

Chapter 12

Introduction

Horacio Arló-Costa, Vincent F. Hendricks, and Johan van Benthem

It is well known that the usual versions of probability kinematics have serious limitations. According to the classical notion of conditioning when one learns a piece of information A its probability raises to its maximum (one). Moreover no further instance of learning will be capable of defeating A . Once a piece of information is learned one should be maximally confident about it and this confidence should remain unaltered forever. It is clear that there are many instances of learning that cannot be accommodated in this Procrustean bed. There are various ways of amending this limited picture by enriching the Bayesian machinery. For example, one can appeal to a notion of primitive conditional probability capable of making sense of conditioning on zero measure events. But the detailed consideration of this alternative leads to similar limitations: the picture of learning that thus arises continues to be cumulative. There are many ways of overcoming these important limitations. Williamson considers one possible way of doing so in his essay reprinted in the section on Bayesian epistemology. One of the lessons that have been learned in recent years is that there is no apparent way of circumventing this rigidity of Bayesianism without introducing in some way a qualitative doxastic or epistemic notion as a primitive alongside probability. Here are two examples:

Horacio Arló-Costa was deceased at the time of publication.

H. Arló-Costa (deceased)
Carnegie Mellon University, Pittsburgh, PA, USA

V.F. Hendricks (✉)
Center for Information and Bubble Studies, University of Copenhagen, Copenhagen, Denmark
e-mail: vincent@hum.ku.dk

J. van Benthem
University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
Stanford University, Stanford, United States
e-mail: johan@science.uva.nl

Williamson proposes a model where knowledge is a primitive, while Levi appeals to a primitive notion of *full belief*.

The traditional work in belief revision has followed Levi in adopting some doxastic notion as a primitive (which does not need to be full belief, it could be Spohn's *plain belief*, etc). But if such a qualitative notion should be introduced primitively how should one go about representing it? And how one would go about characterizing changes of the represented notion? There are many possible options from sentential representations (that can be finite or not) to propositional representations that can, in turn, incorporate more fine-grained non-probabilistic structure than mere beliefs (entrenchment or some notion of plausibility, for example).

The traditional approach derived from the seminal work of Carlos Alchourrón, Peter Gärdenfors and David Makinson (AGM) opted for a sentential representation. They were ambivalent between using finite representations (usually called belief bases) and using more structured representations of belief (theories in logical parlance or belief sets). Isaac Levi proposed a philosophical justification for the use of theories: they are supposed to represent not belief but commitments to belief.

So, how should one go about contracting a sentence A from a theory K ? Merely deleting A will not do, given that the sentence might be entailed by other sentences in K . The AGM solution to the problem is simple. Enter the *remainder set* of K with A (denoted $K \perp A$). This is the set of all maximal subsets of K that fail to entail A . One then makes a selection from the members of this set via the use of a selection function γ . This gives us $\gamma(K \perp A)$. Finally one takes the intersection of the resulting set. This is usually called a *partial meet* contraction. Of course there are many possible selection functions that can be used in the previous account. AGM would in turn require that the selection can be rationalized in the usual sense common in rational choice, i.e. by assuming that the selection function selects the *best* elements of $K \perp A$ (with respect to an underlying transitive preference relation). Then AGM manage to offer a set of rationality postulates that completely characterizes this operation. This is the main result presented in the article reprinted here. This article single handedly created an entire field of research by combining ideas from rational choice and classical proof techniques from philosophical logic.

If one reads the original AGM article carefully one sees that much of the inspiration from the use of selection functions in their model comes from the seminal work of Amartya Sen in rational choice. Exploiting results from the theory of choice Hans Rott (and previously Sten Lindström) systematically studies the relationship between functional constraints placed on selection functions and postulates of belief change. Among other things, Rott shows that certain functional constraints placed on propositional selection functions correspond in a one-to-one fashion to postulates of belief change. Rott's results forge a useful bridge between the mathematical theories of belief change and rational choice.

Still one might inquire why the feasible set from which one chooses rationally via a selection function should be limited to $K \perp A$. One common defense of this idea is in terms of minimizing information loss. The elements of $K \perp A$, usually

called maxi-choice sets, do satisfy this requirement. But then if they are optimal until this point of view why not to use a selection function that picks singletons from $K \perp A$? AGM showed that this type of contraction is badly behaved. So is the opposite idea of taking directly the intersection of the entire $K \perp A$. So, partial meet appears as an Aristotelian middle ground that happens to satisfy a set of intuitive postulates. Or so argued AGM. Nevertheless the subsequent discussion focused on some controversial AGM postulates like *recovery* (requiring that if one contracts K with A and then adds A to the result of this contraction one returns to K). There are many putative counterexamples to recovery and this generated the interest in defining notions of contraction that fail to satisfy recovery. Isaac Levi is a well-known defender of this line of thought and in his article he characterizes a notion of contractions that does fail to satisfy recovery. The central idea he proposes is that what is minimized in contraction is not information loss but loss of *informational value*. The notion of information value is some sort of epistemic utility obeying basic structural postulates like:

(Weak Monotony) If $X \subset Y$, then $V(X) \leq V(Y)$.

This is an intuitive principle that makes permissible that two sets carry equal informational value even when one the sets carries more information than the other. The additional information might not be valuable at all and therefore the level of informational value of the larger set might remain equal to the informational value of the smaller set. What other constraints one should impose on information value? In the article reprinted here Levi presents a very specific form of information value that he uses to characterize a particular notion of withdrawal (some rational notion of contraction where recovery fails) that he calls *mild contraction*. Rott and Pagnucco offered an alternative model of the same notion that they call *severe withdrawal*. It is clear that when epistemic utility satisfies the constraints proposed by Levi this particular form of contraction obtains. What seems to be missing is a pre-systematic explanation of why epistemic utility should satisfy these constraints or a justification of some controversial properties of severe withdrawal (like the postulate of *antitony*). It is a true thought that the introduction of epistemic utility in models of belief change opens up an insightful research strategy that at the moment remains relatively unexplored.

Sven Ove Hansson offers another account of contraction that fails to obey recovery. Nevertheless he arrives at this conclusion in a completely different way. In fact, Hansson is one of the most prominent defenders of finite models of belief in terms of belief bases (finite sets of sentences that are one of the possible axiomatic bases of a given theory). It is easy to characterize a version of partial meet contraction for bases by using the obvious variant of the definition used for theories. Then one can proceed as follows: an operation of contraction on a belief set K is generated by a partial meet base contraction if and only if there is a belief base B for K and an operator \sim of partial meet contraction for B such that the contraction of K with A yields the logical consequences of $(B \sim A)$ for all sentences A in the underlying language. Hansson shows that if an operation on a belief set is generated by some partial meet base contraction, then it satisfies the classical AGM postulates

for contraction except recovery. In addition the operation satisfies other postulates encoding a specific notion of *conservativity*.

The article by Spohn articulates an important epistemological idea, namely that one should focus on changes of entire *epistemic states* endowed with more structure than mere belief. This approach, in a more general setting, is also independently pursued by Adnan Darwiche and Judea Pearl in Darwiche and Pearl (1996). Spohn focuses on a particular type of epistemic state that now is usually called a *ranking function*. Roughly a ranking function is a function from the set of propositions (= sets of possible worlds) to the set of natural, real, or ordinal numbers, similar to a probability measure. Epistemologically one can see such functions as numerical (but non-probabilistic) representations of a notion of *plausibility*. In the presence of a new input the current ranking is mapped to a new ranking incorporating the incoming information (in revision). This is an ideal setting to study the structure of iterated changes of view and as a matter of fact both articles offer the best existing articulation of principles regulating iterated change. This is an important area of research in this field that still remains relatively open.

Suggested Further Reading

There are a number of recent surveys and books that complement the reprinted papers appearing here. Regarding surveys the two most recent surveys are: Logic of Belief Revision, in *Stanford Encyclopedia of Philosophy*, 2006, by Sven Ove Hansson; and: Belief Revision in *The Continuum Companion to Philosophical Logic*, (eds.) L. Hornsten and R. Pettigrew, by Horacio Arló-Costa and Paul Pedersen. These surveys contain references to previous surveys in the field. A classic book in this area that continues to be useful is Peter Gärdenfors's monograph: *Knowledge in Flux: Modeling the Dynamic of Epistemic States*, College Publications (June 2, 2008). A very useful textbook presentation of some of the main results in the theory of belief change is: *A Textbook of Belief Dynamics: Theory Change and Database Updating*, Springer 2010, by Sven Ove Hansson. The book focuses mainly on syntactic presentations of belief change and it contains a very detailed presentation of belief base updating. Some more recent topics like iterated belief change are not treated in detail though.

Decision theoretic foundations for belief change are provided in various books by Hans Rott and Isaac Levi (independently). A book-length argument articulating Rott's account (and extending the content of the article reprinted here) appears in: *Change, Choice and Inference: A Study of Belief Revision and Non-monotonic Reasoning*, Oxford Logic Guides, 2001. Some challenges to this type of foundational strategy are considered by Arló-Costa and Pedersen in: "Social Norms, Rational Choice and Belief Change," in *Belief Revision Meets Philosophy of Science*, (eds.) E.J. Olsson and S. Enqvist, Springer, 2011. Isaac Levi has also published various essays where he presents decision theoretic foundations for belief change (but his account is rather different than Rott's). The most recent book presenting Levi's current views about belief change is: *Mild Contraction: Evaluating Loss of Information Due to Loss of Belief*, Oxford, 2004. Further references to his work can be found in this book.

The previous accounts tried to justify principles of belief change in the broader context of Bayesian or neo-Bayesian theory. An almost orthogonal view consists in deriving principles of belief change by taking some form of formal learning theory as an epistemological primitive. While all the previous accounts focused on justifying the next step of inquiry (or a finite and proximate sequence of steps) this second strategy focuses on selecting belief change methods capable of learning the truth in the long run. One important paper in this tradition is Kevin Kelly's: [Iterated](#)

[Belief Revision, Reliability, and Inductive Amnesia](#), *Erkenntnis*, 50, 1998 pp. 11–58. Daniel Osherson and Eric Martin present a similarly motivated account that nevertheless is formally quite different from Kelly’s theory in: *Elements of Scientific Inquiry*, MIT, 1998.

There are various attempts to extend the theory of belief revision to the multi-agent case and to present a theory of belief change as some form of *dynamic epistemic logic*. The idea in this case is to use traditional formal tools in epistemic logic to represent the process of belief change. Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi have recently published a textbook with some basic results in this area: *Dynamic Epistemic Logic*, Springer, 2011. Krister Segerberg has developed his own brand of dynamic doxastic logic in a series of articles since at least the mid 1990’s. One recent paper including rather comprehensive results in this area is: “Some Completeness Theorems in the Dynamic Doxastic Logic of Iterated Belief Revision,” *Review of Symbolic Logic*, 3(2):228–246, 2010.

The notion of *relevance* is quite central for a representation of belief and belief change. In a Bayesian setting there are standard ways of articulating relevance. But there is recent work that has used proof theoretic techniques to deal with relevance rather than probability theory. Rohit Parikh initiated this area of research with an article published in 1999: Beliefs, belief revision, and splitting languages, *Logic, language, and computation* (Stanford, California) (Lawrence Moss, Jonathan Ginzburg, and Maarten de Rijke, editors), vol. 2, CSLI Publications, pp. 266–278. Recently David Makinson has contributed as well an important article in collaboration with George Kourousias,: [Parallel interpolation, splitting, and relevance in belief change](#), *Journal of Symbolic Logic* 72 September 2007 994–1002. This article contains a detailed bibliography of recent work in this area.

One recent paper including rather comprehensive results in this area is: “Some completeness theorems in the dynamic doxastic logic of iterated belief revision,” *Review of Symbolic Logic*, 3, 02, 2010. For more on iterated belief revision please refer to: Darwiche and Pearl (Darwiche, A., & Pearl, J. (1996). On the logic of iterated belief revision. *Artificial Intelligence*, 89, 1–29) appears in: *Change, choice and inference: A study of belief revision and non-monotonic reasoning*, Oxford Logic Guides, 2001.

And there is also more to be found in Pagnucco and Rott (Pagnucco, M., & Rott, H. (1999). Severe withdrawal – and recovery. *Journal of Philosophical Logic*, 28, 501–547. See publisher’s “Erratum” (2000), *Journal of Philosophical Logic*, 29, 121) and Lindström (Lindström, S. (1991). A semantic approach to nonmonotonic reasoning: Inference operations and choice. *Uppsala Prints and Preprints in Philosophy*, no. 6/1991, University of Uppsala).