuipers' account of nomic truthlikeness. Second, in contrast to Niiniluoto, Renardel and Zwart don't assume that the distance between constituents is always the Hamming-distance—in fact, their final theorem presupposes that this distance may sometimes be non-Hamming. Third, Renardel and Zwart's comparative definition of their *refined* verisimilitude ordering—which is the product of a content- and a likeness-ordering—has some untypical properties which are not shared by the papers discussed so far. In particular, it follows from their “content condition” (3) that their (refined) verisimilitude increases with logical content even for false theories. So for example, if p_1 and p_2 are true, then the theory $\neg p_1 \wedge \neg p_2$ has strictly more refined verisimilitude in the sense of Renardel and Zwart than the theory $\neg p_1$, although intuitively $\neg p_1 \wedge \neg p_2$ makes one mistake more than $\neg p_1$.¹

An immediate consequence of this property of Renardel and Zwart's definition of refined verisimilitude is that the explained negative results about theory-expansion and verisimilitude do not hold. On the contrary, the authors demonstrate (below their equation (13)) that when the true evidence e is T-compatible (in their words, when the revision is “conservative”), it holds trivially in their account that expansion by e increases T 's verisimilitude. However, this is not so in the case of the revision of T with T-incompatible e . At hand of a nice example which goes back to Hansson the authors illustrate that there exist natural preference functions p_T which revise the theory T in the direction of decreasing refined verisimilitude. Let us give an example with three propositional variables, where $\tau = p_1 \wedge p_2 \wedge p_3$, $T = \neg p_1 \wedge \neg p_2 \wedge \neg p_3$ and $e = (p_1 \wedge p_2 \wedge p_3) \vee (\neg p_1 \wedge \neg p_2 \wedge \neg p_3)$. Using the Hamming-distance for BR ($*_1$) one obtains $T^*_{*1}e = \neg p_1 \wedge \neg p_2 \wedge \neg p_3$, which has less verisimilitude than T (also for Renardel and Zwart); only a non-Hamming BR-operation $*_2$ leads to $T^*_{*2}e = p_1 \wedge p_2 \wedge p_3$ which has more verisimilitude than T . Renardel and Zwart finally formulate a condition under which the revision of T by true evidence is guaranteed to increase T 's verisimilitude: when (a) the number of constituents of $e \wedge \neg \tau$ is less than or equal as the number of constituents of $T \wedge \neg \tau$, and (b) there

¹ Renardel and Zwart use inverted ordering symbols; i.e., they write “ $T_1 <^v T_2$ ” for “ T_1 has *more* verisimilitude than T_2 ”.

exists a t -increasing injection of the former to the latter. In particular, when t and p are identical, revising with true evidence always increases verisimilitude.

The two final papers differ from the five previous ones in three respects. First, they do not start from the problems of providing adequate definitions of verisimilitude or belief revision. They rather widen the perspective of this volume to two new areas: the role of doxastic meta-belief and social opinion dynamics for success in truth approximation. Second, the previous five papers have focused on *one-shot* results about the effect of revising a theory T by a piece of new evidence e . In contrast, the next two papers concentrate on the effect of *sequences of revisions* in the long run. Of course, long-run results may be derived from one-shot results—but while this remains a future task for the accounts of the first five papers,² it constitutes the focus of the final two papers.

The paper of Alexandru Baltag and Sonja Smets opens the perspective towards the area of belief revision with *higher-level doxastic information*. Their approach is based on the logical framework of so-called *belief revision friendly dynamic logic*. The authors start from so-called finite *plausibility models*, which represent the graded beliefs of a given agent and have the form $(S, \leq, \text{val}, s_0)$. Here S is a finite set of (epistemically possible) worlds (or states), $\leq$ is a total plausibility pre-ordering over S (equivalent with a Grove-type pre-ordering relation), s_0 is the actual world (state), and val a standard valuation function for a propositional language that is enriched with a knowledge operator K , an unconditional belief operator B and a conditional belief operator B^Q . The truth condition for KP (for “proposition P is known”) is that P is true in all S -worlds (and hence also in s_0); the truth condition for BP (for “ P is believed”) that P is true in all $\leq$ -lowest (i.e., most plausible) worlds, and that for $B^Q P$ (for “ P is believed under condition Q ”) that P is true in all $\leq$ -lowest Q -worlds. While ordinary ‘ontic’ propositions are assumed to be expressible in the object-language, Baltag and Smets also allow for higher-level doxastic propositions such as “If the agent believes P , then Q ” which are represented as S -subsets which need not be expressible in the object language.

Baltag and Smets then introduce three revision operations over plausibility models, so-called “upgrade operations”, which upon receiving an input proposition P transform a given model S into an upgraded model. The first operation is simply called *update* of S by input proposition P , written as $S!P$. Update corresponds to revision satisfying the AGM-axiom of success: i.e., P is assumed to be true ($s_0 \in P$) and after updating with P , P is known and hence true in all worlds of the updated plausibility model. Thus the updated model $S!P$ results from S by simply deleting all S -worlds in which P is false and restricting $\leq$ and val to the so contracted world set. Baltag and Smets then present two further revision operations which don’t obey the axiom of success and are called *radical upgrade* and *conservative upgrade*. Upgrades keep all worlds in S but change the plausibility ordering $\leq$. In the radical upgrade of S by P , denoted by $S \uparrow P$, P is believed with high certainty: this is described as a transformation which shifts all P -worlds below all $\neg P$ -worlds and keeps the ordering relations among P -worlds and among $\neg P$ -worlds as in S . Conservative upgrade, abbreviated as $S \uparrow P$, only shifts all most plausible P -worlds

² Niiniluoto and Kuipers briefly mention some long-run consequences.

below all most plausible $\neg P$ -worlds and keeps the other ordering relations as they were in S .

Next, Baltag and Smets turn to *upgrade streams* which are infinite sequences of upgrades of a given plausibility model S by a corresponding sequence of propositions (P_i) , resulting in an infinite sequence of upgraded models (S_i) . Such an upgrade stream is said to stabilize all (simple or conditional) beliefs of an initial model $S = S_0$ if after some discrete time n all transformed models S_m ($m > n$) verify the same (simple or conditional) beliefs. As a first result Baltag and Smets report that all update streams stabilize all beliefs of a given (initial) model $S = S_0$, and all upgrade streams stabilize all knowledge of an initial model S_0 . However, not all upgrade streams stabilize all beliefs, even not if the upgrade streams are *truthful*, i.e., consist only of true propositions. The authors illustrate this negative result by a nice example that involves higher-order doxastic propositions about the agent's beliefs. What one can positively prove is that every radical—but not necessarily every conservative—upgrade stream stabilizes all simple (but not necessarily all conditional) beliefs of a given model.

Baltag and Smets finally ask what these results imply for the question of truth approximation in the long run. They discover two results on this question. First, they show that for every truthful update or radical upgrade stream (P_i) there exists a time n after which no subsequent upgrades are informative, i.e., for every $m > n$ S_m entails a true answer to the question “is P_m true or false?”. Second, they show that if an upgrade stream (P_i) that is expressible in the given object language is truthful and complete, i.e., the union of all propositions P_i determines the complete truth τ , then every update and radical upgrade stream stabilizes the beliefs of a given plausibility model to the full truth τ . The results of Baltag and Smets depend on the assumption that the set S of possible worlds is finite; extensions to infinite world-sets are left for the future.

The final paper of Igor Douven and Christoph Kelp highlights connections between truth approximation (TA) and certain results from *social epistemology*, concerning the role of *belief exchange among peers* for the success in TA. Building upon a model of Hegselmann and Krause, Douven and Kelp assume in their first model that a number of agents $1, 2, \dots$ (who form a social collective) formulate theories or hypotheses. These hypotheses simply consist in estimations $x_i(t)$ of the value of a real-valued parameter, where $x_i(t)$ is the estimation of agent i at the discrete time or round t . The distance of the estimation $x_i(t)$ to the true parameter value τ is given by the linear difference $|x_i(t) - \tau|$. An agent i is called ε -close to an agent j at time t if their estimations at time t deviate from each other by not more than ε , i.e., $|x_i(t) - x_j(t)| \leq \varepsilon$.

Douven and Kelp first report some major results of the model of opinion dynamics by Hegselmann and Krause. In this model, the opinion of each agent i for the next time, $x_i(t+1)$, is given as a weighted average of a *social opinion component* $s_i(t+1)$ and a bias towards the truth τ : $x_{i+1}(t) = \alpha \cdot s_i(t+1) + (1 - \alpha) \cdot \tau$ (with $\alpha \in [0,1]$). The social opinion component depends on the previous opinions of all “trustworthy” members of the agent's collective: $s_i(t+1)$ is defined as the average of the opinions $x_j(t)$ of those agents j whose opinion was ε -close to agent i 's opinion at time t . The *truth-bias* “ $(1 - \alpha) \cdot \tau$ ” intends to capture the

assumption that each round each agent makes some new observations which bring him or her “a little bit” closer to the truth, independently of the influence of the other trustworthy agents on his or her opinion. Hegselmann and Krause have demonstrated that under very weak conditions on the truth-bias (e.g., when $(1 - \alpha)$ is just slightly greater than zero, or when only a few agents have a high α), the opinions of all members of the collective will converge to the truth τ .

Douven and Kelp ask under which conditions the social component of the opinion dynamics brings advantages compared to a purely individualistic truth-biased opinion dynamics (i.e., $\alpha = 0$). Based on computer simulations the authors find out that when the observations involve a random deviation from the true value τ , the social component increases the verisimilitude of the final opinions of the agents. On the other hand, the disadvantage of the social component is that it decreases the speed of convergence. Douven and Kelp then present an extension of their results in a direction which brings them close to the work in the previous papers: they assume that the conjectures of the agents are propositional theories, interpreted as disjunctions of propositional constituents, and the distance between a theory and the true constituent is given as the minimal Hamming-distance. The latter distance function is not only used to define verisimilitude but also to define the set of agents whose opinion is ε -close to the opinion of a given agent. The major results of their simulation of this more complicated scenario confirm the results of the previous simple scenario: the social component increases the verisimilitude of the final opinions if a random error term is assumed, and at the same time it decreases the speed of convergence.

References

- Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change. *Journal of Symbolic Logic*, 50, 510–530.
- Gärdenfors, P. (1988). *Knowledge in flux*. Cambridge/Mass.: MIT Press.
- Grove, A. (1988). Two modelling for theory change. *Journal of Philosophical Logic*, 17, 157–170.
- Hansson, S. O. (1999). *A textbook of belief dynamics*. Dordrecht: Kluwer.
- Hansson, S. O. (2004). Belief revision from an epistemological point of view. In I. Niiniluoto, M. Sintonen, & J. Wolenski (Eds.), *Handbook of epistemology* (pp. 255–279). Dordrecht: Kluwer.
- Hansson, S. O. (2008). Logic of belief revision. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, URL = <http://plato.stanford.edu/archives/win2008/entries/logic-belief-revision/>.
- Kuipers T. A. F. (Ed.) (1987) *What is closer-to-the-truth?* Amsterdam: Rodopi.
- Kuipers, T. A. F. (2000). *From instrumentalism to constructive realism*. Dordrecht: Kluwer.
- Niiniluoto, I. (1987). *Truthlikeness*. Dordrecht: Reidel.
- Niiniluoto, I. (1998). Verisimilitude: The third period. *The British Journal for the Philosophy of Science*, 49, 1–29.
- Niiniluoto, I. (1999). Belief revision and truthlikeness. In B. Hansson et al. (Eds.), *Internet festschrift for Peter Gärdenfors*. Retrieved February 20, 2009, from <http://www.lucs.lu.se/spinning>.
- Oddie, G. (1986). *Likeness to truth*. Dordrecht: Reidel.
- Oddie, G. (2008). Truthlikeness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, E. N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/truthlikeness/>.
- Popper, K. (1963). *Conjectures and refutations*. London: Routledge.

- Rott, H. (2001). *Change, choice and inference: A study in belief revision and nonmonotonic reasoning*. Oxford: Clarendon Press.
- Schurz, G., & Weingartner, P. (2010). Zwart and Franssen's impossibility theorem holds for possible-world-accounts but not for consequence-accounts to verisimilitude. *Synthese*, 172, 415–436.