

# Logic in general philosophy of science: old things and new things

Hannes Leitgeb

Received: 23 March 2009 / Accepted: 7 September 2009 / Published online: 8 October 2010  
© Springer Science+Business Media B.V. 2010

**Abstract** This is a personal, incomplete, and very informal take on the role of logic in general philosophy of science, which is aimed at a broader audience. We defend and advertise the application of logical methods in philosophy of science, starting with the beginnings in the Vienna Circle and ending with some more recent logical developments.

**Keywords** Logic · Philosophy of science · Vienna circle

## 1 Old things: the Vienna circle and other success stories

... the logical clarification of scientific concepts, sentences, and methods releases us from inhibitory prejudices. Logical and epistemological analysis does not want to impose constraints on scientific research; on the contrary: it provides science with a scope of formal possibilities that is as complete as possible and from which the possibility that corresponds to the experience in question is to be selected<sup>1</sup> (*Wissenschaftliche Weltauffassung. Der Wiener Kreis*, 1929)

These days, a quote such as the above one typically prompts the following reaction: *nice but naive*—philosophy of science cannot be faithful to science if it is done in a manner that is intended to be abstract, idealizing, simplifying, and ahistorical. Haven't we known this at least since the late 1960s or so when logical empiricism had finally

---

<sup>1</sup> I will be using my own translations from the original German text in [Verein Ernst Mach \(1929\)](#) throughout this article.

---

H. Leitgeb (✉)  
Departments of Philosophy and Mathematics, University of Bristol, 9 Woodland Road,  
Bristol BS8 1TB, UK  
e-mail: Hannes.Leitgeb@bristol.ac.uk

died out? Isn't it clear by now that logic is just as helpful to philosophers of science as it is to historians or sociologists, i.e., at no great assistance at all?

Hold your horses! First of all, philosophy of science in the modern sense would not even exist without our heroes from the Vienna and Berlin Circles, their descendants and relatives. Secondly, the logically minded analysis of *logical*, *mathematical*, *computational*, and *semantic* concepts and theories was an overwhelming success: think of results such as the completeness of first-order logic or the incompleteness of any sound axiomatic system for second-order logic; consider the incompleteness theorems for Peano arithmetic and our modern proof theoretic investigations of arithmetical systems; bear in mind the systematic investigation of relative consistency and independence statements in axiomatic set theory; take our modern theory of computation—whether spelled out in terms of Turing machines or the lambda calculus or in some other way—and its unsolvability results; or look at the definability and undefinability results for truth. (You say that these were not the Vienna Circle's achievements? Sure, but that's not the point: the point is that these achievements are exactly what the Vienna Circle had in mind: the logical analysis of scientific concepts, sentences, and methods.) Finally: not even the logico-philosophical analysis of *empirical* theories was necessarily too abstract, idealizing, simplifying, and ahistorical *at all times*. Example: Reichenbach's logical reconstruction(s) of the theory of relativity; it would be quite desperate to argue that books and articles such as Reichenbach's

- *Relativitätstheorie und Erkenntnis apriori* from 1920  
(Reichenbach 1965)
- *Der gegenwärtige Stand der Relativitätsdiskussion* from 1922  
(Reichenbach 1959)
- *Axiomatik der relativistischen Raum-Zeit-Lehre* from 1924  
(Reichenbach 1969)
- *Philosophie der Raum-Zeit-Lehre* from 1928  
(Reichenbach 1957)

were not indeed philosophically important (e.g., Michael Friedman's recent work on the relativised *a priori* derives from Reichenbach's), informed by actual science (Reichenbach attended Einstein's lectures on the topic), and timely (we are speaking of the 1920s here). Even in more controversial cases, such as, say, the analysis of scientific explanation in terms of logically valid arguments according to the classic DN, DS, and IS models, the application of logical methods fares better than usually acknowledged: amongst others, the analysis maintains that explaining is closely related to inferring, explanations are tied to laws and hence, in some way, to causal affairs, and explanations come with pragmatic components (since something is an explanation only at a time *t*). While these inferential, causal, and pragmatic aspects of the models are not sophisticated enough to do justice to the complexity of actual scientific, or even commonsense, explanations—e.g., the pragmatic context of an explanation exhibits *much* more structure than just a temporal one—this is not so bad for a start

after all, is it? No wonder every decent course on scientific explanation still starts with these models, and it would be wrong to think that this were so merely because these models were historically influential: it is also because they are extremely clear, they put all their cards on the table, and they do a very nice job as paradigmatic schemata which may help organising our thoughts even when we are convinced that they will ultimately prove to be inadequate.

So here is the counterclaim: just as in any scientific endeavour, abstractness, simplicity and idealisation are often exactly what we want in philosophy, and there is nothing wrong about this as long as we keep in mind that, for the sake of the general view, we forgot about certain properties and we simplified or idealised others.<sup>2</sup> And, yes, to state the obvious: historical considerations by themselves are not yet philosophical ones (even though it never can hurt to be historically informed).

At the end of the day, logical model building is still likely to be one of the most successful methods in general philosophy of science. Indeed, even some of the “old” logical insights into science from the 1920s and 1930s are just as relevant to modern philosophy of science as they were back then. Let me give an example.

## 2 Old things and new things: a classic example of logical analysis

*Structural Realism* (Worrall 1989) holds that when one of our empirically successful theories of the past got replaced by one that was even more successful empirically, the mathematical structure—and more or less only the mathematical structure—of the original theory was preserved in the transition. Hence, we should commit ourselves only to this mathematical-structural content of our scientific theories.

This highly influential position in the recent debates on Realism, Antirealism, and their alternatives, is not without predecessors; e.g.:

... we get the result that principally every scientific statement can be transformed in the way that it becomes a structure statement. But this transformation is not only possible, it is also asked for. For science wants to talk of the objective; however, everything that does not belong to the structure but to the material, everything concrete, is ultimately subjective. In physics we observe this desubjectivization which already has transformed almost all physical concepts into structural concepts. (Carnap's *Der logische Aufbau der Welt* from 1928; cf. (Carnap 1967))

But how can we determine the structural content of a theory or a statement? Here is a suggestion (Maxwell 1970): By *Ramsification*.

- Assume we are given a theory  $\Phi$  of the form  $\Phi[T_1, \dots, T_m, O_1, \dots, O_n]$ , where  $T_1, \dots, T_m$  are theoretical terms and  $O_1, \dots, O_n$  are empirical terms;

<sup>2</sup> See Kuipers (2007) for an excellent review of, amongst others, the roles of the desideratum of simplicity and of the strategy of idealisation in the explication of concepts.

- then the structural content of  $\Phi$  can be expressed by the Ramsey sentence of  $\Phi$ , i.e.,

$$\exists X_1 \dots \exists X_m \Phi[X_1, \dots, X_m, O_1, \dots, O_n]$$

the idea being that the original terms  $T_1, \dots, T_m$  are not actually taken to refer to anything but rather they function as structural “nodes” in a conceptual network that holds together the empirical predictions of the original theory.

Such transformations by Ramsification are backed up by the well-known logical fact that the Ramsey sentence of  $\Phi$  has the same empirical consequences as  $\Phi$  itself.

While this move turns the “structural content” of a theory into a well-defined notion, it also leads to a new problem—a dilemma: If the range of second-order quantifiers in the Ramsey sentence above is unrestricted, then the structural content of a theory amounts to not much more than the empirical content of the theory plus a cardinality constraint on the underlying class of individuals; cf. (Newman 1928; Demopoulos and Friedman 1985; Ketland 2004). On the other hand, if the range of the second-order quantifiers is restricted in some way (as, say, to *natural* properties and relations), then Structural Realism seems to be not much more than just a brand of traditional Realism.

Once again, this is not exactly a new problem, as can be seen from the publication date of Newman’s paper cited above: e.g., take Carnap’s *Aufbau* again;

- at some point in his construction system, Carnap wants to define the previously primitive relation sign *Er* (for *resemblance recollection*);
- he suggests to do so in a way that amounts to

$$Er =_{df} \iota R(\Phi[R] \wedge R \text{ is founded in experience})$$

where  $\Phi[R]$  is the result of substituting the second-order variable  $R$  for  $Er$  within the true high-level statement  $\Phi[Er]$  of Carnap’s construction system, while ‘founded in experience’ is nothing but the Carnapian version of ‘natural’ (which is thus made explicit by Carnap, rather than leaving it implicit in the understanding of the second-order quantifiers).

Carnap thereby anticipates Lewis’ “How to define theoretical terms” (Lewis 1970) by more than forty years. At the same time, he ends up facing a similar sort of trouble as the one encountered above: is ‘founded in experience’ still a structural term or do scientific sentences have some irreducibly non-structural content?

I am going to stop here, although this is obviously not the end of the story.<sup>3</sup> The two points that I wanted to illustrate with this are: firstly, the level of quality of the early discussions in philosophy of science was very high indeed, and there is still a

<sup>3</sup> For a recent paper on Structural Realism from the viewpoint of the Carnapian tradition, see (Schurz 2010). On a personal note: I am working on a monograph in which I revive Carnap’s *Aufbau* from a modern point of view and in which problems as the above play a major role. In that book, the story continues with Hilbertian–Carnapian epsilon terms, the mathematical structure of experience, and an “enlightened” structuralist version of phenomenalism; see (Leitgeb forthcoming) for a preview of what is to come. By the way, there is also an analogous debate on structuralism in philosophy of mathematics: see e.g. (Leitgeb and Ladyman 2008).

lot to be learned from them; secondly, the logical analysis of issues in general philosophy of science can lead to actual progress. You say: how can a dilemma such as the one described before constitute progress? We still don't know what to think about Structural Realism! True, but without even an attempt of logically analysing the structural content of theories in some way, the problems that are to be solved by the Structural Realist would remain concealed, and the overall discussion would be confined to metaphorical takes on non-metaphorical issues.<sup>4</sup>

### 3 New things: the new logic and new topics

The execution of such investigations shows all too soon that the traditional Aristotelian-scholastic logic is completely inadequate for this purpose. Only in the modern symbolic logic (“logistic”) it becomes possible to make conceptual definitions and propositions sufficiently sharp and to formalise the intuitive inferential processes of ordinary reasoning. . . . (*Wissenschaftliche Weltauffassung. Der Wiener Kreis, 1929*)

Much of the application of logical methods in general philosophy of science has been dominated by the methods of mathematical or “deductive” logic, whether in its traditional syntactic-axiomatic guise or in its later semantic-model-theoretic guise, and this has been so for a good reason: When philosophy of science came into being, what we now call standard mathematical logic was the state of the art, which is why it made perfect sense to first concentrate on exploiting the methods of mathematical logic in philosophy. In more recent times, the Sneed-Stegmüller-. . . school, which originated from Pat Suppes’ work, is perhaps the most mature and richest example of this approach to the logic of science, even though the level of sophistication of the model theoretic approach in philosophy of science never managed to keep pace with the standard of mathematical model theory.

In the meantime, logic has again progressed into other areas and directions. This is not to say that the syntactic or semantic deductive reconstruction of scientific concepts, theories, and methods has become obsolete: there is still a lot of room for progress here (see (van Benthem 1982) on “The Logical Study of Science” both for a retrospective survey and for an advertisement for new lines of investigation). But it is also clear that the original programme of analysing science in purely deductive terms has been running out of steam. Following the tenets of the Vienna Circle, we should expect new logical theories to take over and to build their new insights into science on top of the old ones. This might well happen.

Example 1: In the theory of belief revision, belief sets or scientific theories are considered—quite traditionally—as deductively closed sets  $K$  of formulas. When new evidence comes along in terms of a formula  $\alpha$ , then  $K * \alpha$  is supposed to be the “minimal” revision of  $K$  in light of the new evidence. Generally, the result of such a revision cannot be pinned down uniquely, but as Alchourrón, Gärdenfors, Makinson

<sup>4</sup> Many more stories of a similar kind could be told, e.g., the story of formal learning theory, which would start with Reichenbach, and which would include protagonists such as Putnam, Glymour, Kelly, Schulte, Hendricks. . . . Though logicians love completeness, I do not aim for completeness in this article.

(Alchourrón et al. 1985) have shown, it is possible to axiomatise the properties that *any* such belief revision operator  $*$  ought to have:

**K\*1**  $K * \alpha$  is a belief set

**K\*2**  $\alpha \in K * \alpha$

**K\*3**  $K * \alpha \subseteq K + \alpha$

**K\*4** If  $\neg\alpha \notin K$ , then  $K + \alpha \subseteq K * \alpha$


**K\*5**  $K * \alpha$  is inconsistent iff  $\models \neg\alpha$

**K\*6** If  $\models \alpha \leftrightarrow \beta$ , then  $K * \alpha = K * \beta$

**K\*7**  $K * (\alpha \wedge \beta) \subseteq (K * \alpha) + \beta$

**K\*8** If  $\neg\beta \notin K * \alpha$ , then  $(K * \alpha) + \beta \subseteq K * (\alpha \wedge \beta)$

(Don't worry, we will not discuss these axioms in any detail; they are just here to remind philosophers of science of how neat formal axiomatic systems look sometimes.) Later on, (Grove 1988) proved a representation theorem which showed that belief revision operators  $*$  which satisfy the AGM axioms stand in one-to-one correspondence with ranked "sphere" models of possible worlds of the form



in which  $K * \alpha$  is given by minimizing the ranks of worlds that satisfy  $\alpha$ . So, according to this theory, what a rational scientist does when he confronts his current theory by a new piece of evidence is to solve a minimisation problem; rational theory revision follows a preference ordering of the ways the world might be like.

The next step was to approach *iterated* belief revision as being triggered by a *sequence* or *stream* of evidence: It soon became clear that in order to attack this problem it was necessary to count the preferential order of worlds that is underlying Grove's sphere models as belonging to an agent's state of belief itself; mere deductively closed sets of formulas would no longer do as formalisations of belief states. Where to go from this insight is still a hot and unresolved question: see e.g. (Hild and Spohn 2008; Rott 2008; van Benthem 2007). It might well be that in order to understand the logic

of iterated belief revision properly it is actually necessary to turn belief revision into a logical theory proper, i.e., to go for *more* logic: first of all, formalise  $*$  above not as a metalinguistically expressed function on belief sets and formulas but as an object-linguistically expressed modal dynamic operator; secondly, rewrite the axioms of AGM as modal axioms of this operator; observe that it is only possible to state reduction axioms for this dynamic operator, as one usually has in dynamic epistemic logic for such operators, if the results of belief revision are expressed by a *conditional* belief operator rather than an unconditional one (which is nothing but the logical manifestation of the “insight” into the insufficiency of mere belief sets that was mentioned before); finally, add *eigenaxioms* which formalise the chosen method of iterated belief revision; see (van Benthem 2007) for details, where this line of reasoning is actually executed.<sup>5</sup>

Once translated into the terminology of philosophy of science, this means: One should expect scientific theories to come with a preferential ordering of theoretic fallback positions that govern the dynamics of theory change and which themselves can change in light of new evidence. Quite obviously, this ought to be good news for philosophers of science who crave for a way of dealing with theory change in a precise and rational manner. Immediately, lots of obvious and pressing questions present themselves; e.g.: presumably, only theory change in the normal phase of a scientific paradigm is subject to AGM-type rationality constraints, or can the framework be extended to paradigm changes as well? Or: is it possible to prove that if a method of iterated belief revision satisfies certain constraints, then iterated revision by evidence will approximate the truth in the limit? (Kevin T. Kelly has worked on this neglected topic; see (Kelly 1998).) *Vice versa*, moving from philosophy of science back to logic, it should be possible to refine the logical representation of belief states by importing theory components that had been isolated and analysed before by philosophers of science. (This is actually happening: see (Olsson and Westlund 2006), for a recent example). This also answers a question that was left open in the introductory section: Why is it good for a logically minded philosopher of science to be historically informed as well as to be informed of the scientific state of the art? Because it might pay off in terms of the structural depth, finesse, and applicability of our logical representations for science and its dynamics.

Example 2: Nonmonotonic reasoning was created by theoretical computer scientists as a formal means of dealing with non-deductive inferences of the sort

$$\alpha \sim \beta : \text{from } \alpha \text{ one may plausibly conclude } \beta$$

This is simply the classic topic of inductive logic presented in modern clothes. As Kraus et al. (1990) have shown, any such relation  $\sim$  may reasonably be taken to satisfy the following logical rules<sup>6</sup>:

$$- \frac{}{\varphi \sim \varphi} \text{ (Reflexivity)}$$

<sup>5</sup> More on the advantages of the modal-dynamic approach to belief revision over the traditional one can be found in (Leitgeb and Segerberg 2007).

<sup>6</sup> Cumulativity, i.e., Cautious Cut and Cautious Monotonicity taken together, go back to (Gabbay 1985).

- $\frac{\vdash \varphi \leftrightarrow \psi, \varphi \vdash \rho}{\psi \vdash \rho}$  (Left Equivalence)
- $\frac{\varphi \vdash \psi, \vdash \psi \rightarrow \rho}{\varphi \vdash \rho}$  (Right Weakening)
- $\frac{\varphi \vdash \psi, \varphi \wedge \psi \vdash \rho}{\varphi \vdash \rho}$  (Cautious Cut)
- $\frac{\varphi \vdash \psi, \varphi \vdash \rho}{\varphi \wedge \psi \vdash \rho}$  (Cautious Monotonicity)
- $\frac{\varphi \vdash \rho, \psi \vdash \rho}{\varphi \vee \psi \vdash \rho}$  (Disjunction)

(The main purpose of stating these rules here is again to impress philosophers of science. Philosophers of language might notice that this logical system is nothing but Ernest Adams' logic of indicative conditionals. Philosophical logicians should not need to notice this—they should know it.)<sup>7</sup>

There is a great variety of representation theorems for nonmonotonic consequence relations  $\vdash$ , amongst which are ones which are very similar to Grove's from above. In fact, consequence relations  $\vdash$  and revision operators  $*$  can be shown to be interdefinable. A more unusual representation theorem proceeds in terms of artificial neural networks: If patterns of activation in neural nets receive a "natural" interpretation in terms of formulas, one can prove that the resulting interpreted neural networks stand in a one-to-one correspondence to nonmonotonic consequence relations, such that in ' $\varphi \vdash \psi$ ' the formula  $\varphi$  represents the external input to the network, the formula  $\psi$  represents the stable state that the network reaches given that input, and  $\vdash$  represents the system trajectory that connects the one to the other; cf. (Leitgeb 2004, 2005, 2007a,b). So perhaps inductive logic can be viewed as a logical description of the dynamics of neural networks? In a related manner, (Churchland 1989) suggests to view scientific theories as being given by configurations of connection weights in neural networks, (Flach 1996) defines relations of scientific confirmation and explanation in terms of  $\vdash$ , and (Schurz 2001) takes the logic of  $\vdash$  to be the logic of normic laws, i.e., of the "soft laws" that are characteristic of all scientific disciplines which deal with living systems. So logical developments in computer science build bridges between inductive logic, neural networks, and scientific theorising, which is both unexpected and exciting.

These two simple examples should suffice to show that new logical methods, sometimes in combination with findings from cognitive science, have a lot to offer to general philosophy of science, and that in turn concepts and theses from philosophy of science can stimulate new logical discoveries. Indeed, there is so much yet to be done. To give a more contentious example: how about building up a logical methodology for the humanities? Impossible? Why not develop a logical account of hermeneutics (see Mantzavinos 2005)? Is it abductive logic? Why not think about the logical dos or don'ts of, say, the introduction of literary concepts (such as '*Gothic novel*')?

<sup>7</sup> It would be very easy to mention other logical success stories which are just as relevant to modern philosophy of science as the two examples I mention: just think of the progress in the logic of social science based on recent developments such as the logic of public announcement, logical accounts of game theory, social software, . . . No completeness claim, remember!



Shouldn't there be a formally precise theory of definition by means of prototypes (say, Mary Shelley's *Frankenstein* or Bram Stoker's *Dracula*) rather than by necessary and sufficient conditions? And so forth. We just have to do it.

#### 4 New things and new threats?

Some will—glad of being isolated—live a seclusive life on the icy firm of logic (*Wissenschaftliche Weltauffassung. Der Wiener Kreis, 1929*)

Is the application of new logical methods in philosophy of science threatened by new misfortune?

In particular: Will the mundane attractions of computer science deflect our brave logicians' attention from the more foundational concerns of philosophy of science? *No*. Take the case of nonmonotonic reasoning again: As it turned out, theoretical computer scientists eventually ran into similar questions and problems as philosophers of science and philosophers of language had done before them. It is true that some logicians are now getting paid by computer science departments or computer companies to work on something they would have published just a few years back in the *Journal of Philosophical Logic* or in *Philosophy of Science*, which may be bad for philosophy departments—and if so, it is *bad*—but this won't diminish the future relevance of logical methods for philosophy of science. It doesn't really matter if these methods are presented as the theoretical underpinnings of some fancy AI application or as a new take on the problem of induction, does it? As Clark Glymour (1992), p.367 pointed out, “Carnap wrote the first artificial intelligence program” in his *Aufbau*, and indeed in many ways Artificial Intelligence is a natural continuation of the Vienna Circle's old research programme; see (van Benthem 1989) for further support of this claim. If anything, the hands-on spirit of computer scientists, who are generally very good at building toy models, can only be helpful. So expect well-known problems in the logic of science to be tackled by research groups with funny acronym names—it's computer science after all.

Another worry, this time internal to philosophy: Is probability theory about to take over the role that logic had occupied previously in general philosophy of science? *No*. Probability theory will complement logic in peaceful and mutually beneficial coexistence. Two examples:

Example 1: Let  $\mathfrak{P}_w$  be the primitive conditional (Popperian) probability measure that belongs to world  $w$ ; define

$$\varphi > \psi \text{ is true in } w \text{ iff } \mathfrak{P}_w(\psi|\varphi) = 1$$

Then the following theorem holds:

**Theorem 1** (Leitgeb 2007a,b) *The system V of conditional logic is sound and complete with respect to this Popper function semantics for conditionals.*

V is almost Lewis' standard system for counterfactuals, and the flat fragment of V coincides with the logic of  $\vdash_{\sim}$  above. Moral: Within probabilistic structures one can

find the typical qualitative structures that are studied by logicians; one cannot fully understand the former without understanding the latter.

Example 2: To the best of my knowledge, there is no established theory of probability for indexical information<sup>8</sup> (witness the clash of intuitions brought to light by Elga's Sleeping Beauty problem), and there is no established theory of probability for introspective probabilistic statements (take van Fraassen's Judy Benjamin problem). The open question is: how shall we update probabilistically on information such as 'I woke up today' or 'My (his, her) subjective conditional probability of  $\psi$  given  $\varphi$  is so and so'? In both areas, logic has been developed further than probability theory: not that dynamic epistemic logic can claim to have any definitive answers at this point—how shall one's beliefs be updated in light of a Moore-type sentence?—but at least such questions are in the centre of discussion (see e.g. van Ditmarsch et al. 2007), and indexicality and nesting of modalities are the modal logician's daily bread. We should generously offer to help when our friends from probability theory will eventually come asking about methods of model building for statements with iterated modalities that involve epistemic modalities for more than one agent or which include both epistemic and ontic modalities: they might be in need to find ways of building models for statements with iterated probabilities that invoke epistemic probability measures of more than one agent or which mix epistemic and objective probability measures (as in Lewis' Principal Principle and related probabilistic reflection principles).

One last point: Who cares whether the original conception of the logic of science is being transformed into formal philosophy of science understood more broadly? Didn't Pat Suppes teach us quite some time ago that proper philosophy of science needed mathematical methods in general, rather than "only" logical ones? As long as formal philosophy booms—and it does!<sup>9</sup>—there should be more than enough room for everyone. Logic will afford others to enter.

The scientific conception of the world answers life and life takes it up.  
(*Wissenschaftliche Weltauffassung. Der Wiener Kreis, 1929*)

**Acknowledgements** I would like to thank Johan van Benthem, Vincent Hendricks, Leon Horsten, Theo Kuipers, and the organisers of the E.W. Beth Centenary Conference 2008 in Amsterdam for their invaluable comments and suggestions.

## References

Alchourrón, C. E., Gärdenfors, P., Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50, 510–530.

<sup>8</sup> Though people are working on it: see e.g. Mike Titelbaum's work at <http://sites.google.com/site/michaeltitelbaum/> and Darren Bradley's work at <http://faculty.arts.ubc.ca/dbradley/>.

<sup>9</sup> More particularly: It even does so in formal *philosophy of science*; cf. the ESF Exploratory Workshop on "Applied Logic in the Methodology of Science" in Bristol, September 2006; the special volume of *Studia Logica* on Formal Methods in the Philosophy of Science, which came out in 2008; (some of the subprojects of) the current ESF Research Networking Programme on "The Philosophy of Science in a European Perspective"; and the groups in Formal Epistemology at Tilburg, Leuven, Konstanz, Bristol, and other places.

- Carnap, R. (1967). *The logical structure of the world. Pseudoproblems in philosophy*. California: University of California Press.
- Churchland, P. M. (1989). *A neurocomputational perspective: The nature of mind and the structure of science*. Cambridge, MA: The MIT Press.
- Demopoulos, W., & Friedman, M. (1985). The concept of structure in Russell's the analysis of matter. *Philosophy of Science*, 52(4), 621–639.
- Flach, P. A. (1996). Rationality postulates for induction. In Shoham, Y. (Eds.), *Proceedings of theoretical aspects of rationality and knowledge, TARK'96* (pp. 267–281). San Mateo: Morgan Kaufmann.
- Gabbay, D. M. (1985). Theoretical foundations for non-monotonic reasoning in expert systems. In K. R. Apt (Ed.), *Logics and models of concurrent systems* (pp. 439–458). Berlin: Springer.
- Glymour, C. (1992). Android epistemology: Computation, artificial intelligence, and the philosophy of science. In M. H. Salmon (Ed.), *Introduction to the philosophy of science* (pp. 364–403). Indianapolis/Cambridge: Hackett.
- Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, 17, 157–170.
- Hild, M., & Spohn, W. (2008). The measurement of ranks and the laws of iterated contraction. *Artificial Intelligence*, 172(10), 1195–1218.
- Kelly, K. T. (1998). The learning power of iterated belief revision. In Gilboa, I. (Eds.), *Proceedings of the seventh TARK conference* (pp. 111–125). Illinois: Evanston.
- Ketland, J. (2004). Empirical adequacy and ramsification. *British Journal for the Philosophy of Science*, 55, 287–300.
- Kraus, S., Lehmann, D. J., & Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1–2), 167–207.
- Kuipers, T. A. F. (2007). Introduction. Explication in philosophy of science. In T. A. F. Kuipers (Ed.), *General philosophy of science—focal issues, handbook of the philosophy of science* (Vol. 1, pp. vii–xxiii). Amsterdam: Elsevier.
- Leitgeb, H. (2004). *Inference on the low level. An investigation into deduction, nonmonotonic reasoning, and the philosophy of cognition*. Dordrecht: Kluwer Applied Logic Series.
- Leitgeb, H. (2005). Interpreted dynamical systems and qualitative laws: From inhibition networks to evolutionary systems. *Synthese*, 146, 189–202.
- Leitgeb, H. (2007). Neural network models of conditionals: An introduction. In Arrazola, X. Larrazabal, J. M. et al. (Eds.), *LogKCA-07, proceedings of the first ILCLI international workshop on logic and philosophy of knowledge, communication and action* (pp. 191–223). Bilbao: University of the Basque Country Press.
- Leitgeb, H. (2007). *A probabilistic semantics for counterfactuals* (Unpublished draft).
- Leitgeb, H. (forthcoming). New life for Carnap's *Aufbau*? *Synthese*. doi:10.1007/s11229-009-9605-x.
- Leitgeb, H., & Ladyman, J. (2008). Criteria of identity and structuralist ontology. *Philosophia Mathematica*, 16, 388–396.
- Leitgeb, H., & Segerberg, K. (2007). Dynamic doxastic logic: Why, how, and where to? *Synthese*, 155(2), 167–190.
- Lewis, D. (1970). How to define theoretical terms. *The Journal of Philosophy*, 67(13), 427–446.
- Mantzavinos, C. (2005). *Naturalistic hermeneutics*. Cambridge: Cambridge University Press.
- Maxwell, G. (1970). Structural realism and the meaning of theoretical terms. In S. Winokur & M. Radner (Eds.), *Analyses of theories and methods of physics and psychology, Minnesota studies in the philosophy of science* (Vol. 4, pp. 181–192). Minneapolis: University of Minnesota Press.
- Newman, M. H. A. (1928). Mr. Russell's 'causal theory of perception'. *Mind*, 37(146), 137–148.
- Olsson, E., & Westlund, D. (2006). On the role of the research agenda in epistemic change. *Erkenntnis*, 65(2), 165–183.
- Reichenbach, H. (1957). *The philosophy of space and time*. New York: Dover.
- Reichenbach, H. (1959). The present state of the discussion on relativity. In H. Reichenbach (Ed.), *Modern philosophy of science: Selected essays by Hans Reichenbach* (pp. 1–45). London: Routledge & Kegan Paul.
- Reichenbach, H. (1965). *The theory of relativity and a priori knowledge*. California: University of California Press.
- Reichenbach, H. (1969). *Axiomatization of the theory of relativity*. California: University of California Press.
- Rott, H. (2008). Shifting priorities: Simple representations for twenty seven iterated theory change operators. In D. Makinson (Ed.), *Towards mathematical philosophy* (pp. 1–45). Dordrecht: Springer.

- Schurz, G. (2001). Normische Gesetzhypothesen und die wissenschaftsphilosophische Bedeutung des nichtmonotonen Schliessens. *Journal for General Philosophy of Science*, 32, 65–107.
- Schurz, G. (forthcoming). Structural correspondence between theories and the justification of scientific realism. *Synthese*.
- van Benthem, J. (1982). The logical study of science. *Synthese*, 51, 431–472.
- van Benthem, J., et al. (1989). Semantic parallels in natural language and computation. In H.-D. Ebbinghaus (Ed.), *Logic colloquium, granada 1987*. (pp. 331–375). Amsterdam: North Holland.
- van Benthem, J. (2007). Dynamic logic of belief revision. *Journal of Applied Non-Classical Logics*, 17(2), 129–155.
- van Ditmarsch, H. P., van der Hoek, W., & Kooi, B. P. (2007). *Dynamic epistemic logic* (Synthese Library Vol. 337). New York: Springer.
- Verein Ernst Mach (Eds.). (1929). *Wissenschaftliche Weltauffassung. Der Wiener Kreis*. Vienna: Artur Wolf Verlag.
- Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica*, 43, 99–124.