

Lecture Notes in Artificial Intelligence

2087

Subseries of Lecture Notes in Computer Science

Edited by J. G. Carbonell and J. Siekmann

Lecture Notes in Computer Science

Edited by G. Goos, J. Hartmanis, and J. van Leeuwen

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Singapore

Tokyo

Gabriele Kern-Isberner

Conditionals in Nonmonotonic Reasoning and Belief Revision

Considering Conditionals as Agents



Springer

Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA
Jörg Siekmann, University of Saarland, Saarbrücken, Germany

Author

Gabriele Kern-Isberner
FernUniversität Hagen
Fachbereich Informatik, LG Praktische Informatik VI
Postfach 940, 58084 Hagen
E-mail: gabriele.kern-isberner@fernuni-hagen.de

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Kern-Isberner, Gabriele:
Conditionals in nonmonotonic reasoning and belief revision : considering
conditionals as agents / Gabriele Kern-Isberner. - Berlin ; Heidelberg ; New
York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ; Singapore ; Tokyo :
Springer, 2001
(Lecture notes in computer science ; 2087 : Lecture notes in artificial
intelligence)
ISBN 3-540-42367-2

CR Subject Classification (1998): I.2.3, F.4.1, I.2

ISBN 3-540-42367-2 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York
a member of BertelsmannSpringer Science+Business Media GmbH

<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2001
Printed in Germany

Typesetting: Camera-ready by author, data conversion by PTP Berlin, Stefan Sossna
Printed on acid-free paper SPIN 10839346 06/3142 5 4 3 2 1 0

Preface

Relationships amongst propositions are crucial pieces of knowledge. They express causal or plausible connections, bring isolated facts together, and help us obtain a coherent image of the world. Such relationships may be represented in a most general form by *if-then-conditionals*.

Conditionals are omnipresent, in everyday life as well as in scientific environments. We make use of conditional knowledge when we avoid puddles on sidewalks (being aware of “If you step into a puddle, then your feet might get wet”) and when we expect high wheat prices from observing cold and rainy weather in spring and summer (due to “If the growing weather is poor then there will be an increase in the price of wheat”). Conditionals represent generic knowledge, acquired inductively from experience or learned from books. They tie a flexible and highly interrelated network of connections along which reasoning is possible and which can be applied to different situations.

Therefore, conditionals are most important, but also quite problematic objects in knowledge representation. They are not simply “true” or “false”, like classical logical entities. In a particular situation, a conditional is applicable (you actually step into a puddle) or not (you simply walk around), it can be found confirmed (you step into a puddle and indeed, your feet get wet) or violated (you step into a puddle, but your feet remain dry because you are wearing rain boots). So the central problem in representing and modeling conditional knowledge is to handle adequately, on the one hand, inactive (or neutral, respectively) behavior, and, on the other hand, active as well as polarizing behavior.

This book presents a new approach to conditionals which captures this dynamic, non-propositional nature of conditionals peculiarly well. Conditionals are considered as agents shifting possible worlds in order to establish relationships and beliefs. This understanding of conditionals yields a rich methodological theory, which makes complex interactions between conditionals transparent and operational. Moreover, it provides a unifying and enhanced framework for knowledge representation, nonmonotonic reasoning, and belief revision, and even for knowledge discovery. In separating structural from numerical aspects, the basic techniques for conditionals introduced in this book

are applied both in a qualitative and in a numerical setting, elaborating fundamental lines of reasoning.

The novel theory of conditionals is at the heart of this work, from which its other major topics – revising epistemic states, probabilistic and nonmonotonic reasoning, and knowledge discovery – are developed. So central concerns of Artificial Intelligence research are dealt with in a uniform and homogeneous way by investigating structures of conditional knowledge. Such structures are substantial, for instance, in abductive as well as in predictive reasoning, or for simulation tasks.

Several persons contributed to the making of this book which is a revised version of my habilitation thesis at the FernUniversität Hagen, Department of Computer Science. In the first place, I would like to thank Christoph Beierle for accompanying this work with his criticism and his support, and for refereeing the thesis. I am also very grateful to the other referees, Gerhard Brewka and Dov Gabbay, and to Wilhelm Rödder who infected me with his enthusiasm for probabilistic conditionals and the principle of maximum entropy.

Thanks to Jeff Paris, Gerhard Brewka, and Karl Schlechta for discussing and sharing new ideas with me. Special thanks to Jeff for improving my English.

Parts of the results presented in this book were obtained while I was supported by a Lise-Meitner-scholarship, Department of Science and Research, North-Rhine-Westfalia, Germany.

I dedicate this book to my husband, Klaus, for encouraging me all the time, and to my children Silja, Maj-Britt, and Malte, for their creativity.

Table of Contents

1. Introduction	1
1.1 “Believe It or Not” – The Handling of Uncertain Knowledge	1
1.2 Overview	5
1.3 Basic Definitions and Notations	8
1.3.1 Propositional and Conditional Expressions	8
1.3.2 Probabilistic Logics	9
2. Belief Revision and Nonmonotonic Reasoning – State of the Art	11
2.1 Nonmonotonic Reasoning	12
2.2 Belief Revision	13
2.3 Nonmonotonic Reasoning and Belief Revision – Two Sides of the Same Coin?	17
2.4 Iterated Revision, Epistemic States, and Conditionals	17
2.5 Probabilistic Reasoning – The ME-Approach	23
3. Conditionals	27
3.1 Conditionals and Epistemic States	27
3.2 Conditional Valuation Functions	31
3.3 Conditional Valuation Functions and Beliefs	34
3.4 A Dynamic View on Conditionals	36
3.5 Conditional Structures	39
3.6 Conditional Indifference	48
4. Revising Epistemic States by Conditional Beliefs	53
4.1 Postulates for Revising by Conditional Beliefs	54
4.2 Representation Theorems	56
4.3 Conditional Valuation Functions and Revision	57
4.4 A Revision Operator for Ordinal Conditional Functions	60
4.5 The Principle of Conditional Preservation	62
4.6 C-Revisions and C-Representations	66

5. Characterizing the Principle of Minimum Cross-Entropy .	73
5.1 Conditional Preservation	75
5.2 The Functional Concept	80
5.3 Logical Coherence	85
5.4 Representation Invariance	87
5.5 Uniqueness and the Main Theorem	89
6. Reasoning at Optimum Entropy	91
6.1 Probabilistic Consequence and Probabilistic Inference	92
6.2 Basic Properties of the ME-Inference Operation	93
6.3 ME-Logic and Conditionals	95
6.4 ME-Deduction Rules	97
6.4.1 Chaining Rules	98
6.4.2 Cautious Monotonicity and Cautious Cut	99
6.4.3 Conjoining Literals in Antecedent and Consequent	100
6.4.4 Reasoning by Cases	101
7. Belief Revision and Nonmonotonic Reasoning – Revisited	103
7.1 Universal Inference Operations	104
7.2 Simultaneous and Successive Revision	108
7.3 Separating Background from Evidential Knowledge	109
7.4 Revising Epistemic States by Sets of Conditionals	112
7.5 Revision versus Updating	114
7.6 Focusing in a Probabilistic Framework	116
8. Knowledge Discovery by Following Conditional Structures	119
8.1 Probabilistic Knowledge Discovery	120
8.2 Discovering Conditional Structures	122
8.3 Examples – ME-Knowledge Discovery	131
8.4 Open Problems and Concluding Remarks	134
9. Algorithms and Implementations	137
9.1 ME-Reasoning with Probabilistic Conditionals	137
9.2 Probabilistic Knowledge Discovery	138
9.3 Possibilistic Belief Revision	139
10. Conclusion	141
A. Proofs	145
Bibliography	175
Index	187

Dedicated to my husband, Klaus,
and to my children, Silja, Maj-Britt, and Malte.

Bibliography

- [Acz61] J. Aczél. *Vorlesungen über Funktionalgleichungen und ihre Anwendungen*. Birkhaeuser Verlag, Basel, 1961.
- [Ada66] E.W. Adams. Probability and the logic of conditionals. In J. Hintikka and P. Suppes, editors, *Aspects of inductive logic*, pages 265–316. North-Holland, Amsterdam, 1966.
- [Ada75] E.W. Adams. *The Logic of Conditionals*. D. Reidel, Dordrecht, 1975.
- [AGGR98] R. Agrawal, J. Gehrke, D. Gunopulos, and P. Raghavan. Automatic subspace clustering of high dimensional data for data mining applications. In *Proceedings of the ACM SIGMOD Conference on Management of Data*, Seattle, Washington, 1998.
- [AGM85] C.E. Alchourrón, P. Gärdenfors, and P. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [AIS93] R. Agrawal, T. Imielinski, and A. Swami. Mining association rules between sets of items in large databases. In *Proceedings of the ACM SIGMOD Conference on Management of Data*, pages 207–216, Washington, DC., 1993.
- [AMS⁺96] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A.I. Verkamo. Fast discovery of association rules. In U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in knowledge discovery and data mining*, pages 307–328. MIT Press, Cambridge, Mass., 1996.
- [Ant97] G. Antoniou. *Nonmonotonic reasoning*. MIT Press, Cambridge, Massachusetts, 1997.
- [Bac90] F. Bacchus. *Representing and Reasoning with Probabilistic Knowledge: a Logical Approach to Probabilities*. MIT Press, Cambridge, Mass., 1990.
- [BG93] C. Boutilier and M. Goldszmidt. Revision by conditional beliefs. In *Proceedings 11th National Conference on Artificial Intelligence (AAAI'93)*, pages 649–654, Washington, DC., 1993.
- [Bol96] T. Bollinger. Assoziationsregeln – Analyse eines Data Mining Verfahrens. *Informatik-Spektrum*, 19:257–261, 1996.
- [Bou93] C. Boutilier. Revision sequences and nested conditionals. In *Proceedings International Joint Conference on Artificial Intelligence (IJCAI'93)*, pages 519–525, 1993.

- [Bou94] C. Boutilier. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence*, 68:33–85, 1994.
- [BP99] R.A. Bourne and S. Parsons. Maximum entropy and variable strength defaults. In *Proceedings Sixteenth International Joint Conference on Artificial Intelligence, IJCAI'99*, pages 50–55, 1999.
- [Bre89] G. Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *Proceedings Eleventh Joint Conference on Artificial Intelligence*, volume 2, pages 1043–1048, San Mateo, Ca., 1989.
- [Bre94] G. Brewka. Reasoning about priorities in default logic. In *Proceedings of the 12th National Conference on Artificial Intelligence (AAAI'94)*, pages 940–945. AAAI/MIT Press, 1994.
- [Bre96] G. Brewka, editor. *Principles of Knowledge Representation*. CSLI Publications, 1996.
- [BS84] B.G. Buchanan and E.H. Shortliffe. *Rule-based expert systems. The MYCIN experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, MA, 1984.
- [Cal91] P.G. Calabrese. Deduction and inference using conditional logic and probability. In I.R. Goodman, M.M. Gupta, H.T. Nguyen, and G.S. Rogers, editors, *Conditional Logic in Expert Systems*, pages 71–100. Elsevier, North Holland, 1991.
- [Cox46] R.T. Cox. Probability, frequency and reasonable expectation. *American Journal of Physics*, 14(1):1–13, 1946.
- [Csi75] I. Csiszár. I-divergence geometry of probability distributions and minimization problems. *Ann. Prob.*, 3:146–158, 1975.
- [DeF74] B. DeFinetti. *Theory of Probability*, volume 1,2. John Wiley and Sons, New York, 1974.
- [Dem67] A.P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Ann. Math. Stat.*, 38:325–339, 1967.
- [DGC94] D. Dubois, I.R. Goodman, and P.G. Calabrese. Special issue on the conditional event algebra. *IEEE Transactions on Systems, Man, and Cybernetics*, 24(12), 1994.
- [DLP94] D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In D.M. Gabbay, C.H. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3. Oxford University Press, 1994.
- [DP91a] D. Dubois and H. Prade. Conditional objects and non-monotonic reasoning. In *Proceedings 2nd Int. Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, pages 175–185. Morgan Kaufmann, 1991.
- [DP91b] D. Dubois and H. Prade. Conditioning, non-monotonic logic and non-standard uncertainty models. In I.R. Goodman, M.M. Gupta, H.T. Nguyen, and G.S. Rogers, editors, *Conditional Logic in Expert Systems*, pages 115–158. Elsevier, North Holland, 1991.

- [DP91c] D. Dubois and H. Prade. Epistemic entrenchment and possibilistic logic. *Artificial Intelligence*, 50:223–239, 1991.
- [DP92] D. Dubois and H. Prade. Belief change and possibility theory. In P. Gärdenfors, editor, *Belief revision*, pages 142–182. Cambridge University Press, 1992.
- [DP94] D. Dubois and H. Prade. A survey of belief revision and updating rules in various uncertainty models. *Intern. Journal of Intelligent Systems*, 9:61–100, 1994.
- [DP96] D. Dubois and H. Prade. Non-standard theories of uncertainty in plausible reasoning. In G. Brewka, editor, *Principles of Knowledge Representation*. CSLI Publications, 1996.
- [DP97a] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
- [DP97b] D. Dubois and H. Prade. Focusing vs. belief revision: A fundamental distinction when dealing with generic knowledge. In *Proceedings First International Joint Conference on Qualitative and Quantitative Practical Reasoning, ECSQARU-FAPR'97*, pages 96–107, Berlin Heidelberg New York, 1997. Springer.
- [DPT90] D. Dubois, H. Prade, and J.-M. Toucas. Inference with imprecise numerical quantifiers. In Z.W. Ras and M. Zemankova, editors, *Intelligent Systems - state of the art and future directions*, pages 52–72. Ellis Horwood Ltd., Chichester, England, 1990.
- [Dub86] D. Dubois. Belief structures, possibility theory and decomposable confidence measures on finite sets. *Computers and Artificial Intelligence*, 5:403–416, 1986.
- [FH94] N. Friedman and J.Y. Halpern. Conditional logics of belief change. In *Proceedings 12th National Conference on Artificial Intelligence, AAAI-94*, 1994.
- [FH96] N. Friedman and J.Y. Halpern. Plausibility measures and default reasoning. In *Proceedings 13th National Conference on Artificial Intelligence, AAAI-96*, volume 2, 1996.
- [FH99] N. Friedman and J.Y. Halpern. Belief revision: a critique. *Journal of Logic, Language and Information*, 8, 1999.
- [FPSS96] U. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. From data mining to knowledge discovery: An overview. In U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in knowledge discovery and data mining*. MIT Press, Cambridge, Mass., 1996.
- [FPSSU96] U. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy. *Advances in knowledge discovery and data mining*. MIT Press, Cambridge, Mass., 1996.
- [FR99] B. Fine and G. Rosenberger. *Algebraic Generalizations of Discrete Groups*. Dekker, New York, Basel, 1999.

- [Fre98] M. Freund. Preferential orders and plausibility measures. *J. Logic Computat.*, 8:147–158, 1998.
- [FS96] V.G. Fischer and M. Schramm. Tabl – a tool for efficient compilation of probabilistic constraints. Technical Report TUM-19636, Technische Universität München, 1996.
- [Fy⁺96] U. Fayyad, R. Uthurusamy, et al. Data mining and knowledge discovery in databases. *Communications of the ACM*, 39(11):24–64, 1996.
- [Gab85] D. Gabbay. Theoretical foundations for nonmonotonic reasoning in expert systems. In K. Apt, editor, *Logics and models of concurrent systems*. Springer, Berlin, 1985.
- [Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, Mass., 1988.
- [Gär92] P. Gärdenfors. Belief revision and nonmonotonic logic: Two sides of the same coin? In *Proceedings European Conference on Artificial Intelligence, ECAI'92*, pages 768–773. Pitman Publishing, 1992.
- [Gef92] H. Geffner. *Default Reasoning: Causal and Conditional Theories*. MIT Press, Cambridge, Mass., 1992.
- [Ger97] K. Gerhards. Verwendung probabilistischer Logik in einem medizinischen Expertensystem. Master's thesis, FernUniversität Hagen, 1997. (Diplomarbeit).
- [GGNR91] I.R. Goodman, M.M. Gupta, H.T. Nguyen, and G.S. Rogers, editors. *Conditional Logic in Expert Systems*. Elsevier, North Holland, 1991.
- [GH98] A.J. Grove and J.Y. Halpern. Updating sets of probabilities. In *Proceedings Fourteenth Conference on Uncertainty in AI*, pages 173–182, 1998.
- [GHK94] A.J. Grove, J.Y. Halpern, and D. Koller. Random worlds and maximum entropy. *J. of Artificial Intelligence Research*, 2:33–88, 1994.
- [GHR94] D.M. Gabbay, C.H. Hogger, and J.A. Robinson, editors. *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3. Oxford University Press, 1994.
- [GM94] P. Gärdenfors and D. Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65:197–245, 1994.
- [GMP90] M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. In *Proceedings AAAI-90*, pages 646–652, Boston, 1990.
- [GMP93] M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3):220–232, 1993.
- [Gol94] M. Goldszmidt. Research issues in qualitative and abstract probability. *AI Magazine*, pages 63–66, Winter 1994.
- [Goo63] I.J. Good. Maximum entropy for hypothesis formulation, especially for multidimensional contingency tables. *Ann. Math. Statist.*, 34:911–934, 1963.

- [GP92] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. In *Proceedings Third International Conference on Principles of Knowledge Representation and Reasoning*, pages 661–672, Cambridge, Mass., 1992.
- [GP96] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84:57–112, 1996.
- [GR94] P. Gärdenfors and H. Rott. Belief revision. In D.M. Gabbay, C.H. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, pages 35–132. Oxford University Press, 1994.
- [Gra91] G. Grahne. Updates and counterfactuals. In *Proceedings Second International Conference on Principles of Knowledge Representation and Reasoning, KR-91*, pages 269–276, 1991.
- [Gro88] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [Hal99a] J.Y. Halpern. A counterexample to theorems of cox and fine. *Journal of AI Research*, 10:76–85, 1999.
- [Hal99b] J.Y. Halpern. Cox’s theorem revisited. *Journal of AI Research*, 11:429–435, 1999.
- [Han89] S.O. Hansson. New operators for theory change. *Theoria*, 55:114–132, 1989.
- [Han91] S.O. Hansson. *Belief base dynamics*. PhD thesis, Uppsala University, 1991.
- [Har77] W.L. Harper. Rational conceptual change. *PSA 1976*, 2:462–494, 1977.
- [Hec96] D. Heckerman. Bayesian networks for knowledge discovery. In U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in knowledge discovery and data mining*. MIT Press, Cambridge, Mass., 1996.
- [Her91] H. Herre. Nonmonotonic reasoning and logic programs. In *Proceedings First International Workshop on Nonmonotonic and Inductive Logic, Karlsruhe, Germany, 1990*, pages 38–58. Springer, Lecture Notes in Artificial Intelligence 543, 1991.
- [HHJ92] P. Hajek, T. Havranek, and R. Jirousek. *Uncertain Information Processing in Expert Systems*. CRC-Press, Boca Raton, Florida, 1992.
- [Hoc99] S. Hoche. Possibilistic belief revision – foundations and implementation. Master’s thesis, FernUniversität Hagen, 1999. (Diplomarbeit).
- [Isr87] D.J. Israel. What’s wrong with nonmonotonic logic? In M.L. Ginsberg, editor, *Readings in Nonmonotonic Reasoning*, pages 53–57. Morgan Kaufman, 1987.
- [Jay83a] E.T. Jaynes. *Papers on Probability, Statistics and Statistical Physics*. D. Reidel Publishing Company, Dordrecht, Holland, 1983.
- [Jay83b] E.T. Jaynes. Where do we stand on maximum entropy? In *Papers on Probability, Statistics and Statistical Physics*, pages 210–314. D. Reidel Publishing Company, Dordrecht, Holland, 1983.

- [Jef83] R. Jeffrey. *The logic of decision*. University of Chicago Press, Chicago, IL, 1983.
- [Jen96] F.V. Jensen. *Introduction to Bayesian networks*. UCL Press, London, 1996.
- [JS83] R.W. Johnson and J.E. Shore. Comments on and correction to “Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy”. *IEEE Transactions on Information Theory*, IT-29(6):942–943, 1983.
- [KGK93] R. Kruse, J. Gebhardt, and F. Klawonn. *Fuzzy-Systeme*. Teubner, Stuttgart, 1993.
- [KI96a] G. Kern-Isberner. Characterizing the principle of minimum cross-entropy within a conditional logical framework. Informatik Fachbericht 206, FernUniversität Hagen, 1996.
- [KI96b] G. Kern-Isberner. Conditional logics and entropy. Informatik Fachbericht 203, FernUniversität Hagen, 1996.
- [KI97a] G. Kern-Isberner. A conditional-logical approach to minimum cross-entropy. In *Proceedings 14th Symposium on Theoretical Aspects of Computer Science STACS’97*, pages 237 – 248, Berlin Heidelberg New York, 1997. Springer.
- [KI97b] G. Kern-Isberner. A logically sound method for uncertain reasoning with quantified conditionals. In *Proceedings First International Conference on Qualitative and Quantitative Practical Reasoning, ECSQARU-FAPR’97*, pages 365 – 379, Berlin Heidelberg New York, 1997. Springer.
- [KI97c] G. Kern-Isberner. The principle of minimum cross-entropy and conditional logic. In *Proceedings of the Third Dutch/German Workshop on Nonmonotonic Reasoning Techniques and their Applications, DGNMR-97*, pages 73–82, Saarbrücken, Germany, 1997. MPI for Computer Science.
- [KI98a] G. Kern-Isberner. Characterizing the principle of minimum cross-entropy within a conditional-logical framework. *Artificial Intelligence*, 98:169–208, 1998.
- [KI98b] G. Kern-Isberner. Nonmonotonic reasoning in probabilistics. In *Proceedings European Conference on Artificial Intelligence, ECAI’98*, pages 580 – 584, West Sussex, UK, 1998. Wiley & Sons.
- [KI98c] G. Kern-Isberner. A note on conditional logics and entropy. *International Journal of Approximate Reasoning*, 19:231–246, 1998.
- [KI99a] G. Kern-Isberner. Following conditional structures of knowledge. In *KI-99: Advances in Artificial Intelligence, Proceedings of the 23rd Annual German Conference on Artificial Intelligence*, pages 125–136. Springer Lecture Notes in Artificial Intelligence LNAI 1701, 1999.
- [KI99b] G. Kern-Isberner. Postulates for conditional belief revision. In *Proceedings Sixteenth International Joint Conference on Artificial Intelligence, IJCAI-99*, pages 186–191. Morgan Kaufmann, 1999.

- [KI99c] G. Kern-Isberner. Revising by conditional beliefs. In *Proceedings Fourth Dutch-German Workshop on Nonmonotonic Reasoning Rechniques And Their Applications, DGNMR-99*, University of Amsterdam, 1999. Institute for Logic, Language and Computation.
- [KI01] G. Kern-Isberner. Revising and updating probabilistic beliefs. In M.-A. Williams and H. Rott, editors, *Frontiers in belief revision*, pages 329–344. Kluwer Academic Publishers, Dordrecht, 2001. (to appear).
- [KIR96] G. Kern-Isberner and H.P. Reidmacher. Interpreting a contingency table by rules. *International Journal of Intelligent Systems*, 11(6), 1996.
- [KLM90] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
- [KM91a] H. Katsuno and A. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.
- [KM91b] H. Katsuno and A.O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings Second International Conference on Principles of Knowledge Representation and Reasoning, KR'91*, pages 387–394, San Mateo, Ca., 1991. Morgan Kaufmann.
- [KS91] H. Katsuno and K. Satoh. A unified view of consequence relation, belief revision and conditional logic. In *Proceedings Twelfth International Joint Conference on Artificial Intelligence, IJCAI-91*, pages 406–412, 1991.
- [Kul68] S. Kullback. *Information Theory and Statistics*. Dover, New York, 1968.
- [Lau82] S.L. Lauritzen. *Lectures on Contingency Tables*. Aalborg University Press, Denmark, 1982.
- [Lem82] J.F. Lemmer. Efficient minimum information updating for bayesian inferencing in expert systems. In *Proceedings of the National Conference on Artificial Intelligence, AAAI-82*, 1982.
- [Lem83] J.F. Lemmer. Generalized bayesian updating of incompletely specified distributions. *Large Scale Systems*, 5:51–68, 1983.
- [Lev77] I. Levi. Direct inference. *The journal of philosophy*, 74:5–29, 1977.
- [Lev88] I. Levi. Iteration of conditionals and the Ramsey test. *Synthese*, 76:49–81, 1988.
- [Lew73] D. Lewis. *Counterfactuals*. Harvard University Press, Cambridge, Mass., 1973.
- [Lew76] D. Lewis. Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85:297–315, 1976.
- [LKI99] T. Lukasiewicz and G. Kern-Isberner. Probabilistic logic programming under maximum entropy. In *Proceedings ECSQARU-99*, volume 1638, pages 279–292. Springer Lecture Notes in Artificial Intelligence, 1999.
- [LM92] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.

- [LMS01] D. Lehmann, M. Magidor, and K. Schlechta. Distance semantics for belief revision. *Journal of Symbolic Logic*, 2001. (to appear).
- [LS77] R.C. Lyndon and P.E. Schupp. *Combinatorial group theory*. Springer, Berlin Heidelberg New York, 1977.
- [LS88] S.L. Lauritzen and D.J. Spiegelhalter. Local computations with probabilities in graphical structures and their applications to expert systems. *Journal of the Royal Statistical Society B*, 50(2):415–448, 1988.
- [Mak89] D. Makinson. General theory of cumulative inference. In M. Reinfrank et al., editors, *Non-monotonic Reasoning*, pages 1–18. Springer Lecture Notes on Artificial Intelligence 346, Berlin, 1989.
- [Mak94] D. Makinson. General patterns in nonmonotonic reasoning. In D.M. Gabbay, C.H. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 35–110. Oxford University Press, 1994.
- [Mal89] F.M. Malvestuto. Computing the maximum-entropy extension of given discrete probability distributions. *Computational statistics and Data analysis*, 8:299–311, 1989.
- [McC80] J. McCarthy. Circumscription - a form of nonmonotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.
- [MD80] D. McDermott and J. Doyle. Non-monotonic logic I. *Artificial Intelligence*, 13:41–72, 1980.
- [Mey98] C.H. Meyer. *Korrektes Schliessen bei unvollständiger Information*. Peter Lang Verlag, 1998.
- [MG91] D. Makinson and P. Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In *Proceedings Workshop The Logic of Theory Change, Konstanz, Germany, 1989*, pages 185–205, Berlin Heidelberg New York, 1991. Springer.
- [Moo88] R. Moore. Autoepistemic logic. In P. Smets, E.H. Mamdani, D. Dubois, and H. Prade, editors, *Non-standard logics for automated reasoning*, pages 105–136. Academic Press, London, UK, 1988.
- [MR92] C.H. Meyer and W. Rödder. Propagation in Inferenznetzen unter Berücksichtigung des Prinzips der minimalen relativen Entropie. In *Proceedings of the Annual Conference of the DGOR/ÖGOR 1992*, pages 446–453, Berlin Heidelberg New York, 1992. Springer.
- [MS98] N. Megiddo and R. Srikant. Discovering predictive association rules. In *Proceedings of the 4th International Conference on Knowledge Discovery in Databases and Data Mining*, 1998.
- [Mül98] W. Müller. Systemkonzept und Prototyp für probabilistisches Data Mining. Master’s thesis, FernUniversität Hagen, 1998. (Diplomarbeit).
- [Nea90] R.E. Neapolitan. *Probabilistic Reasoning in expert systems*. Wiley, New York, 1990.
- [Nil86] N.J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–87, 1986.

- [Nut80] D. Nute. *Topics in Conditional Logic*. D. Reidel Publishing Company, Dordrecht, Holland, 1980.
- [Par94] J.B. Paris. *The uncertain reasoner's companion – A mathematical perspective*. Cambridge University Press, 1994.
- [Pea86] J. Pearl. Fusion, propagation and structuring in belief networks. *Artificial Intelligence*, 29:241–288, 1986.
- [Pea88] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, Ca., 1988.
- [Pea89] J. Pearl. Probabilistic semantics for nonmonotonic reasoning: A survey. In G. Shafer and J. Pearl, editors, *Readings in uncertain reasoning*, pages 699–710. Morgan Kaufmann, San Mateo, CA., 1989.
- [Poo88] D. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.
- [PV90] J.B. Paris and A. Vencovská. A note on the inevitability of maximum entropy. *International Journal of Approximate Reasoning*, 14:183–223, 1990.
- [PV92] J.B. Paris and A. Vencovská. A method for updating that justifies minimum cross entropy. *International Journal of Approximate Reasoning*, 7:1–18, 1992.
- [PV97] J.B. Paris and A. Vencovská. In defence of the maximum entropy inference process. *International Journal of Approximate Reasoning*, 17:77–103, 1997.
- [PV98] J.B. Paris and A. Vencovska. Proof systems for probabilistic uncertain reasoning. *Journal of Symbolic Logic*, 63(3):1007–1039, 1998.
- [Ram50] F.P. Ramsey. General propositions and causality. In R.B. Braithwaite, editor, *Foundations of Mathematics and other logical essays*, pages 237–257. Routledge and Kegan Paul, New York, 1950.
- [Rei80] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [RKI93] H.P. Reidmacher and G. Kern-Isberner. Unsichere logische Regeln in Expertensystemen mit probabilistischer Wissensbasis. Fachbereich Wirtschaftswissenschaften, Diskussionsbeitrag 206, FernUniversität Hagen, 1993.
- [RKI97a] W. Rödder and G. Kern-Isberner. Léa Sombé und entropie-optimale Informationsverarbeitung mit der Expertensystem-Shell SPIRIT. *OR Spektrum*, 19/3, 1997.
- [RKI97b] W. Rödder and G. Kern-Isberner. Representation and extraction of information by probabilistic logic. *Information Systems*, 21(8):637–652, 1997.
- [RM96] W. Rödder and C.-H. Meyer. Coherent knowledge processing at maximum entropy by SPIRIT. In E. Horvitz and F. Jensen, editors, *Proceedings 12th Conference on Uncertainty in Artificial Intelligence*, pages 470–476, San Francisco, Ca., 1996. Morgan Kaufmann.
- [Rot91] H. Rott. A nonmonotonic conditional logic for belief revision, part I: Semantics and logic of simple conditionals. In A. Fuhrmann and M. Morreau, editors, *The logic of theory change*, pages 135–181. Springer, Berlin, Heidelberg, New York, 1991.

- [SA95] R. Srikant and R. Agrawal. Mining generalized association rules. In *Proceedings of the 21st VLDB Conference*, Zürich, Switzerland, 1995.
- [Sch91] K. Schlechta. Theory revision and probability. *Notre Dame Journal of Formal Logic*, 32(2):45–78, 1991.
- [Sch97] K. Schlechta. *Nonmonotonic logics: basic concepts, results, and techniques*, volume 1187. Springer Lecture Notes in Artificial Intelligence, Berlin Heidelberg New York, 1997.
- [Sch98] M. Schweitzer. Wissensfindung in Datenbanken auf probabilistischer Basis. Master's thesis, FernUniversität Hagen, 1998. (Diplomarbeit).
- [SE99] M. Schramm and W. Ertel. Reasoning with probabilities and maximum entropy: the system PIT and its application in LEXMED. In *Symposium on Operations Research, SOR'99*, 1999.
- [SF97] M. Schramm and V. Fischer. Probabilistic reasoning with maximum entropy – the system PIT. In *Proceedings of the 12th Workshop on Logic Programming*, 1997.
- [SGS93] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction and Search*. Number 81 in Lecture Notes in Statistics. Springer, New York Berlin Heidelberg, 1993.
- [Sha76] G. Shafer. *A mathematical theory of evidence*. Princeton University Press, Princeton, NJ, 1976.
- [Sha86] G. Shafer. Probability judgment in Artificial Intelligence. In L.N. Kanal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 127–135. North-Holland, Amsterdam, 1986.
- [Sho86] J.E. Shore. Relative entropy, probabilistic inference and AI. In L.N. Kanal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 211–215. North-Holland, Amsterdam, 1986.
- [SJ80] J.E. Shore and R.W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, IT-26:26–37, 1980.
- [SJ81] J.E. Shore and R.W. Johnson. Properties of cross-entropy minimization. *IEEE Transactions on Information Theory*, IT-27:472–482, 1981.
- [Som92] Léa Sombé. *Schließen bei unsicherem Wissen in der Künstlichen Intelligenz*. Vieweg, Braunschweig, 1992.
- [Spi91] M. Spies. Combination of evidence with conditional objects and its application to cognitive modeling. In I.R. Goodman, M.M. Gupta, H.T. Nguyen, and C.S. Rogers, editors, *Conditional logic in expert systems*, pages 181–209. Elsevier, North-Holland, 1991.
- [Spo88] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics, II*, pages 105–134. Kluwer Academic Publishers, 1988.
- [SW76] C.E. Shannon and W. Weaver. *Mathematische Grundlagen der Informationstheorie*. Oldenbourg, München Wien, 1976.

- [TGK91] H. Thöne, U. Güntzer, and W. Kiessling. Probabilistic reasoning with facts and rules in deductive databases. In *Proceedings European Conference on Symbolic and Quantitative Approaches for Uncertainty (ECSQAU)*, Marseille, pages 333–337, Berlin, Heidelberg, New York, 1991. Springer.
- [TGK92] H. Thöne, U. Güntzer, and W. Kiessling. Towards precision of probabilistic bounds propagation. In D. Dubois, M.P. Wellmann, B. D'Ambrosio, and P. Smets, editors, *Proceedings 8th Conference on Uncertainty in Artificial Intelligence*, pages 315–322, San Mateo, Ca., 1992. Morgan Kaufmann.
- [Thi89] H. Thiele. Monotones und nichtmonotones Schliessen. In J. Grabowski, K.P. Jantke, and H. Thiele, editors, *Grundlagen der Künstlichen Intelligenz*, pages 80–160. Akademie-Verlag, Berlin, 1989.
- [Tou86] D. Touretzky. *The Mathematics of Inheritance Systems*. Pitman and Morgan Kaufman, London and Los Altos, California, 1986.
- [Voo96a] F. Voorbraak. Probabilistic belief expansion and conditioning. Technical Report LP-96-07, Institute for Logic, Language and Computation, University of Amsterdam, 1996.
- [Voo96b] F. Voorbraak. Reasoning with uncertainty in AI. In L. Dorst, M. van Lambalgen, and F. Voorbraak, editors, *Proceedings Reasoning with Uncertainty in Robotics (RUR'95)*, LNCS/LNAI 1093, pages 52–90, Berlin, 1996. Springer.
- [Whi90] J. Whittaker. *Graphical models in applied multivariate statistics*. John Wiley & Sons, New York, 1990.
- [Wil94] M.-A. Williams. Transmutations of knowledge systems. In *Proceedings Fourth International Conference on Principles of Knowledge Representation and Reasoning, KR-94*, pages 619–629. Morgan Kaufman, 1994.
- [Yag85] R.R. Yager. Inference in a multivalued logic system. *Int. J. Man-Machine Studies*, 23:27–44, 1985.
- [Zad83] L.A. Zadeh. The role of fuzzy logic in the management of uncertainty in expert systems. *Fuzzy Sets Syst.*, 11:199–227, 1983.

Index

\top	8	$Mod(\Psi)$	21
\perp	8	$\min(A; \Psi)$	21
Ω	8	$Bel(\Psi)$	18
$\widehat{\Omega}$	43	$Bel(P)$	29
$\widehat{\Omega}_0$	46	$Bel(\kappa)$	29
Ω_+	48	$Th(P)$	10
$\widehat{\Omega}_+$	48	$Th^*(\Psi)$	31
$\widehat{\Omega}_+^*$	63	$(B A)(\omega)$	28
\widehat{A}	47	$(B A)$	9, 27
$form(\omega_1, \omega_2, \dots)$	9	$(B A)[x]$	9
\mathcal{AP}	81	$(B A)[n]$	31, 35
\mathcal{R}	10	$(B A)^+$	36
\mathcal{B}	123	$(B A)^-$	36
\mathcal{V}	8	$A \rightarrow B[x]$	96
\mathcal{L}	8	$P(A)$	9, 29
\mathcal{L}^*	12	$P(B A)$	10, 29
$(\mathcal{L} \mathcal{L})$	9	$\kappa(A)$	29
$(\mathcal{L} \mathcal{L})^{prob}$	9, 31	$\kappa(B A)$	30
$(\mathcal{L} \mathcal{L})^{prob, \rightarrow}$	96	$\Pi(A)$	30
$(\mathcal{L} \mathcal{L})^{OCF}$	31	$\Pi(B A)$	30
$(\mathcal{L} \mathcal{L})^A$	35	$V(A)$	33
$(\mathcal{L} \mathcal{L})^*$	31	$V(B A)$	32
$\mathcal{E}_{\mathcal{V}}^*$	30	P_0	24
\mathcal{E}^*	104	P_{ME}	25, 76
$\mathcal{E}^*(C\Psi)$	105	$P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-]$	77
$C(\mathcal{A})$	12	$P[\alpha_1, \dots, \alpha_n]_F$	84
$Cn(\mathcal{A})$	13	$H(P)$	24
$Cn_{prob}(\mathcal{R})$	92	$R(Q, P)$	24
$Cn^*(\mathcal{R})$	105	$wf(P^*)$	78
\mathbf{C}	105	$WF(P, \mathcal{R})$	78
C_P^{ME}	93	$WQ_F(P, \mathcal{R})$	84
C_{Ψ}	105	$\psi_{\omega, \omega'}$	37
$C_{v, l}$	123	$\psi_{v, l}$	123
$C(P, \mathcal{R})$	76	\models	9, 10, 58
$Mod(A)$	8	\models^{prob}	92
$Mod(\mathcal{R})$	10	\models^*	104
$Mod^*(\mathcal{R})$	104	\equiv	9, 28
$Mod(\kappa)$	29	$\equiv_{\mathcal{R}}$	45
		\equiv_{\sqsubseteq}	110

\equiv_g 126
 $[\hat{\omega}]_{\equiv_{\mathcal{R}}}$ 45
 \vdash 12
 \vdash_P^{ME} 93
 \leq_K 16
 \leq_{Ψ} 20, 22
 $\leq_{\mathcal{A}}$ 32
 \ll 76
 $+$ 14
 $*$ 14
 $*_S$ 60
 $*_{DP}$ 60
 $*_0$ 61
 $*_{ME}$ 89, 93
 $*_F$ 84
 $-$ 15
 \diamond 16
 \circ 110
 \oplus 31
 \odot 31
 $0^{\mathcal{A}}$ 31
 $1^{\mathcal{A}}$ 31
 \sqsubseteq 36, 110
 \sqsubseteq_{Ψ} 106
 \sqcap 37
 \sqcup 37
 \perp 39
 \mathbf{a}_i^+ 41
 \mathbf{a}_i^- 41
 $\mathbf{b}_{v,l}^+$ 123
 $\mathbf{b}_{v,l}^-$ 123
 $\mathbf{b}_{v,l}$ 125
 $\mathcal{F}_{\mathcal{R}}$ 41
 $\mathcal{F}_{\mathcal{B}}$ 123
 σ_i 42
 $\sigma_{(B|A)}$ 41
 $\sigma_{\mathcal{R}}$ 42
 $\sigma_{\mathcal{B}}$ 123
 $\sigma_{v,l}$ 123
 $\ker \sigma$ 44
 $\ker_0 \sigma$ 47

adjustment condition 71
 affirmative set 36
 AGM-postulates
 – for revising epistemic states 19
 – for contraction 15
 – for expansion 14
 – for revision 14
 antecedent 9

antecedent conjunction problem 86,
 101
 association rule 120
 – confidence 120
 – support 120
 atomicity condition 82
 atomicity principle 79

 belief base 109
 belief set 14

c-adaptations 67
 categorical specificity 98
 cautious cut 100
 cautious monotonicity 12, 99
 condition
 – adjustment 71
 – atomicity 82
 – continuity 82
 – positivity 71
 – relevance 81
 – smoothness 21
 – uniqueness 85
 conditional 9, 18, 27, 28
 – basic 37
 – contradictory 36
 – infimum of conditionals 37
 – OCF-conditional 35
 – perpendicular conditionals 39
 – probabilistic 9, 35
 – quantified 34
 – single-elementary 121
 – subconditional 36
 – supremum of conditionals 37
 – tautological 36
 – validity of 34
 conditional function
 – ordinal 29
 – simple 18
 conditional independence 33
 conditional indifference
 – strict conditional indifference with
 respect to \mathcal{R} 49
 – strict conditional indifference with
 respect to \mathcal{R} and V 63
 – weak conditional indifference with
 respect to \mathcal{R} 49
 – with respect to \mathcal{R} and V 63
 conditional inference operation 105
 – complete 105
 conditional preservation 53

- principle of conditional preservation 63
- conditional structure
 - of ω (representation of) 42
 - of $\hat{\omega}$ (representation of) 43
 - with respect to \mathcal{R} 45
- conditional valuation function 32
 - (strictly) faithful with respect to \mathcal{R} 50
 - uniform 33
- conditionalization 13, 59, 95
 - (A, m) -conditionalization 60
- conflicting evidence 99
- conflicting set 36
- conjunction
 - complete 8
 - elementary 8
- conjunction left 100
- conjunction right 100
- consequence 9
- consequence operation 13
- consistent 10
 - P-consistent 76
 - strictly V-consistent 63
 - V-consistent 63
- continuity condition 82
- contraction 15
- c-representation 67
 - faithful 68
- c-revision 67
 - faithful 68
- cross-entropy 24
- cumulativity 12
 - strong 107
- cut 12, 92
- data mining 5, 119
- distribution 13
 - probability 28
- DP-postulates 22
- entropy 24
- epistemic state 17
- erasure 16
- expansion 14
- fact 36
- faithful assignment 16
 - for epistemic states 20
- focusing 116
- founded 106
- free abelian group 41
- full absorption 13
- function
 - conditional valuation 32
 - ordinal conditional 29
 - probability 28
 - ranking 29
 - relative change 63
 - simple conditional 18
- Harper identity 15
- idempotence 92
- identity
 - Harper 15
 - Levi 15
- inclusion 12, 92
- inference operation
 - conditional 105
 - ME-inference operation 93
 - nonmonotonic 12
 - probabilistic 93
 - universal 105
- inference relation
 - ME-inference relation 93
 - nonmonotonic 12
 - probabilistic 93
- infimum of conditionals 37
- informational economy 25
- inverse maxent problem 131
- inverse representation problem 4
- Jeffrey conditionalization 23
- Jeffrey’s rule 23
- kernel (of a homomorphism) 44
- KM-postulates (for updating) 16
- knowledge discovery 5, 119
- left absorption 13
- left logical equivalence 13
- LEG-networks 137
- Levi identity 15
- limit assumption 21
- literal 8
- logical coherence 86
- loop 13
- ME-inference 93
 - categorical specificity 98
 - cautious cut 100

- cautious monotonicity 99
- conjunction left 100
- conjunction right 100
- reasoning by cases 101
- transitive chaining 98
- ME-inference operation 93
- ME-inference relation 93
- ME-principle 25
- monotonicity 11, 13, 92
 - cautious 12
 - rational 13
- neighboring worlds 8
- neutral set 36
- nonmonotonic inference operation 12
 - cumulative 12
 - idempotent 12
 - reflexive 12
 - supraclassical 13
- nonmonotonic inference relation 12
- OCF *see* ordinal conditional function
- OCF-conditional 35
- ordinal conditional function 29
- P*-consistency 76
- plausibility (pre-)ordering 21
- plausibility measure 33
- positivity condition 71
- possibilistic logic 139
- possibility distribution 30
- possibility measure 30
- postulates
 - for revising by sets of conditionals: 113
 - AGM-postulates 14
 - DP-postulates 22
 - for conditional base revision 111
 - for conditional revision 55
 - KM-postulates 16
- premise 9
- principle of conditional preservation 63
- principle of maximum entropy 25
- principle of minimum cross-entropy 25
- probabilistic conditional 9, 35
- probabilistic equivalence 87
- probabilistic fact 10
- probabilistic inference operation 93
 - complete 93
- probabilistic inference relation 93
- probabilistic rule 9
- probabilistically equivalent 10
- probability distribution 9, 28
- probability function 28
- Ramsey test 18
- ranking 21
- rational monotonicity 13
- reasoning by cases 101
- reciprocity 12
- relative change function 63
- relevance condition 81
- representation problem 67
- revision 14
- right weakening 13
- simple conditional function 18
- Simpson's paradox 39
- smoothness condition 21
- SPIRIT 137
- standard probabilistic consequence operation 92
- standard probabilistic consequence relation 92
- strong cumulativity 107
- subconditional 36
- supremum of conditionals 37
- system P 23
- system-Z 69
- transitive chaining 98
- transmutations 53
- triviality result 20
- uniqueness assumption 31
- uniqueness condition 85
- universal inference operation 105
 - cumulative 105
 - faithful 106
 - founded 106
 - idempotent 105
 - preserving consistency 105
 - reflexive 105
 - strongly cumulative 107
- update 16
- V-consistency 63, 106
- weight factors 78
- weight quotients 84

1. Introduction

Conditionals are closely connected with reasoning – typically, they suggest plausible, yet often defeasible conclusions from what is evident. Therefore, studying conditionals means primarily to overcome the strict framework of classical logic and to enter into the world of defeasible reasoning, nonmonotonic logics and uncertain knowledge.

1.1 “Believe It or Not” – The Handling of Uncertain Knowledge

From the beginnings of Artificial Intelligence, commonsense and expert-like reasoning has been modelled in two basically different ways: The *quantitative approaches* using certainty factors [BS84, Voo96b], fuzzy logic [Zad83, Yag85, KGK93], belief functions [Dem67, Sha76, Sha86], possibilities [DP92, DLP94], and probabilities [DeF74, Pea88, Bac90, Par94], and the *symbolic approaches* like circumscription [McC80], autoepistemic logic [Moo88], default logic [Rei80] and many others (the references given are only examples). While the former methods have proved to be successful in practical reasoning, the latter have helped to reveal and model structures of uncertain reasoning in a qualitative way.

Probability theory here occupies an outstanding position: Devised and developed to perform sound reasoning in a quantitative setting, it not only became an important benchmark in the area of quantified reasoning in general, but also provides a semantics for qualitative default reasoning by considering infinitesimal probabilities (see [Ada75, Pea89]) or orders of magnitudes of probabilities (see [GP96]). Within the last decades, knowledge representation and reasoning based upon probability theory has received increasing attention in the area of artificial intelligence. Probability theory provides a solid foundation for nonmonotonic reasoning methods ([Ada75, Bac90, Gef92, Gol94, Pea88]), and probabilistic networks allow a consistent computation of (quantified) uncertainty ([Pea86, LS88, RKI97b, RM96]). Probabilities are particularly appropriate to quantify conditional statements “*If*

A then B ” which are of major interest in the areas of *nonmonotonic reasoning* and *belief revision* dealing with the dynamics of belief (see, for instance, [Nut80, Cox46, Cal91, DP91b, DP97a, LM92]).

The property of monotonicity is undoubtedly crucial for classical deduction: Adding new facts does not invalidate previously derived conclusions, so the set of conclusions may only increase monotonically. Thus it establishes a very solid fundament for exact sciences like mathematics. From daily experiences, however, we know that monotonicity is not an appropriate guideline for human reasoning: New information often makes us revising our beliefs, i.e. some beliefs which turn out to be false are given up, and other beliefs supported by the new information are accepted. Due to the fact that most of our knowledge is uncertain or incomplete, this process is so typical and so successful in our everyday lives that the then harsh debates attacking nonmonotonicity as detrimental to logics (cf. [Isr87]) appear quite whimsical nowadays. Just to the contrary – Dubois and Prade [DP96] state *nonmonotonicity* as one of the most salient features an exception-tolerant inference system has to possess.

Nevertheless, nonmonotonic reasoning is still a challenge – which beliefs are to be given up, which to be established? In their early paper [MD80], McDermott and Doyle tried to specify the scope of their “*Nonmonotonic logic*” between arbitrariness and rigidity:

The purpose of non-monotonic inference rules is not to add certain knowledge where there is none, but rather to guide the selection of tentatively held beliefs in the hope that fruitful investigations and good guesses will result. This means that one should not a priori expect non-monotonic rules to derive valid conclusions independent of the monotonic rules. Rather one should expect to be led to a set of beliefs which while perhaps eventually shown incorrect will meanwhile coherently guide investigations.

Belief revision, on the other hand, deals with the *dynamics of belief* – how should currently held beliefs be modified in the light of new information? Results in this area are mainly influenced by the so-called *AGM theory*, named after Alchourron, Gärdenfors and Makinson who set up a framework of postulates for a reasonable change of beliefs ([AGM85, Gär88]).

This book exploits the crucial relationship between plausible uncertain reasoning and conditionals to obtain a unified and enhanced framework for studying nonmonotonic reasoning and belief revision: We will investigate how to revise epistemic states by (sets of) conditionals in different settings, including a purely qualitative, a probabilistic, and an intermediate environment using ordinal conditional functions for representation. Epistemic states repre-

sent the cognitive state of an intelligent agent at a given time. They permit us to graduate beliefs according to their plausibility and thus allow a more appropriate studying of belief change than plain propositional *belief sets* which are the objects of interest in AGM theory. While AGM theory only observes the results of revisions, considering epistemic states under change focuses on the mechanisms underlying that change, taking conditional beliefs as *revision policies* explicitly into account. So the work presented here meets a crucial demand raised in Friedman and Halpern’s *Critique* [FH99] to AGM revision: “... *whatever we take to be our representation of the epistemic state, it seems appropriate to consider how these representations should be revised.*”

The idea to consider conditionals, nonmonotonic reasoning and belief revision from a common point of view is not new. Indeed, the crucial role of “conditional objects” has been recognized for many years (see [DP91a, KS91, Bou94, FH94]). In this book, however, conditionals are not considered as logical entities, but as dynamic agents shifting worlds in order to establish beliefs. This understanding of conditionals has far-reaching consequences and yields a theory which is quite different from the ones raised by logical considerations.

In separating structural from numerical aspects when handling conditionals, the basic notions for conditionals developed here may be applied to yield important results both in a qualitative and in a numerical setting. Indeed, conditionals are at the heart of this book from which the other major topics – revising epistemic states, extended nonmonotonic reasoning and knowledge discovery – will be developed. *Conditional valuation functions* will be introduced as abstract (numerical) representations of epistemic states covering probability functions, ordinal conditional functions and possibility distributions. The notion of a *conditional structure* defined for (multi-)sets of possible worlds allows us to formalize correctly the idea of *indifference* of conditional valuation functions with respect to sets of conditionals, resulting in the statement of a *principle of conditional preservation* for revisions. Within a purely qualitative environment, we set up postulates describing what it means for a revision to preserve conditional beliefs. Moreover, we show that the (numerical) principle of conditional preservation actually generalizes these postulates. Thus a thorough axiomatization of this principle is obtained which constitutes an important paradigm when revising epistemic states, similar to the paradigm of *minimal propositional change* guiding AGM-revisions.

Besides conditionals, a second focus of this book is on *probabilistic reasoning at optimum entropy* where the principles of maximum and minimum entropy (*ME-principles*), respectively, will be used to represent incomplete probabilistic knowledge in an information-theoretically sound way. It can easily be seen that each ME-revision satisfies the principle of conditional preservation and therefore fits the formal framework sketched above. We will

investigate what other postulates are necessary to characterize ME-revisions as a “best” probabilistic revision function, identifying an appropriate *functional concept* and the properties of *logical coherence* and *representation invariance* as responsible for the special form of ME-revisions. This establishes reasoning at optimum entropy as a most fundamental inference method when using probabilistic conditionals to represent knowledge.

Integrating ME-inference and ME-revision, respectively, into the frameworks of nonmonotonic reasoning and belief revision turns out to be particularly fascinating and fruitful: On the one hand, the abstract properties of nonmonotonic inference operations like *cumulativity*, *left logical equivalence* etc. ([Gab85, KLM90, Mak94]) and the axioms for belief revision ([Gär88, DP94]) help to classify the ME-techniques from a formal point of view, thereby raising the reputation of this powerful, but sometimes seemingly obscure method. On the other hand, however, studying ME-inference may give important impetus to the field of nonmonotonic logics and belief revision in general. For a long time, both areas have been concentrating on handling only propositional beliefs in a one-step manner, without basing inferences explicitly on a theory and without considering iterated revisions. In contrast to this, the principles of optimum entropy provide a comprehensive frame to realize iterated revisions of (probabilistic) epistemic states by sets of conditionals, thus generalizing classical AGM-revision in nearly all aspects. And indeed, the property of *logical coherence*, mentioned already in [SJ81] and here used as one of the postulates to characterize ME-revision, may be read as a set-theoretical version of Darwiche and Pearl’s axiom (C1) for iterated revision ([DP97a]).

The axiom of logical coherence may be formulated easily for general epistemic states, and we interpret it as a kind of “*cumulativity with respect to epistemic states*” phrased for *universal inference operations*. Universal inference operations provide a global setting to study inference and revision, and logical coherence proves to be an important means to link up inferences based on different epistemic states. Considering revision operators in the enhanced framework of revising epistemic states by sets of conditionals allows us to differentiate between *simultaneous* and *successive revision*, and to separate *background* from *evidential knowledge*. This permits us to distinguish more clearly between different belief change operations like (genuine) *revision*, *updating* and *focusing* which may be realized, however, by the same (binary) revision operator.

Furthermore, the principle of conditional preservation presented here was first developed for ME-revision in [KI98a], but it can also be applied in a more general framework. The conditionals to be incorporated impose a specific structure on the resulting probability distribution. This structure may be used to elaborate sets of probabilistic conditionals that represent a given distribution (*inverse representation problem*). This is particularly important

when one wants to extract probabilistic knowledge from statistics in order to present relevant connections between the considered variables (*knowledge discovery/data mining*) and/or to design a probabilistic network for knowledge representation and reasoning. The first steps in this direction will be undertaken in this book.

The principal concern of this book is to develop a common approach to symbolic and numerical uncertain reasoning via conditionals. Most of the topics to be treated here, however, were actually addressed when working on a concrete computational system, namely the expert system shell SPIRIT realizing maximum entropy propagation ([RKI97b, RKI97a, RM96]). Here the following questions arose: Besides respecting (conditional) independence, what are the mechanisms underlying ME-techniques? How can ME-inference and ME-adaptation be compared to other methods? Conditionals are generally considered to be very important for knowledge representation and reasoning, but how can their meaning and effects be made explicit? And, last not least, a crucial problem in designing expert systems: Where do all the conditionals representing substantial knowledge come from? How should we use experimental data?

This book aims at answering all these questions by presenting a general framework for nonmonotonic reasoning and belief revision that features conditionals and ME-methods as particularly meaningful both to qualitative and quantitative approaches.

1.2 Overview

The organization of this paper is as follows: Fixing *basic definitions and notations* in the next section will conclude this introduction.

Chapter 2 outlines the state of the art of *belief revision and nonmonotonic reasoning*. Several properties of nonmonotonic inference operations which will be used in this book are listed here. In the area of belief revision, the standard AGM-theory dealing with expansion, revision and contraction of propositional beliefs is recalled, and we explain the difference between revision and updating in the sense of Katsuno and Mendelzon. Then we discuss how to extend this framework of propositional belief change by studying epistemic states and conditionals, allowing us to perform iterative revisions. Finally, we present a picture of belief revision and nonmonotonic reasoning from a probabilistic point of view, featuring revisions and inferences based on the principles of optimum entropy as particularly sophisticated and powerful methods.

Chapter 3 focuses on *conditionals* and starts with studying the connections between conditionals and epistemic states. Conditional valuation functions are introduced, and we discuss how the acceptance of (conditional) beliefs in epistemic states may be modelled by these functions. Then we turn to more formal things: In Section 3.4, we define two relations, *subconditionality*, \sqsubseteq , and *perpendicularity*, $\perp\!\!\!\perp$, on conditionals, describing quite extreme ways of conditionals interacting with one another. These relations will prove to be useful especially in a qualitative setting of belief change and in the field of (conditional) knowledge discovery. An even more important notion is introduced in Section 3.5: We represent conditionals by generators of a (free-abelian) group and define the *conditional structure* of a world. Using group theoretical structures makes it possible to calculate with conditionals, or with their effects on worlds, respectively. These formal means provide a framework adequate to phrase exactly, what *conditional indifference* of a conditional valuation function with respect to a set of conditionals means (see Section 3.6). We show that probability functions and ordinal conditional functions which are indifferent with respect to some set of conditionals follow quite a simple conditional-logical pattern.

Conditional indifference will prove to be of crucial importance when *revising epistemic states by conditional beliefs* in Chapter 4 in that it is the essential ingredient to formalizing a quantitative *principle of conditional preservation* for revising conditional valuation functions in Section 4.5. Revisions by sets of conditionals and representations of sets of conditionals will be called *c-revisions* and *c-representations*, respectively. But first, we will describe what it means to preserve conditional beliefs in a purely qualitative setting by stating postulates for revising an epistemic state by a conditional in Section 4.1. Representation theorems for these postulates will be given in Section 4.2, and they will be exemplified by presenting a revision operator for ordinal conditional functions in Section 4.4. We investigate the meaning of conditional valuation functions for qualitative revisions in Section 4.3, and in Section 4.5, we show that both approaches to the principle of conditional preservation developed here, the qualitative one and the quantitative one, are compatible.

The idea of revisions obeying the principle of conditional preservation is pursued further in a probabilistic environment in Chapter 5. We elaborate three more postulates such a revision should satisfy: First, a *functional concept* should establish a clear and unique connection between prior knowledge, new information and the revised probability distribution (see Section 5.2). Second, we present the postulate for *logical coherence* in Section 5.3. This postulate claims that revised probability distributions can be used unambiguously as priors for further revisions and thus is of crucial importance particularly for *iterated revisions*. Finally, the postulate for *representation invariance* states that revisions should not depend on the syntactical re-

presentation of probabilistic knowledge (cf. Section 5.4). Following these four postulates, we arrive at a new characterization of revisions based on the principle of minimum cross-entropy (*ME-revisions*) (cf. Theorem 5.5.1 in Section 5.5).

We investigate how *ME-reasoning* works in Chapter 6. First we check which of the properties relevant to nonmonotonic inference operations are satisfied by ME-inference. In particular, we show that ME-inference is cumulative and fulfills the loop-property (cf. Section 6.2). Then we present some ME-deduction schemes in Section 6.4 to illustrate ME-reasoning in simple, but typical and informative situations such as *transitive chaining*, *cautious monotonicity* and *reasoning by cases*.

In Chapter 7, we return to general belief revision and nonmonotonic reasoning. Due to the formal manner in which numerical ME-inference is handled in this book, it is possible to transfer some crucial insights provided by this powerful inference operation to the general theory. ME-reasoning exemplifies effectively how a more comprehensive and unified view on this area is opened by considering revisions and inferences in an extended framework using epistemic states and conditionals. *Universal inference operations* are introduced in Section 7.1 as a proper counterpart to revision operators in nonmonotonic reasoning, allowing us to take a basic epistemic state into account. We show how to distinguish between *simultaneous* and *successive revision* (cf. Section 7.2), and how to separate clearly between *background* and *evidential knowledge* (cf. Section 7.3). This allows us in particular to overcome the conceptual difference between (genuine) revision (in the AGM-sense) and updating (in the sense of Katsuno and Mendelzon) by considering them not as different change operators, but as applying the same change operator in different ways (cf. Sections 7.4 and 7.5). Moreover, *focusing* may be realized as different from revision by the ME-revision operator (cf. Section 7.6). *Iterated revisions* may be dealt with adequately in that framework, too. The postulate for logical coherence used for ME-characterization proves to be of crucial importance to control iterated revision and to link up inference operations.

Chapter 8 brings a brief sketch of some results in *probabilistic knowledge discovery* and then turns to its main part, the *discovery of structures of knowledge by following conditional patterns* within conditional valuation functions. Revisions of such functions which obey the principle of conditional preservation are necessarily indifferent with respect to the revising set of conditionals. So discovering “conditional structures” means in particular finding a set of conditionals with respect to which the given conditional valuation function, e.g. a probability distribution, is indifferent. In Section 8.2, we develop an approach to accomplish this task by using the group theoretical representations of conditionals, developed in Chapter 3. Part of an algorithm is presented that allows us to calculate such a set of conditionals by studying numerical

relationships between the values given. This method is applied in Section 8.3 to illustrate how an ME-optimal representation of a probability distribution by a set of conditionals can be computed.

Chapter 9 presents briefly a selection of various computational approaches to *ME-reasoning with probabilistic conditionals*, to *probabilistic knowledge discovery* and to *possibilistic belief revision*.

Finally, Chapter 10 summarizes the results of this book.

Preliminary versions of various parts of this book have already been published in [RKI93, KIR96, KI96b, KI96a, KI97a, KI97c, KI97b, KI98c, KI98a, RKI97b, RKI97a, KI01, KI98b, KI99c, KI99b], and in [KI99a].

In order to improve the readability of the text, the full proofs of lemmata, propositions, corollaries and theorems have been moved to Appendix A.

1.3 Basic Definitions and Notations

1.3.1 Propositional and Conditional Expressions

We consider a propositional language $\mathcal{L} = \mathcal{L}(\mathcal{V})$ over a finite alphabet $\mathcal{V} = \{a, b, c, \dots\}$. Uppercase roman letters $A, B, C \dots$ will denote atoms or formulas in \mathcal{L} . \mathcal{L} is equipped with the usual logical connectives \wedge (*and*), \vee (*or*) and \neg (*negation*). We will largely avoid material implication in order not to get confused with conditional implication (see below). To simplify notations, we will replace a conjunction by juxtaposition and indicate the negation of a proposition by barring it, i.e.

$$AB = A \wedge B \quad \text{and} \quad \overline{A} = \neg A$$

\dot{A} will denote one of the formulas A, \overline{A} . *Elementary conjunctions* are conjunctions of *literals*, i.e. of positive or negated atoms. *Complete conjunctions* are elementary conjunctions which contain each atom either in positive or negated form. Tautologies and contradictions will be denoted by \top and \perp , respectively.

Let Ω denote the set of possible worlds, i.e. Ω is a complete set of interpretations of \mathcal{L} . Two worlds $\omega, \omega' \in \Omega$ are called *neighbors* if they differ with respect to exactly one atom.

Given a propositional formula $A \in \mathcal{L}$, we denote by $Mod(A)$ the set of all A -worlds,

$$Mod(A) = \{\omega \in \Omega \mid \omega \models A\}$$

Definition 1.3.1. For a set of worlds $\{\omega_1, \omega_2, \dots\} \subseteq \Omega$, we define

$$\text{form}(\omega_1, \omega_2, \dots) \in \mathcal{L}$$

to be that proposition in \mathcal{L} which has $\omega_1, \omega_2, \dots$ as its models:

$$\text{Mod}(\text{form}(\omega_1, \omega_2, \dots)) = \{\omega_1, \omega_2, \dots\}$$

Sometimes, worlds will be identified simply with their corresponding complete conjunction

$$\omega = \bigwedge_{v:\omega \models v} v \quad (1.1)$$

where the conjunction is taken over all atoms v in \mathcal{L} .

If $A, B \in \mathcal{L}$ are two propositional formulas in \mathcal{L} , then $A \leq B$ iff $A \models B$, i.e. iff $\text{Mod}(A) \subseteq \text{Mod}(B)$. \equiv means classical logical equivalence, that is $A \equiv B$ iff $\text{Mod}(A) = \text{Mod}(B)$.

\mathcal{L} is extended to a conditional language $(\mathcal{L} \mid \mathcal{L})$ by introducing a conditional operator \mid :

$$(\mathcal{L} \mid \mathcal{L}) = \{(B \mid A) \mid A, B \in \mathcal{L}\}$$

A is called the *antecedent* or the *premise* of $(B \mid A)$, and B is the *consequence* of the conditional $(B \mid A)$. $(\mathcal{L} \mid \mathcal{L})$ is taken to include \mathcal{L} by identifying a proposition A with the conditional $(A \mid \top)$.

1.3.2 Probabilistic Logics

The atoms in \mathcal{L} may be looked upon as (binary) propositional variables, and possible worlds or complete conjunctions, respectively, correspond to elementary events. So, given a probability distribution P over \mathcal{V} , a probability can be assigned to each propositional formula $A \in \mathcal{L}(\mathcal{V})$ via

$$P(A) = \sum_{\omega \models A} P(\omega)$$

In this way, a probabilistic interpretation of \mathcal{L} is obtained.

We extend $\mathcal{L}(\mathcal{V})$ to a *probabilistic conditional language* $(\mathcal{L} \mid \mathcal{L})^{\text{prob}}$ by attaching a probability $x \in [0, 1]$ to each conditional $(B \mid A) \in (\mathcal{L} \mid \mathcal{L})$:

$$(\mathcal{L} \mid \mathcal{L})^{\text{prob}} = \{(B \mid A)[x] \mid (B \mid A) \in (\mathcal{L} \mid \mathcal{L}), x \in [0, 1]\}$$

$(B \mid A)[x]$ is called a *probabilistic conditional*, or sometimes a *probabilistic rule*, too. It is to represent syntactically non-classical conditional assertions $(B \mid A)$ weighted with a degree of certainty x . Probabilistic conditionals are interpreted via conditional probabilities: If P is a distribution, we write

$$P \models (B|A)[x] \quad \text{iff} \quad P(A) > 0 \text{ and } P(B|A) = \frac{P(AB)}{P(A)} = x$$

A *probabilistic fact* $A[x]$ is regarded as equivalent to the probabilistic conditional $(A|\top)[x]$ with tautological antecedent, so $P \models A[x]$ iff $P(A) = x$.

In general, we have

$$x = P(B|A) \quad \text{iff} \quad P(A) > 0 \text{ and } (1-x)P(AB) = xP(\overline{AB}),$$

so the quotient $\frac{P(AB)}{P(\overline{AB})}$ determines the probability of the conditional $(B|A)$.

It represents the proportion of individuals or objects with property A which also have property B to those that do not. Thus it is crucial for accepting the conditional, not only within a probabilistic framework (cf. [Nut80]).

Probabilistic facts and conditionals will also be denoted by small Greek letters ϕ, ψ etc. So for a distribution P and for a set $\mathcal{R} \subseteq (\mathcal{L} | \mathcal{L})^{prob}$ of probabilistic conditionals, we write $P \models \mathcal{R}$ iff $P \models \phi$ for all $\phi \in \mathcal{R}$. *Models* of probabilistic conditionals are probability distributions that fulfill them, hence

$$Mod(\mathcal{R}) = \{Q \mid Q \text{ distribution over } \mathcal{V}, Q \models \mathcal{R}\}$$

for $\mathcal{R} \subseteq (\mathcal{L} | \mathcal{L})^{prob}$, $\mathcal{L} = \mathcal{L}(\mathcal{V})$. A set of probabilistic conditionals $\mathcal{R} \subseteq (\mathcal{L} | \mathcal{L})^{prob}$ is *consistent* iff it has a probabilistic model, i.e. iff there is a distribution Q such that $Q \models \mathcal{R}$. Two sets $\mathcal{R}_1, \mathcal{R}_2 \subseteq (\mathcal{L} | \mathcal{L})^{prob}$ are *probabilistically equivalent* iff $Mod(\mathcal{R}_1) = Mod(\mathcal{R}_2)$.

For a distribution P over \mathcal{V} , let

$$Th(P) = \left\{ (B|A)[x] \in (\mathcal{L} | \mathcal{L})^{prob} \mid P \models (B|A)[x] \right\}$$

denote the set of all probabilistic conditionals which are valid in P . $Th(P)$ explicitly represents the conditional knowledge embodied in P .

Two distributions P_1, P_2 are identical iff $P_1(\omega) = P_2(\omega)$ for all $\omega \in \Omega$, that is, $P_1 \models \omega[x]$ iff $P_2 \models \omega[x]$. So there is a one-to-one correspondence between distributions P and their theories $Th(P)$.

2. Belief Revision and Nonmonotonic Reasoning – State of the Art

The capability of revising knowledge and giving up conclusions in the light of conflicting evidence is one of the most outstanding features of commonsense reasoning. Though it seems to be practised in everyday life in a most natural and self-evident way, it challenges knowledge representation and inference procedures in AI because it clashes with the classical property of monotonicity. Therefore *defeasible* or *nonmonotonic reasoning*, as it is usually called, requires new formalisms to be realized adequately, and to date, a number of approaches to “nonmonotonic logic” have been proposed. Makinson and others ([Gab85, Mak94, KLM90]) set forth formal properties and axiom systems to judge and classify inference relations lacking monotonicity. Makinson’s work also covers quite general inference procedures not being based on classical structures.

The topic of *belief revision* is to investigate knowledge bases in change. The great variety of approaches that have been proposed to date, usually each method coming along with a descriptive axiom scheme (for a survey, cf. [GR94]), corresponds to the many different interpretations and names the term *change* has been given. Gärdenfors [Gär88] identified three fundamental types of belief change, *revision*, *expansion* and *update*. Katsuno and Mendelzon [KM91b] recommend *updating* to handle knowledge in a *changing world*. *Conditioning* has been regarded as an adequate method for revising probabilistic beliefs (see, for instance, [Par94, Gär88]), but Dubois and Prade [DP97b] emphasize that actually, conditioning does not correspond to revision but rather to *focusing*.

Nonmonotonic reasoning and belief revision are closely related but have different focusses: While studying and realizing *nonmonotonic inference relations* is the principal topic of the first, the latter is mainly concerned with investigating the resulting *changes in the belief sets* (or belief states).

2.1 Nonmonotonic Reasoning

A number of different methods to realize nonmonotonic reasoning have been proposed up to now, among others Reiter's default logic [Rei80, Ant97] and its variants ([Bre94, Ant97]), circumscription [McC80], autoepistemic logic [MD80, Moo88], inheritance diagrams [Tou86], maxiconsistent sets [Poo88], preferential models [KLM90], and, using quantified uncertainty, probabilistic, possibilistic and fuzzy logics (cf. [Pea88, DLP94, KGK93]); for a survey, see [Som92] or [GHR94]. In spite of the diversity of representation forms and inference techniques used it is possible to compare different nonmonotonic logics by focussing on the inference relations induced. This idea goes back to Gabbay [Gab85] and was later pursued and elaborated by Kraus, Lehmann, Magidor [KLM90] and Makinson [Mak89, Mak94]. A number of important properties to describe reasonable nonmonotonic inference relations have been identified, e.g. *cumulativity*, *loop* and *rational monotonicity*. In the sequel, we will give an overview so as to cover the scope of this book.

Let \mathcal{L}^* be any language used to represent relevant knowledge appropriately, e.g. in this book, \mathcal{L}^* may be one of \mathcal{L} , $(\mathcal{L} \mid \mathcal{L})$ or $(\mathcal{L} \mid \mathcal{L})^{prob}$. If \vdash is a (nonmonotonic) *inference relation* between sets of formulas and single formulas of \mathcal{L}^* ,

$$\vdash \subseteq 2^{\mathcal{L}^*} \times \mathcal{L}^*,$$

then the corresponding *inference operation* C is defined via

$$C(\mathcal{A}) = \{\phi \in \mathcal{L}^* \mid \mathcal{A} \vdash \phi\}$$

for sets of formulas $\mathcal{A} \subseteq \mathcal{L}^*$. Conversely, each inference operation $C : 2^{\mathcal{L}^*} \rightarrow 2^{\mathcal{L}^*}$ induces an inference relation \vdash by setting $\mathcal{A} \vdash \phi$ iff $\phi \in C(\mathcal{A})$. We will use both notations simultaneously, depending on which appears more intuitive. In generalizing the inference relation from single formulas on its right side to sets of formulas, we will write $\mathcal{A} \vdash \mathcal{B}$ for $\mathcal{B} \subseteq C(\mathcal{A})$.

Let \mathcal{A}, \mathcal{B} be sets of formulas of \mathcal{L}^* , ϕ be a single formula. A (nonmonotonic) inference operation C is called *reflexive* if $\mathcal{A} \subseteq C(\mathcal{A})$, and *idempotent* if $C(C(\mathcal{A})) = C(\mathcal{A})$. C is said to satisfy *cut* if $\mathcal{A} \subseteq \mathcal{B} \subseteq C(\mathcal{A})$ implies $C(\mathcal{B}) \subseteq C(\mathcal{A})$, and fulfills *cautious monotonicity* if $\mathcal{A} \subseteq \mathcal{B} \subseteq C(\mathcal{A})$ implies $C(\mathcal{A}) \subseteq C(\mathcal{B})$. Cut and cautious monotonicity together yield the property of *cumulativity*

$$\mathcal{A} \subseteq \mathcal{B} \subseteq C(\mathcal{A}) \quad \text{implies} \quad C(\mathcal{A}) = C(\mathcal{B}) \quad (2.1)$$

guaranteeing a convenient stability of inference: Taking already inferred knowledge into account does not change inferences. A *cumulative inference operation* is assumed to satisfy reflexivity (synonym: *inclusion*) besides cumulativeness. So a cumulative inference operation also fulfills the condition of *reciprocity*: If $\mathcal{A} \subseteq C(\mathcal{B})$ and $\mathcal{B} \subseteq C(\mathcal{A})$ then $C(\mathcal{A}) = C(\mathcal{B})$ (cf. [Mak94]).

Compared to a classical deduction operation, a cumulative inference operation only differs with respect to monotonicity: Instead of full monotonicity, i.e. $\mathcal{A} \subseteq \mathcal{B}$ implies $C(\mathcal{A}) \subseteq C(\mathcal{B})$, we have cautious monotonicity. A classical inference operation satisfying inclusion, monotonicity and cut is called a *consequence operation* and is denoted by Cn .

If \mathcal{L}^* is a propositional language, $\mathcal{L}^* = \mathcal{L}$, then a (nonmonotonic) inference operation C can be linked to a classical deduction operation Cn via the property of *supraclassicality*, i.e. $Cn(\mathcal{A}) \subseteq C(\mathcal{A})$ for all $\mathcal{A} \subseteq \mathcal{L}$: Any formula, that can be deduced from \mathcal{A} classically, can also be derived nonmonotonically. Other conditions connecting nonmonotonic inference operations to classical logic are the following:

- *left absorption*: $Cn C = C$;
- *full absorption*: $Cn C = C = C Cn$;
- *right weakening*: If $\phi \in C(\mathcal{A})$ and $\psi \in Cn(\phi)$ then $\psi \in C(\mathcal{A})$;
- *left logical equivalence*: $Cn(\mathcal{A}) = Cn(\mathcal{B})$ implies $C(\mathcal{A}) = C(\mathcal{B})$;
- *distribution*: $C(\mathcal{A}) \cap C(\mathcal{B}) \subseteq C(Cn(\mathcal{A}) \cap Cn(\mathcal{B}))$;
- *conditionalization*: If $\psi \in C(\mathcal{A} \cup \phi)$ then $\phi \Rightarrow \psi \in C(\mathcal{A})$ (where \Rightarrow means material implication).

Two more properties are of relevance when studying the behavior of nonmonotonic inference relations:

An inference relation \vdash (or operation C , respectively) is said to satisfy *loop* if

$$\text{whenever } \mathcal{A}_1 \vdash \mathcal{A}_2 \vdash \dots \mathcal{A}_n \vdash \mathcal{A}_1 \text{ then } C(\mathcal{A}_i) = C(\mathcal{A}_j) \text{ for } i, j \leq n \quad (2.2)$$

The loop-property is a weakened form of transitivity which typically does not hold for nonmonotonic inferences.

Finally, we recall the non-Horn condition of *rational monotonicity* (cf. [LM92]):

$$\text{If } \mathcal{A} \vdash \psi \text{ and not } \mathcal{A} \vdash \neg\phi \text{ then } \mathcal{A} \cup \{\phi\} \vdash \psi \quad (2.3)$$

Rational monotonicity is quite a strong condition for nonmonotonic logics, assuming anything as irrelevant to the inference $\mathcal{A} \vdash \psi$ the negation of which cannot be inferred from \mathcal{A} .

2.2 Belief Revision

In general, belief revision means the process of adapting some set of beliefs, or some accepted knowledge, respectively, to new information. Belief revision has many facets, depending on the compatibility of old and new beliefs,

whether belief is added or given up, or if only generic knowledge is applied to evidential information. Following [Gär88, AGM85] we will sketch the principal types of change for propositional beliefs – expansion, revision and contraction – and the corresponding catalogues of postulates, each describing a reasonable change of beliefs; these basic postulates are known today as the *AGM-theory*, named after *Alchourron*, *Gärdenfors* and *Makinson*.

Let K be a propositional *belief set*, that is, $K \subseteq \mathcal{L}$ is a set of propositions which is closed under classical consequence Cn . Let A be some proposition representing the newly acquired information which K is to be revised by.

The simplest type of revision occurs if A is consistent with K , i.e. if A does not contradict any of the beliefs in K . This type of revision is called *expansion*, denoted by $+$. Gärdenfors [Gär88] lists intuitive postulates for an expansion which are apt to characterize expansion uniquely within a classical logical framework:

AGM-postulates for expansion:

(AGM +1) $K + A$ is a belief set.

(AGM +2) $A \in K + A$.

(AGM +3) $K \subseteq K + A$.

(AGM +4) If $A \in K$ then $K + A = K$.

(AGM +5) If $K \subseteq H$ then $K + A \subseteq H + A$.

(AGM +6) $K + A$ is the smallest belief set satisfying (AGM +1) - (AGM +5).

Theorem 2.2.1 ([Gär88]). *The expansion function $+$ satisfies (AGM +1) - (AGM +6) iff $K + A = Cn(K \cup \{A\})$.*

Belief revision in general (operator: $*$) does not presuppose the consistency of the new information A with the belief set K ; if consistency holds, however, revision reduces to expansion:

AGM-postulates for revision:

(AGM *1) $K * A$ is a belief set.

(AGM *2) $A \in K * A$.

(AGM *3) $K * A \subseteq K + A$.

(AGM *4) If $\neg A \notin K$ then $K + A \subseteq K * A$.

(AGM *5) $K * A$ is inconsistent iff A is contradictory.

(AGM *6) If A and B are logically equivalent, then $K * A = K * B$.

(AGM *7) $K * A \wedge B \subseteq (K * A) + B$.

(AGM *8) If $\neg B \notin K * A$ then $(K * A) + B \subseteq K * A \wedge B$.

Unlike for expansion, these postulates are not sufficient to describe uniquely one optimal revision operator; they only outline the scope of reasonable belief revision. Katsuno and Mendelzon [KM91a] rephrased the AGM-revision axioms more concisely for propositional logic (where K is supposed to be a single formula, too):

- (AGM' *1) $K * A$ implies A .
- (AGM' *2) If $K \wedge A$ is satisfiable, then $K * A \equiv K \wedge A$.
- (AGM' *3) If A is satisfiable, then $K * A$ is also satisfiable.
- (AGM' *4) If $K_1 \equiv K_2$ and $A_1 \equiv A_2$, then $K_1 * A_1 \equiv K_2 * A_2$.
- (AGM' *5) $(K * A) \wedge B$ implies $K * (A \wedge B)$.
- (AGM' *6) If $(K * A) \wedge B$ is satisfiable, then $K * (A \wedge B)$ implies $(K * A) \wedge B$.

The third important belief change operation presented in [Gär88] is *contraction* (operator: $-$) dealing with the mere deletion of beliefs.

AGM-postulates for contraction:

- (AGM -1) $K - A$ is a belief set.
- (AGM -2) $K - A \subseteq K$.
- (AGM -3) if $A \notin K$ then $K - A = K$
- (AGM -4) if $\text{not } \vdash A$ then $A \notin K - A$.
- (AGM -5) if $A \in K$ then $K \subseteq (K - A) + A$.
- (AGM -6) if A and B are logically equivalent, then $K - A = K - B$.
- (AGM -7) $(K - A) \cap (K - B) \subseteq K - A \wedge B$.
- (AGM -8) if $A \notin K - A \wedge B$ then $K - A \wedge B \subseteq K - A$.

Following Levi [Lev77], expansion and contraction are more fundamental than revision. As he sees it, revising by A means first contracting $\neg A$ and then adding consistently belief in A ; this is formalized by the so-called *Levi identity*

$$K * A = (K - \neg A) + A \quad (2.4)$$

In this way, revision can be expressed by contraction and expansion. Conversely, a contraction operation in terms of revision is given by the so-called *Harper identity*

$$K - A = K \cap (K * \neg A) \quad (2.5)$$

(cf. [Har77]).

Another interesting aspect of belief change was considered by Katsuno and Mendelzon in [KM91b]. They argued that the AGM-type revision was only adequate to describe a revision of knowledge about a *static* world, but

not for recording changes in an *evolving* world. They called this latter type of change *update* (operator: \diamond), with *erasure* being its inverse operation. Their axioms for belief update, stated below, presuppose the knowledge base K to be representable by a single propositional formula, here also denoted by K :

KM-postulates for updating:

(KM $\diamond 1$) $K \diamond A$ implies A .

(KM $\diamond 2$) If $K \models A$ then $K \diamond A = K$.

(KM $\diamond 3$) If K and A are satisfiable, then $K \diamond A$ is also satisfiable.

(KM $\diamond 4$) If $K_1 \equiv K_2$ and $A_1 \equiv A_2$ then $K_1 \diamond A_1 \equiv K_2 \diamond A_2$.

(KM $\diamond 5$) $(K \diamond A) \wedge B$ implies $K \diamond (A \wedge B)$.

(KM $\diamond 6$) If $K \diamond A_1$ implies A_2 and $K \diamond A_2$ implies A_1 then $K \diamond A_1 \equiv K \diamond A_2$.

(KM $\diamond 7$) If K is complete then $(K \diamond A_1) \wedge (K \diamond A_2)$ implies $K \diamond (A_1 \vee A_2)$.

(KM $\diamond 8$) $(K_1 \vee K_2) \diamond A \equiv (K_1 \diamond A) \vee (K_2 \diamond A)$.

For a thorough discussion of these postulates and a comparison to AGM-revision, cf. [KM91b].

An important representation result characterizing the AGM-revision was established in [KM91a]. It made use of *faithful assignments* which assigns to each propositional formula K a pre-order \leq_K over the set of worlds (or interpretations, respectively) Ω such that the following three conditions are satisfied:

1. If $\omega, \omega' \in \text{Mod}(K)$ then $\omega <_K \omega'$ does not hold;
2. if $\omega \in \text{Mod}(K)$ and $\omega' \notin \text{Mod}(K)$ then $\omega <_K \omega'$;
3. if $K_1 \equiv K_2$ then $\leq_{K_1} = \leq_{K_2}$;

where $\omega <_K \omega'$ means $\omega \leq_K \omega'$ and not $\omega' \leq_K \omega$.

Thus faithful assignments identify the models of K with the \leq_K -minimal worlds without making any differences between these models.

Theorem 2.2.2 (Representation theorem [KM91a]). *The revision operator $*$ defines an AGM-revision in the sense of the postulates (AGM *1) - (AGM *8) iff there exists a faithful assignment that maps each propositional formula K to a total pre-order \leq_K such that*

$$\text{Mod}(K * A) = \min_{\leq_K}(\text{Mod}(A)) \quad (2.6)$$

A similar representation result for update was established in [KM91b].

2.3 Nonmonotonic Reasoning and Belief Revision - Two Sides of the Same Coin?

This question was raised and investigated by Gärdenfors and Makinson in [MG91] and [Gär92]. The similarities between nonmonotonic reasoning and belief revision are evident – both fields are concerned with uncertain reasoning, and the relationship between them is *prima facie* very plausibly established by

$$A \sim B \quad \text{iff} \quad T * A \models B \quad (2.7)$$

where T is some (classical) theory. In particular, each belief revision operation gives rise to a nonmonotonic inference operation (or relation). Gärdenfors and Makinson show in [MG91] how postulates from one field translate via (2.7) into properties of the other. But – although the relationship is close, it is not perfect, due to essentially different focusses: In belief revision, the theory T (cf. (2.7)) which is to be revised, is of central concern, whereas in nonmonotonic reasoning, it is not mentioned at all! This flaw may be partially remedied by considering the relation \sim as describing inferences based on some fixed background knowledge T so that (2.7) now reads

$$A \sim_T B \quad \text{iff} \quad T * A \models B$$

This idea of explicitly representing background knowledge, however, is not dealt with in nonmonotonic reasoning. So belief revision should be considered the more general approach, handling the full *dynamics of belief* more thoroughly.

2.4 Iterated Revision, Epistemic States, and Conditionals

Classical belief revision takes *belief sets*, i.e. deductively closed sets of propositional formulas, or classical consequences of one propositional formula, respectively, as basic representations of knowledge. These sets, however, are only poor reflections of the complex attitudes an individual may hold. The limitation to propositional beliefs severely restricts the frame of the AGM theory, in particular, when iterated revision has to be performed (cf. [DP97a, Bou93, BG93, Sch91]). Instead of belief sets, *epistemic states*, Ψ , should be considered as representations of the cognitive state of some intelligent agent at a given time.¹

¹ An interesting approach to iterated revisions of *belief sets* was proposed quite recently by Lehmann, Magidor and Schlechta in [LMS01]. They base their revision operations upon a formal notion of distance.

Formally, epistemic states may be described in many different ways. A very simple representation of an epistemic state is obtained by focusing on the corresponding belief set, that is, on the propositions believed to be true in that epistemic state, and by means of *simple conditional functions* (cf. [Spo88]) recording its epistemological changes. Here the belief set is thought of as being represented by one propositional formula, the so-called *net content* of the epistemic state (cf. [Spo88]). Identifying propositions with subsets of the set Ω of possible worlds, a simple conditional function f assigns to all non-empty subsets $A \subseteq \Omega$ a non-empty subset of Ω such that

$$f(A) \subseteq A; \quad (2.8)$$

$$\text{if } f(A) \cap B \neq \emptyset \text{ then } f(A \cap B) = f(A) \cap B. \quad (2.9)$$

$f(A)$ is to be interpreted as the net content of the changed epistemic state when accepting the new information A to be true. Simple conditional functions correspond to partitions of the set of worlds in equally plausible sets of worlds (cf. [Spo88]) and thus to plausibility pre-orderings. The notion of simple conditional functions was generalized to *ordinal conditional functions* in [Spo88] to achieve a more adequate representation of epistemological attitudes.

Furthermore, probability distributions are generally considered as a particularly sophisticated means for representing epistemic states ([Gär88, Spo88]). Another numerical approach to epistemic states is provided by *possibility distributions* ([DP91c, DP94, DLP94]).

The crucial difference between belief sets and epistemic states is that besides the set of propositional beliefs, $Bel(\Psi) \subseteq \mathcal{L}$, the individual accepts for certain, an epistemic state Ψ also contains the *revision policies* the individual entertains at that time (cf. [Bou93, BG93]). These revision policies reflect the (propositional) beliefs, B , the individual is inclined to hold if new information, A , becomes obvious. They are adequately represented by conditionals $(B|A)$, i.e. expressions of the form “*If A then B*”, conjoining two propositional formulas A and B for a plausible conclusion. So the conditional $(B|A)$ is accepted in the epistemic state Ψ iff revising Ψ by A yields belief in B . This defines a fundamental relationship between conditionals and the process of revision, known as the *Ramsey test* ([BG93, Gär88, Ram50]):

$$\Psi \models (B|A) \quad \text{iff} \quad Bel(\Psi \star A) \models B \quad (2.10)$$

where \star is a revision operator, taking an epistemic state Ψ and some new belief A as inputs and yielding a revised epistemic state $\Psi \star A$ as output. So in the context of revision, a subjunctive meaning of conditionals fits particularly well, in accordance with the Ramsey test: *If A were true, B would be believed*, implicitly referring to a revision of the actual epistemic state by A .

Hence epistemic states are intimately related with belief revision as well as with conditionals, and so are the axioms of each of the three domains; e.g. the properties (2.8) and (2.9) stated above for simple conditional functions are essentially equivalent to the AGM-postulate (AGM *2) and to a conjunction of (AGM *7) and (AGM *8), respectively. This implies also a close connection between epistemic states and conditionals on one hand and nonmonotonic reasoning on the other. For instance, simple conditional functions may be regarded as so-called *choice functions* which Schlechta bases his *Nonmonotonic Logics* [Sch97] upon. Furthermore, it is easy to show that property (2.9) for simple conditional functions implies the condition

$$\text{if } B \subseteq A \text{ then } f(A) \cap B \subseteq f(B)$$

which is crucial for characterizing minimal preferential structures (cf. [Sch97, p. 13]). The relationships between these different areas have been studied in several papers (see, for instance, [KS91, DP91a, Gra91]).

Extending belief revision to an operation on epistemic states instead of belief sets opens up the framework to investigate *iterated revision* in a fully dynamic system of belief change. As a first step towards this aim, Darwiche and Pearl generalized the AGM-revision postulates for revising epistemic states by conditional beliefs (cf. [DP97a]):

AGM-Postulates for revising epistemic states [DP97a]

Suppose Ψ, Ψ_1, Ψ_2 to be epistemic states and $A, A_1, A_2, B \in \mathcal{L}$;

- (AGM* *1) A is believed in $\Psi \star A$: $Bel(\Psi \star A) \models A$.
- (AGM* *2) If $Bel(\Psi) \wedge A$ is satisfiable, then $Bel(\Psi \star A) \equiv Bel(\Psi) \wedge A$.
- (AGM* *3) If A is satisfiable, then $Bel(\Psi \star A)$ is also satisfiable.
- (AGM* *4) If $\Psi_1 = \Psi_2$ and $A_1 \equiv A_2$, then $Bel(\Psi_1 \star A_1) \equiv Bel(\Psi_2 \star A_2)$.
- (AGM* *5) $Bel(\Psi \star A) \wedge B$ implies $Bel(\Psi \star (A \wedge B))$.
- (AGM* *6) If $Bel(\Psi \star A) \wedge B$ is satisfiable then $Bel(\Psi \star (A \wedge B))$ implies $Bel(\Psi \star A) \wedge B$.

Considered superficially, these postulates are exact reformulations of the AGM-postulates with “belief sets” replaced by “belief sets of epistemic states”. So the postulates above ensure that the revision of epistemic states is in line with the AGM-theory as long as the revision of the corresponding belief sets is considered. The most important new aspect in contrast to propositional belief revision is given by postulate (AGM* *4): Only *identical* epistemic states are supposed to yield equivalent revised belief sets. This is a clear but adequate weakening of the corresponding AGM-postulate (AGM *4) which

only would require the belief sets of Ψ_1 and Ψ_2 to be *equivalent*. This would amount to reducing the revision of epistemic states to propositional belief revision which is inappropriate since two different epistemic states Ψ_1, Ψ_2 may have equivalent belief sets $Bel(\Psi_1) \equiv Bel(\Psi_2)$. Thus an epistemic state is not described uniquely by its belief set, and revising Ψ_1 and Ψ_2 with equivalent belief sets by new information A may result in different revised belief sets $Bel(\Psi_1 \star A) \not\equiv Bel(\Psi_2 \star A)$, as the following example illustrates.

Example 2.4.1. Two physicians have to make a diagnosis when confronted with a patient showing certain symptoms. They both agree that disease A is by far the most adequate diagnosis, so they both hold belief in A . Moreover, as the physicians know, diseases B and C might also cause the symptoms, but here the experts disagree: One physician regards B to be a possible diagnosis, too, but excludes C , thus accepting the conditionals $(B|\neg A)$ and $(\neg C|\neg A)$, whereas the other physician is inclined to take C into consideration, but not B , so holding belief in $(\neg B|\neg A)$ and $(C|\neg A)$.

Suppose now that a specific blood test definitely proves that the patient is not suffering from disease A . So both experts have to change their beliefs, the first physician now takes B to be the correct diagnosis, the second one takes C for granted. Though initially the physicians' opinions may be described by the same belief set, $\{A\}$, they end up with different belief sets after revision.

It is important to note that Gärdenfors' famous triviality result [Gär88] complaining the incompatibility of the Ramsey test with some of the AGM-postulates does not hold if conditional beliefs are considered essentially different from propositional beliefs, as is emphasized here and elsewhere (see, for instance, [DP97a, Lev88]; cf. also [Lew76]). Therefore, obeying the difference between $Bel(\Psi_1) \equiv Bel(\Psi_2)$ and $\Psi_1 = \Psi_2$ makes the Ramsey test compatible with the AGM-theory for propositional belief revision: Whereas $Bel(\Psi_1) \equiv Bel(\Psi_2)$ only means that both epistemic states have equivalent belief sets, $\Psi_1 = \Psi_2$ requires the two epistemic states to be identical, i.e. to incorporate in particular the same conditional beliefs as well as the same propositional beliefs.

Darwiche and Pearl [DP97a] proved a representation theorem for their postulates above which parallels the corresponding theorem in AGM-theory (cf. [KM91a]), using a generalized notion of *faithful assignments*:

Definition 2.4.1. A faithful assignment (for epistemic states) is a function that maps each epistemic state Ψ to a total pre-order \leq_Ψ on the worlds Ω satisfying the following conditions:

- (1) $\omega_1, \omega_2 \models Bel(\Psi)$ only if $\omega_1 =_\Psi \omega_2$;

(2) $\omega_1 \models \text{Bel}(\Psi)$ and $\omega_2 \not\models \text{Bel}(\Psi)$ only if $\omega_1 <_\Psi \omega_2$;

for worlds $\omega_1, \omega_2 \in \Omega$ and epistemic states Ψ .

As usual, $\omega_1 <_\Psi \omega_2$ means $\omega_1 \leq_\Psi \omega_2$ and $\omega_2 \not\leq_\Psi \omega_1$; $\omega_1 =_\Psi \omega_2$ iff $\omega_1 \leq_\Psi \omega_2$ and $\omega_2 \leq_\Psi \omega_1$.

Theorem 2.4.1 ([DP97a]). *A revision operator \star satisfies postulates (AGM* *1) – (AGM* *6) iff there exists a faithful assignment that maps each epistemic state Ψ to a total pre-order \leq_Ψ such that*

$$\text{Mod}(\Psi \star A) = \min(A; \Psi) := \min_{\leq_\Psi}(\text{Mod}(A))$$

where $\text{Mod}(\Psi) := \text{Mod}(\text{Bel}(\Psi))$, i.e. the worlds satisfying $\text{Bel}(\Psi \star A)$ are precisely those worlds satisfying A that are minimal with respect to \leq_Ψ .

This theorem shows an important connection between the pre-ordering \leq_Ψ associated with an epistemic state Ψ and the process of revising Ψ by propositional beliefs. \leq_Ψ may be thought of as a *plausibility (pre-)ordering* (or *ranking*, respectively) providing a representation of the epistemic state, that is, as a total pre-order on the set of worlds satisfying conditions (1)-(2) of Definition 2.4.1 and the so-called *smoothness condition*

$$\min(A; \Psi) \neq \emptyset \quad \text{for any satisfiable } A \in \mathcal{L} \quad (2.11)$$

([BG93]), and such that $\text{Mod}(\Psi) = \min_{\leq_\Psi}(\Omega)$. The smoothness condition is also called *limit assumption*, see [Gro88, Lew73]. Such epistemic states (Ψ, \leq_Ψ) correspond to Boutilier's *revision models*, as described in [BG93]. Because we assume the numbers of possible worlds to be finite, the smoothness condition is trivially fulfilled. A more general approach to nonmonotonic reasoning and belief revision makes use of an indexed set of possible worlds, thus allowing infinitely many possible worlds (cf. [KLM90, Sch97, LMS01]). In [Bou94], Boutilier considers revision based on pre-orders without requiring the limit assumption. Other approaches to epistemic states and belief revision (or nonmonotonic reasoning, respectively) are accomplished by considering *system of spheres* [Gro88], *epistemic entrenchment orderings* [Gär88] and *expectation orderings* [GM94]. Friedman and Halpern [FH99] emphasize the need to clarify the *ontology* underlying a belief change process. They present an approach to model dynamic revisions of epistemic states.

Using the Ramsey test (2.10), Theorem 2.4.1 immediately yields

Lemma 2.4.1. *A conditional $(B|A)$ is accepted in an epistemic state (Ψ, \leq_Ψ) iff all minimal A -worlds satisfy B , i.e.*

$$\Psi \models (B|A) \quad \text{iff} \quad \min(A; \Psi) \subseteq \text{Mod}(B)$$

Thus the pre-order \leq_Ψ encodes the conditional beliefs held in Ψ . For two propositional formulas A, B , we define

$$A \leq_\Psi B$$

iff for all $\omega \in \min(A; \Psi)$, $\omega' \in \min(B; \Psi)$, we have $\omega \leq_\Psi \omega'$, i.e. iff the minimal A -worlds are at least as plausible as the minimal B -worlds. \leq_Ψ is a plausibility relation in the sense of Grove (cf. [Gro88, Bou94]), hence dual to an epistemic entrenchment relation² (see [Gär88]). Using this, the lemma above may be reformulated as

Lemma 2.4.2. *A conditional $(B|A)$ is accepted in an epistemic state (Ψ, \leq_Ψ) iff $AB <_\Psi A\bar{B}$.*

In particular, we have (see Definition 1.3.1)

Corollary 2.4.1. *Let $\omega, \omega' \in \Omega$ be two different worlds, let (Ψ, \leq_Ψ) be a representation of an epistemic state.*

$$\Psi \models (\text{form}(\omega) | \text{form}(\omega, \omega')) \text{ iff } \omega <_\Psi \omega'.$$

Belief revision of an epistemic state, however, should not only deal with the revision of propositional beliefs but also with the modification of the revision strategies maintained in that state ([DP97a, Bou93, BG93]). Therefore, taking these revision strategies as conditionals, revision of epistemic states should be concerned with changes in conditional beliefs and, the other way around, with the preservation of conditional beliefs.

Investigating iterated revision, Darwiche and Pearl [DP97a] explicitly took conditional beliefs into account, and they advanced four postulates in addition to the AGM axioms to model what may be called *conditional preservation* under revision by propositional beliefs:

DP-postulates for conditional preservation:

- (C1) If $C \models B$ then $\Psi \models (D | C)$ iff $\Psi \star B \models (D | C)$.
- (C2) If $C \models \bar{B}$ then $\Psi \models (D | C)$ iff $\Psi \star B \models (D | C)$.
- (C3) If $\Psi \models (B | A)$ then $\Psi \star B \models (B | A)$.
- (C4) If $\Psi \star B \models (\bar{B} | A)$ then $\Psi \models (\bar{B} | A)$.

For discussion of these postulates, see the original paper [DP97a].

² Rott [Rot91] investigates belief revision and conditionals within the framework of epistemic entrenchment relations.

2.5 Probabilistic Reasoning – The ME-Approach

Though probability distributions have long been appreciated as high-quality representations of epistemic states, probabilistic belief change has not really been a major topic of the field. This is mainly due to the following reasons:

- Probabilistic beliefs are much more difficult to deal with than propositional beliefs.
- Probabilistic deduction is complicated, and unfortunately, in many interesting cases quite meaningless (cf. [Nil86, TGK91]).
- Even in probabilistic belief change, the paradigm has generally been to establish beliefs for certain, i.e. with probability 1 ([Gär88, DP94]). Two revision operations have been proposed for this purpose: (Bayesian) conditioning and imaging (cf. [Gär88]).

A thorough treatment of uncertain probabilistic beliefs has not yet taken place within the area of belief revision theory, most of the work in describing and classifying belief change operations has been done for belief sets based on classical logics. Gärdenfors [Gär88] and Dubois and Prade [DP94] developed some axioms for probabilistic belief change, being mostly concerned with revising in the sense of establishing facts for certain. While Gärdenfors, however, claimed that conditioning corresponds to expansion (and thus to a certain case of revision) in a probabilistic framework (cf. [Gär88, pp. 105 ff]), Dubois and Prade emphasize that conditioning is not revising but *focusing*, i.e. applying generic or background knowledge to the reference class describing properly the case under consideration. Paris [Par94] and Voorbraak [Voo96a] also consider probabilistic belief revision in the case of uncertain evidences. An important approach to *default probabilistic reasoning* was developed by Adams [Ada75]. His ϵ -*semantics* satisfies some basic inference schemes for nonmonotonic reasoning and proved to yield proper representations of *system P-inferences* (cf. [KLM90]; see also [Pea89]). For a brief overview of qualitative probabilistic reasoning, see [Gol94].

Nevertheless, probabilistic reasoning and probabilistic belief change has been investigated from different point of views in the long tradition of probability theory (for surveys, see, for instance, [Pea88, Par94]).

As a generalization of standard Bayesian conditioning, Jeffrey conditionalization allows one to modify a probability distribution so as to incorporate a changed propositional probability ([Voo96a, Jef83, Par94]): Let P denote a probability distribution, and let $A \in \mathcal{L}$ be a proposition such that $P(A) \neq 0$. Suppose we learn that the probability of A has changed to x : $P'(A) = x$. Then a revised probability function, P' , taking into account the new information while obviously being related to the prior P is given by *Jeffrey's rule*:

$$P'(B) = xP(B|A) + (1 - x)P(B|\neg A) \quad (2.12)$$

for all propositions $B \in \mathcal{L}$. Note that for $x = 1$, we obtain Bayesian conditioning. The rationale behind Jeffrey's rule was to preserve conditional degrees of belief with respect to A and $\neg A$, i.e. the Jeffrey conditionalization as given in (2.12) satisfies

$$P'(B|A) = P(B|A) \quad \text{and} \quad P'(B|\neg A) = P(B|\neg A)$$

for all $B \in \mathcal{L}$. So, many years before nonmonotonic reasoning and belief revision became important topics in Artificial Intelligence, not only had the issue of belief change been addressed in probability theory, but also the necessity of *conditional preservation* when revising epistemic states had been perceived.

Jeffrey conditionalization, however, is only capable of dealing with one factual uncertain evidence. It is not apt to manage changes in conditional probabilities, nor to adopt a *set* of uncertain facts simultaneously (cf. [PV92]).

A powerful tool to realize general changes in probabilistic beliefs has long been available: the *principle of minimum cross-entropy*. The *entropy*

$$H(P) = - \sum_{\omega \in \Omega} P(\omega) \log P(\omega)$$

(with the convention $0 \log 0 = 0$) of a distribution P first appeared as a physical quantity in statistical mechanics and was later interpreted by Shannon as an information-theoretic measure of the uncertainty inherent to P (see [SW76]; for a historical review, cf. [Jay83a]). It is generalized by the notion of *cross-entropy*

$$R(Q, P) = \sum_{\omega \in \Omega} Q(\omega) \log \frac{Q(\omega)}{P(\omega)}$$

(with $0 \log \frac{0}{0} = 0$ and $Q(\omega) \log \frac{Q(\omega)}{0} = \infty$ for $Q(\omega) \neq 0$) between two distributions Q and P . If P_0 denotes the uniform distribution $P_0(\omega) = 1/m$ for all worlds ω , then

$$R(Q, P_0) = -H(Q) + \log m$$

relates absolute and relative entropy.

Cross-entropy is a well-known information-theoretic measure of dissimilarity between two distributions and has been studied extensively (see, for instance, [Csi75, HHJ92, Jay83a, Kul68]; for a brief, but informative introduction and further references, cf. [Sho86]; see also [SJ81]). Cross-entropy is also called *directed divergence* since it lacks symmetry, i.e. $R(Q, P)$ and $R(P, Q)$ differ in general, so it is not a metric. But cross-entropy is *positive*, that means we have $R(Q, P) \geq 0$, and $R(Q, P) = 0$ iff $Q = P$ (cf. [Csi75, HHJ92, Sho86]).

Consider the probabilistic belief revision problem

- (*_{prob}) Given a (prior) distribution P and some set of probabilistic conditionals $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\} \subseteq (\mathcal{L} | \mathcal{L})^{prob}$, how should P be modified to yield a (posterior) distribution P^* with $P^* \models \mathcal{R}$?

When solving (*_{prob}), the paradigm of *informational economy*, i.e. of minimal loss of information (see [Gär88, p. 49]), is realized in an intuitive way by following the *principle of minimum cross-entropy*

$$\min R(Q, P) = \sum_{\omega \in \Omega} Q(\omega) \log \frac{Q(\omega)}{P(\omega)} \quad (2.13)$$

s.t. Q is a probability distribution with $Q \models \mathcal{R}$

For a distribution P and some set \mathcal{R} of probabilistic conditionals compatible with P (cf. [Csi75] or Definition 5.1.1 for the details) there is a (unique) distribution $P_{ME} = P_{ME}(P, \mathcal{R})$ that fulfills \mathcal{R} and has minimal relative entropy to the prior P (cf. [Csi75]), i.e. P_{ME} solves (2.13) and thereby (*_{prob}). Note that (*_{prob}) exceeds the framework of the classical AGM-theory with regard to several aspects: an *epistemic state* (P) is to be revised by a *set of conditionals* representing *uncertain knowledge*.

Maximizing (absolute) entropy under some given constraints \mathcal{R} is equivalent to minimizing cross-entropy with respect to the uniform distribution, given \mathcal{R} . Therefore, the principle of minimum cross-entropy can be regarded as more general than the *principle of maximum entropy*

$$\max H(Q) = - \sum_{\omega} Q(\omega) \log Q(\omega) \quad (2.14)$$

s.t. Q is a probability distribution with $Q \models \mathcal{R}$.

which solves the problem of representing \mathcal{R} by a probability distribution without adding information unnecessarily. We refer to both principles as the *ME-principle*, where the abbreviation *ME* stands both for *Minimum cross-Entropy* and for *Maximum Entropy*.

So, if \mathcal{R} is a set of conditionals, each associated with a probability, then the “best” distribution to represent \mathcal{R} is the one which fulfills all conditionals in \mathcal{R} and has maximum entropy. By an analogous argument, if prior knowledge given by a distribution P has to be adjusted to new probabilistic knowledge \mathcal{R} , the one distribution should be chosen that satisfies \mathcal{R} and has minimum relative entropy to P .

To justify the ME-principles, a couple of authors have demonstrated their usefulness from points of views outside of information theory. So, optimizing entropy is known to yield best expectation values in statistics (cf.

[GHK94, Jay83a]). Two other papers [PV90, SJ80] are concerned with characterizing the ME-principles as logically consistent inference methods, too. Shore and Johnson [JS83, SJ80], succeeded in proving (cross-)entropy to be the only functional the optimization of which satisfies four (resp. five) fundamental axioms of probabilistic inference. A similar result is attained for entropy in [PV90] by Paris and Vencovská without assuming that inference is performed by optimizing a functional, rather basing their characterization on seven postulates. In a further paper [PV92], they justify the principle of minimum cross-entropy as an appropriate method to update a probability function by new uncertain evidence. Independence and invariance properties are in the first place among the properties these authors used for characterizing ME-reasoning. This justifies ME-inference as an inference procedure of minimal changes, but little was said about the nature or the extent of changes actually occurring under ME-adjustment.

Applying the ME-principles means to use an appropriate notion of distance for choosing a “best” representation of new beliefs. In this context, it is interesting to mention the paper of Lehmann, Magidor and Schlechta [LMS01] supposing propositional revisions to be obtained from a formal (pseudo-)distance.

3. Conditionals

This chapter is dedicated to conditionals as objects of crucial concern for knowledge representation, plausible reasoning and belief revision. The relationship between conditionals, epistemic states and beliefs is studied, and we develop the formal means to handle conditionals in revision and reasoning. In particular, we explain how *conditional structures*, imposed by conditionals on worlds, can be represented appropriately to investigate interrelated effects of conditionals.

Parts of the ideas to be developed here can also be found in other papers (see, for instance, [KI98a, KI99c]), but most of the results presented in this chapter are new.

3.1 Conditionals and Epistemic States

Conditionals ($B|A$) represent statements of the form “*If A then B*” conjoining two propositional formulas A , the *antecedent* or *premise*, and B , the *consequent*. A lot of different approaches to a *logic of conditionals* have been made (see, for instance, [Nut80, Lew73, DGC94, GGNR91]), also aiming at reflecting more general relationships between antecedent and consequent so as to capture the manifold meanings of commonsense conditionals. In general, conditionals are used to describe plausible relationships between antecedent and consequent. Besides such qualitative approaches, the validity of conditionals may be quantified by *degrees of certainty* (see [Cal91]). Cox [Cox46] argued that a logically consistent handling of quantified conditionals is only possible within a probabilistic framework, where the degree of certainty associated with a conditional is interpreted as a conditional probability (which should not be confused with assigning a probability to the conditional as a logical sentence, cf. [Lew76]; for a rigorous discussion of Cox’s Theorem, cf. [Hal99a, Hal99b]).

Conditionals may be given a lot of different interpretations, for instance, as counterfactuals, as indicative, subjunctive or normative conditionals etc. (see [Nut80, Bou94]). Independently of its given meaning, however, a condi-

tional $(B|A)$ can be represented as a *generalized indicator function* on worlds, setting

$$(B|A)(\omega) = \begin{cases} 1 & : \omega \models AB \\ 0 & : \omega \models A\overline{B} \\ u & : \omega \models \overline{A} \end{cases} \quad (3.1)$$

where u stands for *undefined* (cf. [DeF74, Cal91]). Two conditionals are considered *equivalent* iff the corresponding indicator functions are identical, i.e.

$$(B|A) \equiv (D|C) \quad \text{iff} \quad A \equiv C \text{ and } AB \equiv CD$$

(cf. [Cal91])

This definition captures excellently the three-valued, thus non-classical character of conditionals. According to it, a conditional $(B|A)$ is a function that polarizes AB and $A\overline{B}$, leaving \overline{A} untouched. Each possible world ω either confirms $(B|A)$, or refutes it, or is of no relevance for it. So conditionals are evaluated with respect to worlds, but considering only single, isolated worlds is not enough to decide if a conditional (as an entity) is accepted or not. To validate conditionals, we need richer epistemic structures than plain propositional interpretations, at least to compare different worlds with regard to their relevance for a conditional (see, for example, [Nut80, Bou94, DP97a]).

An epistemic notion that turned out to be of great importance for conditionals as well as for belief revision is that of *plausibility*: conditionals are supposed to represent plausible conclusions, and plausibility relations on formulas or worlds, respectively, guide AGM-revisions (cf. [Nut80, KM91a, DP97a]; for investigating the deep connection between conditional logic and belief revision theory, see, for instance, [Gro88, FH94]; cf. also Section 2.4, in particular the Ramsey test (2.10)). So we assume epistemic states Ψ appropriate to study conditionals and belief revision to be at least equipped with a plausibility pre-ordering (ranking) \leq_Ψ on worlds (cf. Section 2.4).

There are others, more sophisticated methods to represent epistemic attitudes; among these appreciated representation forms for epistemic states are *probability functions*, *ordinal conditional functions*, *OCFs* and *possibility distributions* ([Gär88, Spo88, DP97a, DP94]):

Definition 3.1.1. A probability function (or probability distribution) is a map

$$P : \Omega \rightarrow [0, 1]$$

such that

$$\sum_{\omega \in \Omega} P(\omega) = 1$$

Each probability function obviously induces a probability measure on 2^Ω and vice versa, and to each propositional formula $A \in \mathcal{L}$, a probability may

be assigned by setting

$$P(A) = \sum_{\omega \models A} P(\omega)$$

The set of propositional formulas with probability 1 constitutes the belief set of P :

$$Bel(P) = \{A \in \mathcal{L} \mid P(A) = 1\}$$

Probabilities of conditionals are defined via conditional probabilities

$$P(B|A) = \frac{P(AB)}{P(A)}$$

for $P(A) \neq 0$.

Probability distributions are generally considered to be most adequate representations of epistemic states ([Gär88, Spo88]). They use the full scope of real numbers between 0 and 1 to specify knowledge, but they are also burdened with a lot of numbers. As a qualitative abstraction of probability functions, Spohn [Spo88] introduced ordinal conditional functions:

Definition 3.1.2. Ordinal conditional functions (OCF's, ranking functions) are functions κ from worlds to ordinals such that some worlds are mapped to the minimal element 0.

Throughout this book, we will simply assume that OCF's are functions from the set of worlds to the natural numbers, extended by 0 and ∞ :

$$\kappa : \Omega \rightarrow \mathbb{N} \cup \{0, \infty\},$$

where ∞ corresponds to the ordinal ω_0 . Ordinal conditional functions not only induce a plausibility pre-order on the set Ω of worlds by $\omega_1 \leq_{\kappa} \omega_2$ iff $\kappa(\omega_1) \leq \kappa(\omega_2)$, but furthermore, they specify non-negative integers as degrees of plausibility – or, more precisely, as degrees of *disbelief* – to worlds. The smaller $\kappa(\omega)$ is, the more plausible the world ω appears, and what is believed (for certain) in the epistemic state represented by κ is described precisely by the set $Mod(\kappa) := \{\omega \in \Omega \mid \kappa(\omega) = 0\}$, and consequently,

$$Bel(\kappa) = \{A \in \mathcal{L} \mid \omega \models A \text{ for all } \omega \in Mod(\kappa)\}$$

For propositional formulas $A, B \in \mathcal{L}$, we set

$$\kappa(A) = \min\{\kappa(\omega) \mid \omega \models A\},$$

so that

$$\kappa(A \vee B) = \min\{\kappa(A), \kappa(B)\}.$$

In particular, $0 = \min\{\kappa(A), \kappa(\overline{A})\}$, so that at least one of A or \overline{A} is considered mostly plausible. A proposition A is believed iff $\kappa(\overline{A}) > 0$ (which implies $\kappa(A) = 0$), so that A is believed iff $Bel(\kappa) \models A$. We abbreviate this by $\kappa \models A$. Therefore, ordinal conditional functions not only allow us to compare worlds according to their plausibility but also to take the relative distances between them into account. So they can be considered as a refinement of the concept of simple conditional functions (cf. [Spo88]; see Section 2.4). For the connections between ordinal conditional functions and qualitative probabilistic reasoning, cf. [Spo88, DP97a, GP96].

A conditional may be assigned a degree of plausibility via

$$\kappa(B|A) = \kappa(AB) - \kappa(A) = \begin{cases} 0 & : \kappa(AB) \leq \kappa(\overline{A}) \\ \kappa(AB) - \kappa(\overline{A}) & : \kappa(AB) \geq \kappa(\overline{A}) \end{cases}$$

Plausibility relations are dual to epistemic entrenchment relations ([Gro88, Bou94]) which are essentially equivalent to qualitative necessity relations ([DP91c, Hoc99]). The quantitative counterpart of necessity relations are necessity measures ([Dub86]) the dual of which are *possibility measures*, i.e. functions $\Pi : \mathcal{L} \rightarrow [0, 1]$ observing logical equivalence and such that $\Pi(\top) = 1$, $\Pi(\perp) = 0$ and

$$\Pi(A \vee B) = \max(\Pi(A), \Pi(B)) \quad (3.2)$$

(cf. [DLP94]). Each possibility measure Π is determined by its values on possible worlds,

$$\Pi(A) = \max_{\omega \models A} \Pi(\{\omega\}),$$

due to (3.2). Functions $\pi : \Omega \rightarrow [0, 1]$, $\pi(\omega) = \Pi(\{\omega\})$, are also called *possibility distributions*. A *conditional possibility* may be defined by setting

$$\Pi(B|A) = \frac{\Pi(AB)}{\Pi(A)} \quad (3.3)$$

for $\Pi(A) \neq 0$ (cf. [DP94]). For the relationships between possibility theory and both probability theory and ordinal conditional functions, see [DP94].

If \mathcal{V} is the set of atomic propositions under consideration, then let $\mathcal{E}_{\mathcal{V}}^*$ be the set of all representations of epistemic states over \mathcal{V} of a certain type. For example, within a probabilistic framework, $\mathcal{E}_{\mathcal{V}}^*$ would be the set of all probability distributions on \mathcal{V} . As soon as the type of representation for epistemic states is fixed, we will not distinguish between an epistemic state and its representation by an element of $\mathcal{E}_{\mathcal{V}}^*$. Furthermore, two epistemic states will be considered *equivalent* if they are represented by the same element in $\mathcal{E}_{\mathcal{V}}^*$. So the type of representation chosen is assumed to reflect all relevant

knowledge held in an epistemic state. Note that an epistemic state is supposed to be a state of equilibrium which contains all explicitly stated as well as all implicitly derived knowledge ([Gär88]). This then demands a complete representation of epistemological knowledge.

To each type of representation of epistemic states, we choose a language $(\mathcal{L} \mid \mathcal{L})^*$ that allows a suitable notation of conditionals which is in accordance with the intended representation of epistemic states. Therefore, for instance, for probability distributions and ordinal conditional functions, we take $(\mathcal{L} \mid \mathcal{L})^* = (\mathcal{L} \mid \mathcal{L})^{prob}$ and $(\mathcal{L} \mid \mathcal{L})^* = (\mathcal{L} \mid \mathcal{L})^{OCF}$, respectively, where

$$\begin{aligned} (\mathcal{L} \mid \mathcal{L})^{prob} &= \{(B|A)[x] \mid A, B \in \mathcal{L}, x \in [0, 1]\} \\ \text{and } (\mathcal{L} \mid \mathcal{L})^{OCF} &= \{(B|A)[n] \mid A, B \in \mathcal{L}, n \in \mathbb{N} \cup \{0, \infty\}\}. \end{aligned}$$

In a purely qualitative setting, $(\mathcal{L} \mid \mathcal{L})^* = (\mathcal{L} \mid \mathcal{L})$ seems to be appropriate. In correspondence to classical definitions, we set

$$Th^*(\Psi) = \{\phi \in (\mathcal{L} \mid \mathcal{L})^* \mid \Psi \models \phi\}$$

to denote all conditionals accepted in the epistemic state Ψ . Here the acceptance relation \models between epistemic states in $\mathcal{E}_{\mathcal{V}}^*$ and conditionals has to be specified appropriately. We will deal with this in detail in Section 3.3.

We further assume (*uniqueness assumption*) that $Th^*(\Psi)$ describes Ψ uniquely (up to representation equivalence):

$$Th^*(\Phi) = Th^*(\Psi) \quad \text{iff} \quad \Phi \equiv \Psi \text{ in } \mathcal{E}_{\mathcal{V}}^* \quad (3.4)$$

This holds for all the representation types mentioned above, which is easy to see in a quantitative environment, and for plausibility pre-orderings \leq_{Ψ} , we have $\omega <_{\Psi} \omega'$ iff $(\omega|form(\omega, \omega')) \in Th^*(\Psi)$. In general, this assumption is justified taking the view that an epistemic state is describable as a response scheme to changes in belief and by observing the Ramsey test (cf. Section 2.4).

3.2 Conditional Valuation Functions

What is common to probability functions, ordinal conditional functions, and possibility measures is, that they make use of *two different operations* to handle both purely propositional information and conditionals adequately. Therefore, we will introduce the abstract notion of a conditional valuation function to reveal more clearly and uniformly the way in which knowledge may be represented and treated within epistemic states. As an adequate structure, we assume an algebra $\mathcal{A} = (\mathcal{A}, \oplus, \odot, 0^{\mathcal{A}}, 1^{\mathcal{A}})$ of real numbers to be equipped with two operations, \oplus and \odot , such that

- (\mathcal{A}, \oplus) is an associative and commutative structure with neutral element $0^{\mathcal{A}}$;
- $(\mathcal{A} - \{0^{\mathcal{A}}\}, \odot)$ is a commutative group with neutral element $1^{\mathcal{A}}$;
- the rule of distributivity holds, i.e.

$$x \odot (y \oplus z) = (x \odot y) \oplus (x \odot z)$$

for $x, y, z \in \mathcal{A}$;

- \mathcal{A} is totally ordered by $\leq_{\mathcal{A}}$ which is compatible with \oplus and \odot in that

$$x \leq_{\mathcal{A}} y \quad \text{implies} \quad x \oplus z \leq_{\mathcal{A}} y \oplus z \quad (3.5)$$

$$x \leq_{\mathcal{A}} y \quad \text{implies} \quad x \odot z \leq_{\mathcal{A}} y \odot z \quad (3.6)$$

hold for all $x, y, z \in \mathcal{A}$.

So \mathcal{A} is nearly an ordered field, except that the elements of \mathcal{A} need not be invertible with respect to \oplus .

Definition 3.2.1. A conditional valuation function is a function

$$V : \mathcal{L} \rightarrow \mathcal{A}$$

from the set \mathcal{L} of formulas into the algebra \mathcal{A} satisfying the following conditions:

1. $V(\perp) = 0^{\mathcal{A}}, V(\top) = 1^{\mathcal{A}}$, and for exclusive formulas A, B (i.e. $AB \equiv \perp$), we have

$$V(A \vee B) = V(A) \oplus V(B);$$

2. for each conditional $(B|A) \in (\mathcal{L} | \mathcal{L})$ with $V(A) \neq 0^{\mathcal{A}}$,

$$V(B|A) = V(AB) \odot V(A)^{-1}, \quad (3.7)$$

where $V(A)^{-1}$ is the \odot -inverse element of $V(A)$ in \mathcal{A} .

Conditional valuation functions assign degrees of certainty, plausibility, possibility etc. to propositional formulas and to conditionals. Making use of two operations, they provide a framework for considering and treating conditional knowledge as fundamentally different from propositional knowledge, a point that is stressed by various authors and that seems to be indispensable for representing epistemic states adequately (cf. [DP97a]). There is, however, no deep conflict between these two different kinds of knowledge (that would be unintuitive, even disastrous) – conditionals should rather be regarded as extending propositional knowledge by a new dimension, and facts may be considered as conditionals of a degenerate form by identifying A with $(A | \top)$. Note

that conditional valuation functions also take this compatibility into account: According to Definition 3.2.1, we have $V(A|\top) = V(A) \odot (1^V)^{-1} = V(A)$.

For each conditional valuation function V , we have

$$V(A) = \sum_{\omega \models A}^{\oplus} V(\omega)$$

so V is determined uniquely by its values on interpretations or on possible worlds, respectively, and we will also write $V : \Omega \rightarrow \mathcal{A}$. For all $A \in \mathcal{A}$, we have $0^{\mathcal{A}} \leq_{\mathcal{A}} V(A) \leq_{\mathcal{A}} 1^{\mathcal{A}}$.

Plausibility relations, as defined in [Gro88] (see also [Bou94]), are most appropriate for representing epistemic states and conditionals qualitatively. But such relations do not really fit the numerical framework of probability theory. So we need a more general notion to integrate conditional valuation functions in the field of nonmonotonic reasoning and belief revision. It is easy to see that any such function $V : \mathcal{L} \rightarrow \mathcal{A}$ is a *plausibility measure*, in the sense of Friedman and Halpern, ([FH96, Fre98]), that is, it fulfills $V(\perp) \leq_{\mathcal{A}} V(A)$ for all $A \in \mathcal{L}$, and $A \models B$ implies $V(A) \leq_{\mathcal{A}} V(B)$.

Two notions which are well-known from probability theory may be generalized for conditional valuation functions:

Definition 3.2.2. *A conditional valuation function V is said to be uniform if $V(\omega) = V(\omega')$ for all worlds ω, ω' .*

So uniform conditional valuation function assign the same degree of plausibility to each world.

Definition 3.2.3. *Let V be a conditional valuation function, let $A, B, C \in \mathcal{L}$ such that $V(C) \neq 0^{\mathcal{A}}$. A and B are called conditionally independent given C (with respect to V) if $V(A|BC) = V(A|C)$.*

Some important examples will help to illustrate the newly introduced notion of a conditional valuation function:

Example 3.2.1. Each probability function P may be taken as a conditional valuation function

$$P : \Omega \rightarrow (\mathbb{R}^+, +, \cdot, 0, 1)$$

where \mathbb{R}^+ denotes the set of all non-negative real numbers. Conversely, each conditional valuation function $V : \Omega \rightarrow (\mathbb{R}^+, +, \cdot, 0, 1)$ is a probability function.

Similarly, each ordinal conditional function κ is a conditional valuation function

$$\kappa : \Omega \rightarrow (\mathbb{Z} \cup \{\infty\}, \min, +, \infty, 0)$$

where \mathbb{Z} denotes the set of all integers, and any *possibility measure* Π can be regarded as a conditional valuation function

$$\Pi : \Omega \rightarrow (\mathbb{R}^+, \max, \cdot, 0, 1).$$

Conditional valuation functions not only provide an abstract means to quantify epistemological attitudes. Their extended ranges allow us to calculate and compare arbitrary proportions of values attached to single worlds. This will prove quite useful to handle complex conditional interrelationships.

Probability functions and ordinal conditional functions will serve as our standard examples for conditional valuation functions.

3.3 Conditional Valuation Functions and Beliefs

By means of a conditional valuation function $V : \mathcal{L} \rightarrow \mathcal{A}$, we are able to validate propositional as well as conditional beliefs. We may say, for instance, that proposition A is *believed in* V , $V \models A$, iff $V(A) = 1^{\mathcal{A}}$, or that the conditional $(B|A)$ is *valid* or *accepted in* V , $V \models (B|A)$, iff $V(A) \neq 0^{\mathcal{A}}$ and $V(A\overline{B}) <_{\mathcal{A}} V(AB)$, i.e. iff AB is more plausible (probable, possible etc.) than $A\overline{B}$.

If $V = P$ is a probability function, then saying that A is believed is usually associated with $P(A) = 1$. On the other hand, the qualitative statement “the conditional $(B|A)$ is valid in P ” might be accepted iff $P(A\overline{B}) < P(AB)$. But both these points of views prove to be not compatible with one another when considering propositional beliefs, A , as degenerate conditional beliefs $(A|\top)$.

If $V = \kappa$ is an ordinal conditional function, then stating that $\kappa(A) = 0$ is not sufficient for establishing belief in A – we must have $\kappa(\overline{A}) > 0$, which is equivalent to stating $\kappa(A) < \kappa(\overline{A})$ and thus compatible to the acceptance of A as a degenerate conditional belief.

In a quantitative framework, it is possible – and sometimes appears to be more adequate (see, for instance, [Spo88, FH99]) – to specify knowledge more exactly by making use of *quantified facts* $A[x]$, and *quantified conditionals* $(B|A)[x]$, respectively, where x is an element of \mathcal{A} .

In accordance with the remarks above and with generally agreed points of views, we say that a conditional $(B|A)[x]$ with $x \in [0, 1]$ is *valid in* or *accepted by* a probability function P ,

$$P \models (B|A)[x] \quad \text{iff} \quad P(A) > 0 \text{ and } P(B|A) = x. \quad (3.8)$$

In particular, we obtain for facts $A[x]$, $x \in [0, 1]$,

$$P \models A[x] \quad \text{iff} \quad P(A) = x. \quad (3.9)$$

Observing that ordinal conditional functions κ specify degrees of disbelief, a conditional $(B|A)[n]$ with $n \in \mathbb{N} \cup \{\infty\}$ is *valid in* or *accepted by* κ ,

$$\kappa \models (B|A)[n] \quad \text{iff} \quad \kappa(\overline{B}|A) = n. \quad (3.10)$$

For facts $A[n], n \in \mathbb{N} \cup \{\infty\}$, this means

$$\kappa \models A[n] \quad \text{iff} \quad \kappa(\overline{A}) = n. \quad (3.11)$$

This is exactly what Spohn [Spo88, p. 117] calls “*A is believed in κ with firmness n* ”.

Note that the exact specification of which beliefs and conditionals are accepted depends upon the type of the valuation function used. We will return to this issue later in the context of revision (see Section 4.3).

For probability functions as well as for ordinal conditional functions, it is possible to use the notion of validity more vaguely by using inequalities instead of equalities (see, for example, the system Z^+ by Goldszmidt and Pearl [GP96]). Here we prefer a crisp validity not only for formal reasons (in fact, stronger results hold e.g. for the optimum entropy approach in the case of equality constraints, cf. [SJ81]). If we follow consequently the argumentation of Spohn [Spo88] that not only do qualitative relations matter in representing epistemic attitudes but also distances between degrees of plausibility, then this same argument should apply to conditionals, too, and we should consider assigning numerical degrees of disbelief to conditionals as meaningful. Similarly for probability functions, statements such as “a conditional $(B|A)$ holds with probability 0.8” or “a conditional $(B|A)$ holds with probability 0.99”, respectively, are usually taken as two different pieces of information. Nevertheless, using inequality constraints may be very important and useful in modelling vague or incomplete knowledge.

Let $(\mathcal{L} \mid \mathcal{L})^{\mathcal{A}}$ be the set of all conditionals quantified by elements of \mathcal{A} :

$$(\mathcal{L} \mid \mathcal{L})^{\mathcal{A}} = \{(B|A)[x] \mid (B|A) \in (\mathcal{L} \mid \mathcal{L}), x \in \mathcal{A}\}$$

Sets of quantified conditionals will sometimes be marked by a suitable superscript, for instance, by $\mathcal{R}^{\mathcal{A}}$ in general, or by $\mathcal{R}^{prob} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$, or by $\mathcal{R}^{OCF} \subseteq (\mathcal{L} \mid \mathcal{L})^{OCF}$ etc. in corresponding settings. Quantified conditionals in $(\mathcal{L} \mid \mathcal{L})^{prob}$ will be called *probabilistic conditionals*, and those in $(\mathcal{L} \mid \mathcal{L})^{OCF}$ will be called *OCF-conditionals*.

3.4 A Dynamic View on Conditionals

Independently of its associated meaning – indicative, subjunctive, counterfactual, etc. –, a conditional $(B|A)$ is an object of a three-valued nature, partitioning the set of worlds Ω in three parts: those worlds satisfying $A \wedge B$ and thus confirming the conditional, those worlds satisfying $A \wedge \neg B$, thus refuting the conditional, and those worlds not fulfilling the premise A and so which the conditional may not be applied to at all (cf. representation (3.1), p. 28). So we define the *affirmative set* and the *conflicting set* of the conditional $(B|A)$ to be

$$\begin{aligned}(B|A)^+ &:= \{\omega \in \Omega \mid \omega \models AB\} = \text{Mod}(AB), \\ (B|A)^- &:= \{\omega \in \Omega \mid \omega \models A\bar{B}\} = \text{Mod}(A\bar{B}),\end{aligned}$$

respectively. $\text{Mod}(\bar{A})$ is called the *neutral set* of $(B|A)$. Each of these sets may be empty. If $(B|A)^+ = \emptyset$, $(B|A)$ is called *contradictory*, if $(B|A)^- = \emptyset$, $(B|A)$ is called *tautological*, and if $\text{Mod}(\bar{A}) = \emptyset$, $(B|A)$ is called a *fact*.

Example 3.4.1. $(\bar{A}|A)$ is a contradictory conditional, $(A|A)$ is tautological and $(A|\top)$ is a fact.

The following lemma is only a reformulation of equivalence by using the introduced notions of affirmative and conflicting sets:

Lemma 3.4.1. *Two conditionals $(B|A), (D|C)$ are equivalent iff their corresponding affirmative and conflicting sets are equal, i.e.*

$$(B|A) \equiv (D|C) \quad \text{iff} \quad (B|A)^+ = (D|C)^+ \text{ and } (B|A)^- = (D|C)^-.$$

Definition 3.4.1. *A conditional $(D|C)$ is called a subconditional of $(B|A)$, written as*

$$(D|C) \sqsubseteq (B|A)$$

iff

$$(D|C)^+ \subseteq (B|A)^+ \quad \text{and} \quad (D|C)^- \subseteq (B|A)^-$$

Thus $(D|C) \sqsubseteq (B|A)$ if the effect of the former conditional on worlds is in line with the latter one, but $(D|C)$ applies to fewer worlds. The \sqsubseteq -relation may be expressed using the standard ordering \leq between propositional formulas: $A \leq B$ iff $A \models B$, i.e. iff $\text{Mod}(A) \subseteq \text{Mod}(B)$:

Lemma 3.4.2. *Let $(B|A), (D|C) \in (\mathcal{L} \mid \mathcal{L})$. Then $(D|C)$ is a subconditional of $(B|A)$, $(D|C) \sqsubseteq (B|A)$, iff $CD \leq AB$ and $C\bar{D} \leq A\bar{B}$; in particular, if $(D|C) \sqsubseteq (B|A)$ then $C \leq A$.*

The proof of this lemma is straightforward, and also the following lemmata are easily proved:

Lemma 3.4.3. *For any two conditionals $(B|A)$ and $(D|C)$,*

$$(B|A) \equiv (D|C) \quad \text{iff} \quad (B|A) \sqsubseteq (D|C) \text{ and } (D|C) \sqsubseteq (B|A).$$

Lemma 3.4.4. *The relation \sqsubseteq defines a pre-ordering on $(\mathcal{L} | \mathcal{L})$, i.e. a reflexive and transitive relation, and induces an ordering on the set of equivalence classes $(\mathcal{L} | \mathcal{L})/\equiv$ with minimal element $(\top|\perp)$ and maximal elements $(A|\top)$, $A \in \mathcal{L}$.*

So we have a unique minimal element (an equivalence class of conditionals, respectively) in $(\mathcal{L} | \mathcal{L})$, namely $(\top|\perp)$, whereas all the facts are maximal elements in $(\mathcal{L} | \mathcal{L})$.

Therefore for any two conditionals $(B|A), (D|C) \in (\mathcal{L} | \mathcal{L})$, the infimum in $(\mathcal{L} | \mathcal{L})$ with respect to \sqsubseteq

$$(B|A) \sqcap (D|C) := \inf\{(B|A), (D|C)\}$$

exists. The supremum of both conditionals, however,

$$(B|A) \sqcup (D|C) := \sup\{(B|A), (D|C)\}$$

only exists if

$$ABC\overline{D} \equiv A\overline{B}CD \equiv \perp \quad (3.12)$$

holds, for otherwise $(B|A)^+ \cup (D|C)^+$ and $(B|A)^- \cup (D|C)^-$ would not be disjoint.

Lemma 3.4.5. *The supremum $(B|A) \sqcup (D|C)$ of two conditionals $(B|A)$ and $(D|C)$ exists iff there is a conditional $(F|E)$ such that both $(B|A)$ and $(D|C)$ are subconditionals of $(F|E)$.*

In particular, each non-tautological and non-contradictory conditional can be represented as the supremum of its basic subconditionals:

Definition 3.4.2. *Basic conditionals are conditionals of the form*

$$\psi_{\omega, \omega'} = (\text{form}(\omega) | \text{form}(\omega, \omega')) \quad (3.13)$$

for any two worlds $\omega, \omega' \in \Omega$ (cf. Definition 1.3.1, p. 9).

Lemma 3.4.6. *Let $(B|A) \in (\mathcal{L} | \mathcal{L})$ be a non-tautological and non-contradictory conditional. Then*

$$(B|A) = \bigsqcup_{\psi_{\omega, \omega'} \sqsubseteq (B|A)} \psi_{\omega, \omega'} \quad (3.14)$$

where $\psi_{\omega, \omega'}$ is defined by (3.13).

Lemma 3.4.7. *For any two non-tautological and non-contradictory conditionals $(B|A), (D|C)$, it holds that $(D|C) \sqsubseteq (B|A)$ iff $\psi_{\omega, \omega'} \sqsubseteq (B|A)$ for all basic subconditionals $\psi_{\omega, \omega'}$ of $(D|C)$.*

We omit the straightforward proof of the foregoing lemma and the technical proofs of Lemma 3.4.8 and Proposition 3.4.1.

Lemma 3.4.8. *Let $(B|A), (D|C) \in (\mathcal{L} | \mathcal{L})$. Then*

$$(B|A) \sqcap (D|C) \equiv (BD | AC(BD \vee \overline{B}\overline{D})), \quad (3.15)$$

and if $ABCD \equiv \overline{A}\overline{B}\overline{C}\overline{D} \equiv \perp$, then

$$(B|A) \sqcup (D|C) \equiv (AB \vee CD | A \vee C). \quad (3.16)$$

Note that even if $ABCD \equiv \overline{A}\overline{B}\overline{C}\overline{D} \equiv \perp$ does not hold, (3.16) defines an operation on conditionals which, however, does not coincide with the supremum of the corresponding conditionals in this case.

Example 3.4.2. Consider the (non-tautological) conditionals $(B|A)$ and $(\overline{B}|A)$ violating condition (3.12). A blind application of (3.16) yields $(B|A) \sqcup (\overline{B}|A) \equiv (A|A)$, but $(A|A)^- = \emptyset$, so $(A|A)$ cannot be the supremum of $(B|A)$ and $(\overline{B}|A)$.

Proposition 3.4.1. *\sqcup and \sqcap , as defined by (3.16) and (3.15), are associative and distributive operations on $(\mathcal{L} | \mathcal{L})$, if all required suprema exist.*

Though $(\mathcal{L} | \mathcal{L})$, equipped with \sqcup and \sqcap , has some convenient algebraic properties, it is not a lattice, because the supremum $(B|A) \sqcup (D|C)$ is not defined for arbitrary conditionals $(B|A)$ and $(D|C)$. In particular, it is not a Boolean algebra. But this should not be considered a disadvantage. The relation \sqsubseteq is not aiming at managing the logical properties of conditionals but at capturing their dynamic effects on worlds. (That is exactly the reason why the application of \sqcup fails in the example above.) And this effect on worlds is a crucial factor in the framework of (conditional) belief revision. Establishing a conditional $(B|A)$ within an epistemic state Ψ means shifting (some) worlds in $(B|A)^+$ and $(B|A)^-$ appropriately. Therefore the changes brought about by conditional revision depends on a world's being in one of the sets $(B|A)^+, (B|A)^-$ and $Mod(\overline{A})$.

We will now introduce another relation between conditionals that is quite opposite to the subconditional relation and so describes another extreme of possible interaction:

Definition 3.4.3. Suppose $(B|A), (D|C) \in (\mathcal{L} | \mathcal{L})$ are two conditionals. $(D|C)$ is called perpendicular to $(B|A)$,

$$(D|C) \perp\!\!\!\perp (B|A)$$

iff either $\text{Mod}(C) \subseteq (B|A)^+$, or $\text{Mod}(C) \subseteq (B|A)^-$, or $\text{Mod}(C) \subseteq \text{Mod}(\bar{A})$.

The perpendicularity relation symbolizes a kind of *irrelevance* of one conditional for another one. We have $(D|C) \perp\!\!\!\perp (B|A)$ if $\text{Mod}(C)$, i.e. the range of application of the conditional $(D|C)$, is completely contained in exactly one of the sets $(B|A)^+, (B|A)^-$ or $\text{Mod}(\bar{A})$. So for all worlds which $(D|C)$ may be applied to, $(B|A)$ has the same effect and yields no further partitioning. Note, that $\perp\!\!\!\perp$ is not a symmetric relation; $(D|C) \perp\!\!\!\perp (B|A)$ rather expresses that $(D|C)$ is *not affected* by $(B|A)$, or, that $(B|A)$ is *irrelevant* for $(D|C)$.

The following two lemmata provide characteristic properties of the perpendicularity of conditionals (only the second lemma is proved in the Appendix):

Lemma 3.4.9. Let $(B|A), (D|C) \in (\mathcal{L} | \mathcal{L})$. Then $(D|C) \perp\!\!\!\perp (B|A)$ iff either $C \leq AB$, or $C \leq A\bar{B}$, or $C \leq \bar{A}$.

Lemma 3.4.10. Let $(B|A), (D|C) \in (\mathcal{L} | \mathcal{L})$ be conditionals, and let $(D|C)$ be neither tautological nor contradictory. Then $(D|C) \perp\!\!\!\perp (B|A)$ iff $\psi_{\omega, \omega'} \perp\!\!\!\perp (B|A)$ for all basic subconditionals $\psi_{\omega, \omega'} \sqsubseteq (D|C)$ of $(D|C)$.

Because of Lemmata 3.4.7 and 3.4.10, one may say that in most cases, the relations \sqsubseteq and $\perp\!\!\!\perp$ may be checked by considering basic subconditionals.

In the following section, we will pursue the idea of conditionals having effects on worlds further. As an interesting generalization, however, we will deal with sets of conditionals instead of regarding only one conditional at a time.

3.5 Conditional Structures

Due to their non-Boolean nature, conditionals are rather complicated objects. In particular, it is not an easy task to handle the relationships between them so as to preserve conditional dependencies “as far as possible” under adaptation. To make the problem clear and to point out a possible way to solve it, we give an example which is taken from [Whi90] and which illustrates a phenomenon also well-known under the name “Simpson’s paradox”.

Example 3.5.1 (Florida murderers). This example is based on a real life investigation. During the six years period 1973-79, about 5000 murder cases were recorded in the US state of Florida, and the following probability distribution P mirrors the sentencing policy in those years (for further references, cf. [Whi90, pp. 46f]). The propositional variables involved are $V = \underline{Victim}$ (of the murder) is black or white, respectively, $v \in \{v_b, v_w\}$, $M = \underline{Murderer}$ is black or white, respectively, $m \in \{m_b, m_w\}$, and $D = \underline{Murderer}$ is sentenced to Death, $d \in \{d, \bar{d}\}$.

ω	$P(\omega)$	ω	$P(\omega)$
$v_w m_w d$	0.0151	$v_w m_w \bar{d}$	0.4353
$v_w m_b d$	0.0101	$v_w m_b \bar{d}$	0.0502
$v_b m_w d$	0	$v_b m_w \bar{d}$	0.0233
$v_b m_b d$	0.0023	$v_b m_b \bar{d}$	0.4637

Thus P implies

$$P \models (d|m_w)[0.0319] \quad \text{and} \quad P \models (d|m_b)[0.0236],$$

so justice seemingly passed sentences without respect of color of skin. Differences, however, become strikingly apparent if the third variable V , revealing the color of skin of the *victim*, is also taken into account:

$$\begin{aligned} P \models (d|v_w m_w)[0.0335], & \quad P \models (d|v_w m_b)[0.1675], \\ P \models (d|v_b m_w)[0], & \quad P \models (d|v_b m_b)[0.0049]. \end{aligned}$$

If, for instance, the probability of the conditional $(d|m_b)[0.0236]$ is to be changed, the probabilities of the conditionals $(d|v\bar{m})$ containing important information should be preserved in an adequate manner.

This last example illustrates a strange but typical behavior that marginal distributions and the conditionals involved may have (we will continue it later on, see Example 4.5.1 in Section 4.5). Let us look upon this problem in a more abstract environment.

Suppose P is a probability distribution on a set of variables containing a, b , and suppose $P \models (b|a)[x]$. In which way may a third variable, c , affect this conditional, i.e. what can be said about the probability of $(b|a\bar{c})$ in P ?

Roughly, there are two possibilities. In the first case, c does not affect $(b|a)[x]$ at all, that is to say, we have $P(b|a\bar{c}) = P(b|a)$, showing b and c to be *conditionally independent given a* (cf. Definition 3.2.3), and c to be irrelevant for $(b|a)[x]$. By a straightforward calculation, we see that $P(b|a\bar{c}) = P(b|a)$ iff $\frac{P(abc)P(a\bar{b}\bar{c})}{P(a\bar{b}c)P(ab\bar{c})} = 1$. In the second, more usual case, we have $P(b|a\bar{c}) \neq P(b|a)$,

and consequently $\frac{P(abc)P(\bar{a}\bar{b}\bar{c})}{P(a\bar{b}c)P(ab\bar{c})} \neq 1$. Thus departures from conditional independence – and thereby the *extent of relevance* – may be measured by the cross product ratio or *interaction quotient* $\frac{P(abc)P(\bar{a}\bar{b}\bar{c})}{P(a\bar{b}c)P(ab\bar{c})}$. A reasonable demand for a posterior distribution P^* adapted to a changed probability of $(b|a)$ then is that posterior interaction should be the same as prior interaction, i.e.

$$\frac{P^*(abc)P^*(\bar{a}\bar{b}\bar{c})}{P^*(a\bar{b}c)P^*(ab\bar{c})} = \frac{P(abc)P(\bar{a}\bar{b}\bar{c})}{P(a\bar{b}c)P(ab\bar{c})}. \quad (3.17)$$

In statistics, logarithms of such expressions are used to measure the interactions between the variables involved (cf. [Goo63, Whi90]).

In the general case, we consider joint influences of groups of variables (instead of one single variable) on a conditional $(B|A)$, and we have to take a *set* of conditionals into account. Thus the notion of (statistical) interaction quotients has to be generalized, involving more worlds both in the numerators and in the denominators and being based appropriately on \mathcal{R} . The comments above give interaction quotients a logical meaning that fits the intention of this treatise better than a statistical interpretation. It offers a suitable way to carry out the necessary generalization from a conditional-logical point of view:

In (3.17), two sets of worlds are related to each other with respect to P and P^* : $\{abc, \bar{a}\bar{b}\bar{c}\}$ in the numerator, and $\{a\bar{b}c, ab\bar{c}\}$ in the denominator. In both sets, the conditional $(b|a)$ is once confirmed (by abc and by $ab\bar{c}$, respectively) and once refuted (by $\bar{a}\bar{b}\bar{c}$ and $a\bar{b}c$, respectively), so both sets show the same behavior with regard to the revising conditional $(b|a)[x]$. This idea of a *behavior* or *structure with respect to* \mathcal{R} can be formalized appropriately by group-theoretical means, as will be developed in the sequel.

When we consider (finite) sets of conditionals $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\} \subseteq (\mathcal{L} | \mathcal{L})$ we have to modify the representation given in (3.1), p. 28, appropriately to identify the effect of each conditional in \mathcal{R} on worlds in Ω . This leads to introducing the functions $\sigma_i = \sigma_{(B_i|A_i)}$ below (see (3.18)) which generalize (3.1) by replacing the numbers 0 and 1 by abstract symbols. Moreover, we will make use of a group structure to represent the joint impact of conditionals on worlds.

To each conditional $(B_i|A_i)$ in \mathcal{R} we associate two symbols $\mathbf{a}_i^+, \mathbf{a}_i^-$. Let

$$\mathcal{F}_{\mathcal{R}} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \dots, \mathbf{a}_n^+, \mathbf{a}_n^- \rangle$$

be the free abelian group with generators $\mathbf{a}_1^+, \mathbf{a}_1^-, \dots, \mathbf{a}_n^+, \mathbf{a}_n^-$, i.e. $\mathcal{F}_{\mathcal{R}}$ consists of all elements of the form $(\mathbf{a}_1^+)^{r_1}(\mathbf{a}_1^-)^{s_1} \dots (\mathbf{a}_n^+)^{r_n}(\mathbf{a}_n^-)^{s_n}$ with integers $r_i, s_i \in \mathbb{Z}$ (the ring of integers). Each element of $\mathcal{F}_{\mathcal{R}}$ can be identified by its

exponents, so that $\mathcal{F}_{\mathcal{R}}$ is isomorphic to \mathbb{Z}^{2n} (cf. [LS77, FR99]). The commutativity of $\mathcal{F}_{\mathcal{R}}$ corresponds to the fact that the conditionals in \mathcal{R} shall be effective simultaneously, without assuming any order of application.

For each $i, 1 \leq i \leq n$, we define a function $\sigma_i : \Omega \rightarrow \mathcal{F}_{\mathcal{R}}$ by setting

$$\sigma_i(\omega) = \begin{cases} \mathbf{a}_i^+ & \text{if } (B_i|A_i)(\omega) = 1 \\ \mathbf{a}_i^- & \text{if } (B_i|A_i)(\omega) = 0 \\ 1 & \text{if } (B_i|A_i)(\omega) = u \end{cases} \quad (3.18)$$

$\sigma_i(\omega)$ represents the manner in which the conditional $(B_i|A_i)$ applies to the possible world ω . The neutral element 1 of $\mathcal{F}_{\mathcal{R}}$ corresponds to the non-applicability of $(B_i|A_i)$ in case that the antecedent A_i is not satisfied. The function

$$\begin{aligned} \sigma_{\mathcal{R}} &= \prod_{1 \leq i \leq n} \sigma_i : \Omega \rightarrow \mathcal{F}_{\mathcal{R}}, \\ \sigma_{\mathcal{R}}(\omega) &= \prod_{1 \leq i \leq n} \sigma_i(\omega) = \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \mathbf{a}_i^+ \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} \mathbf{a}_i^- \end{aligned} \quad (3.19)$$

describes the all-over effect of \mathcal{R} on ω . $\sigma_{\mathcal{R}}(\omega)$ is called *(a representation of) the conditional structure of ω with respect to \mathcal{R}* . For each world ω , $\sigma_{\mathcal{R}}(\omega)$ contains at most one of each \mathbf{a}_i^+ or \mathbf{a}_i^- , but never both of them because each conditional applies to ω in a well-defined way. The next lemma shows that this property characterizes conditional structure functions:

Lemma 3.5.1. *Let $\sigma : \Omega \rightarrow \mathcal{F}$ be a map from the set of worlds Ω to the free abelian group $\mathcal{F} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \dots, \mathbf{a}_n^+, \mathbf{a}_n^- \rangle$ generated by $\mathbf{a}_1^+, \mathbf{a}_1^-, \dots, \mathbf{a}_n^+, \mathbf{a}_n^-$, such that $\sigma(\omega)$ contains at most one of each \mathbf{a}_i^+ or \mathbf{a}_i^- , for each world $\omega \in \Omega$. Then there is a set of conditionals \mathcal{R} with $\text{card}(\mathcal{R}) \leq n$ such that $\sigma = \sigma_{\mathcal{R}}$.*

Example 3.5.2. Let $\mathcal{R} = \{(c|a), (c|b)\}$, where a, b, c are atoms, and let $\mathcal{F}_{\mathcal{R}} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \mathbf{a}_2^+, \mathbf{a}_2^- \rangle$. We associate \mathbf{a}_1^{\pm} with the first conditional, $(c|a)$, and \mathbf{a}_2^{\pm} with the second one, $(c|b)$. The following table shows the values of the function $\sigma_{\mathcal{R}}$ on worlds $\omega \in \Omega$:

ω	$\sigma_{\mathcal{R}}(\omega)$	ω	$\sigma_{\mathcal{R}}(\omega)$
abc	$\mathbf{a}_1^+ \mathbf{a}_2^+$	$\bar{a}bc$	\mathbf{a}_2^+
$ab\bar{c}$	$\mathbf{a}_1^- \mathbf{a}_2^-$	$\bar{a}b\bar{c}$	\mathbf{a}_2^-
$a\bar{b}c$	\mathbf{a}_1^+	$\bar{a}\bar{b}c$	1
$a\bar{b}\bar{c}$	\mathbf{a}_1^-	$\bar{a}\bar{b}\bar{c}$	1

abc confirms both conditionals, so its conditional structure is represented by $\mathbf{a}_1^+ \mathbf{a}_2^+$. This corresponds to the product (in $\mathcal{F}_{\mathcal{R}}$) of the conditional structures of the worlds $\bar{a}bc$ and $a\bar{b}c$. Two worlds, namely $\bar{a}b\bar{c}$ and $a\bar{b}\bar{c}$, are not affected at all by the conditionals in \mathcal{R} .

The next example illustrates that also multiple copies of worlds may be necessary to relate conditional structures:

Example 3.5.3. Consider the set $\mathcal{R} = \{(d|a), (d|b), (d|c)\}$ of conditionals using the atoms a, b, c, d . Let $\mathbf{a}_1^\pm, \mathbf{a}_2^\pm, \mathbf{a}_3^\pm$ be the group generators associated with $(d|a), (d|b), (d|c)$, respectively. Then we have

$$\begin{aligned} \sigma_{\mathcal{R}}(ab\bar{c}d)\sigma_{\mathcal{R}}(\bar{a}bcd)\sigma_{\mathcal{R}}(\bar{a}bcd) &= (\mathbf{a}_1^+\mathbf{a}_2^+)(\mathbf{a}_1^+\mathbf{a}_3^+)(\mathbf{a}_2^+\mathbf{a}_3^+) \\ &= (\mathbf{a}_1^+)^2(\mathbf{a}_2^+)^2(\mathbf{a}_3^+)^2 \\ &= (\mathbf{a}_1^+\mathbf{a}_2^+\mathbf{a}_3^+)^2 \\ &= \sigma_{\mathcal{R}}(abcd)^2 \end{aligned}$$

Here two copies of $abcd$, or of its structure, respectively, are necessary to match the product of the conditional structures of $ab\bar{c}d, \bar{a}bcd$ and $\bar{a}bcd$.

To compare worlds adequately with respect to their conditional structures, we impose a multiplication on the set of worlds Ω by considering the worlds ω as formal symbols. That means, we introduce the free abelian group $\widehat{\Omega}$ generated by all $\omega \in \Omega$

$$\widehat{\Omega} := \langle \omega \mid \omega \in \Omega \rangle \quad (3.20)$$

and consisting of all products

$$\widehat{\omega} = \omega_1^{r_1} \dots \omega_m^{r_m}, \quad \omega_1, \dots, \omega_m \in \Omega, \text{ and } r_1, \dots, r_m \text{ integers.}$$

Now $\sigma_{\mathcal{R}}$ may be extended to $\widehat{\Omega}$ in a straightforward manner by setting

$$\begin{aligned} \sigma_{\mathcal{R}}(\widehat{\omega}) &= \sigma_{\mathcal{R}}(\omega_1^{r_1} \dots \omega_m^{r_m}) \\ &= \sigma_{\mathcal{R}}(\omega_1)^{r_1} \dots \sigma_{\mathcal{R}}(\omega_m)^{r_m}, \end{aligned}$$

yielding a *homomorphism of groups*

$$\sigma_{\mathcal{R}} : \widehat{\Omega} \rightarrow \mathcal{F}_{\mathcal{R}}$$

For $\widehat{\omega} = \omega_1^{r_1} \dots \omega_m^{r_m} \in \widehat{\Omega}$, we obtain the group element

$$\begin{aligned} \sigma_{\mathcal{R}}(\omega_1^{r_1} \dots \omega_m^{r_m}) &= \prod_{1 \leq i \leq n} (\mathbf{a}_i^+)^{\sum_{k: \sigma_i(\omega_k) = \mathbf{a}_i^+} r_k} \cdot \\ &\quad \cdot \prod_{1 \leq i \leq n} (\mathbf{a}_i^-)^{\sum_{k: \sigma_i(\omega_k) = \mathbf{a}_i^-} r_k} \end{aligned} \quad (3.21)$$

as representation of the conditional structure of $\widehat{\omega}$. We will often use fractional representations for the elements of $\widehat{\Omega}$, that is, for instance, we will write $\frac{\omega_1}{\omega_2}$ instead of $\omega_1\omega_2^{-1}$.

Thus the conditional structure of $\widehat{\omega}$ is represented by a group element which is a product of the generators $\mathbf{a}_i^+, \mathbf{a}_i^-$ of $F_{\mathcal{R}}$, with each \mathbf{a}_i^+ occurring with exponent $\sum_{k:\sigma_i(\omega_k)=\mathbf{a}_i^+} r_k = \sum_{k:\omega_k \models A_i B_i} r_k$, and each \mathbf{a}_i^- occurring with exponent $\sum_{k:\sigma_i(\omega_k)=\mathbf{a}_i^-} r_k = \sum_{k:\omega_k \models A_i \bar{B}_i} r_k$ (note that each of the sums may be empty in which case the corresponding conditional cannot be applied to any of the worlds occurring in $\widehat{\omega}$). So the exponent of \mathbf{a}_i^+ in $\sigma_{\mathcal{R}}(\widehat{\omega})$ indicates the number of worlds in $\widehat{\omega}$ which confirm the conditional $(B_i|A_i)$, each world being counted with its multiplicity, and in the same way, the exponent of \mathbf{a}_i^- indicates the number of elementary events that are in conflict with $(B_i|A_i)$. The elements of $\widehat{\Omega}$ replace the multi-sets considered in [KI98a], allowing a more coherent and elegant handling of conditional structures.

In particular, it is possible to isolate the (positive or negative) net impact of one conditional in \mathcal{R} by considering suitable elements of $\widehat{\Omega}$, as the following example illustrates:

Example 3.5.4 (continued). In Example 3.5.2 above, we have

$$\sigma_{\mathcal{R}}\left(\frac{abc}{\bar{a}\bar{b}\bar{c}}\right) = \frac{\mathbf{a}_1^+ \mathbf{a}_2^+}{\mathbf{a}_2^+} = \mathbf{a}_1^+$$

So $\frac{abc}{\bar{a}\bar{b}\bar{c}}$ reveals the positive net impact of the conditional $(c|a)$ within \mathcal{R} , symbolized by \mathbf{a}_1^+ .

Similarly, in Example 3.5.3, the element $\frac{ab\bar{c}\bar{d} \cdot \bar{a}\bar{b}c\bar{d}}{a\bar{b}\bar{c}\bar{d}}$ isolates the negative net impact of the second conditional, $(d|b)$:

$$\sigma_{\mathcal{R}}\left(\frac{ab\bar{c}\bar{d} \cdot \bar{a}\bar{b}c\bar{d}}{a\bar{b}\bar{c}\bar{d}}\right) = \frac{\mathbf{a}_1^- \mathbf{a}_2^- \cdot \mathbf{a}_3^-}{\mathbf{a}_1^- \mathbf{a}_3^-} = \mathbf{a}_2^-.$$

The generators \mathbf{a}_i^+ are mere symbols, representing the effects of the corresponding conditional on worlds. If we choose different symbols $\mathbf{b}_i^+, \mathbf{b}_i^-$ to be associated with the conditionals in \mathcal{R} , we arrive at a different representation homomorphism $\sigma'_{\mathcal{R}} : \widehat{\Omega} \rightarrow \mathcal{F}'_{\mathcal{R}} = \langle \mathbf{b}_1^+, \mathbf{b}_1^-, \dots, \mathbf{b}_n^+, \mathbf{b}_n^- \rangle$. But for all $\widehat{\omega}_1, \widehat{\omega}_2 \in \widehat{\Omega}$, we have

$$\sigma_{\mathcal{R}}(\widehat{\omega}_1) = \sigma_{\mathcal{R}}(\widehat{\omega}_2) \quad \text{iff} \quad \sigma'_{\mathcal{R}}(\widehat{\omega}_1) = \sigma'_{\mathcal{R}}(\widehat{\omega}_2)$$

as can easily be seen from (3.21). This means

$$\ker \sigma_{\mathcal{R}} = \ker \sigma'_{\mathcal{R}}$$

where $\ker \sigma$ denotes the kernel of a homomorphism σ , i.e.

$$\ker \sigma := \{\widehat{\omega} \in \widehat{\Omega} \mid \sigma(\widehat{\omega}) = 1\}$$

Hence, the kernel of such a representation homomorphism does not depend on the symbols chosen as group generators and therefore, it is an invariant of the set of conditionals \mathcal{R} . The kernel $\ker \sigma_{\mathcal{R}}$ contains exactly all group elements $\hat{\omega} \in \hat{\Omega}$ with a balanced conditional structure, that means, where all effects of conditionals in \mathcal{R} on worlds occurring in $\hat{\omega}$ are completely cancelled.

Having the same conditional structure defines an equivalence relation $\equiv_{\mathcal{R}}$ on $\hat{\Omega}$:

$$\hat{\omega}_1 \equiv_{\mathcal{R}} \hat{\omega}_2 \quad \text{iff} \quad \sigma_{\mathcal{R}}(\hat{\omega}_1) = \sigma_{\mathcal{R}}(\hat{\omega}_2).$$

The equivalence classes are in one-to-one correspondence to the elements of the quotient group $\hat{\Omega} / \ker \sigma_{\mathcal{R}} = \{\hat{\omega} \cdot (\ker \sigma_{\mathcal{R}}) \mid \hat{\omega} \in \hat{\Omega}\}$ by observing that

$$\hat{\omega}_1 \equiv_{\mathcal{R}} \hat{\omega}_2 \quad \text{iff} \quad \sigma_{\mathcal{R}}(\hat{\omega}_1 \hat{\omega}_2^{-1}) = 1.$$

Because the kernel $\ker \sigma_{\mathcal{R}}$ is an invariant of \mathcal{R} , $\equiv_{\mathcal{R}}$ does not depend on the chosen representation either. Therefore, the equivalence class

$$[\hat{\omega}]_{\equiv_{\mathcal{R}}} = \{\hat{\omega}' \in \hat{\Omega} \mid \sigma_{\mathcal{R}}(\hat{\omega}') = \sigma_{\mathcal{R}}(\hat{\omega})\}$$

of an element $\hat{\omega} \in \hat{\Omega}$ is called its *conditional structure with respect to \mathcal{R}* . According to (3.21), we have

$$\begin{aligned} [\omega_1^{r_1} \dots \omega_m^{r_m}]_{\equiv_{\mathcal{R}}} &= [\nu_1^{s_1} \dots \nu_p^{s_p}]_{\equiv_{\mathcal{R}}} \quad \text{iff for all } i, 1 \leq i \leq n, \\ \sum_{k: \omega_k \models A_i B_i} r_k &= \sum_{l: \nu_l \models A_i B_i} s_l \quad \text{and} \quad \sum_{k: \omega_k \models A_i \overline{B_i}} r_k = \sum_{l: \nu_l \models A_i \overline{B_i}} s_l. \end{aligned} \quad (3.22)$$

The kernel of $\sigma_{\mathcal{R}}$ plays an important part in identifying the conditional structure of elements $\hat{\omega} \in \hat{\Omega}$, in particular of worlds ω , with respect to the set of conditionals \mathcal{R} . No nontrivial relations hold between different group generators $\mathbf{a}_1^+, \mathbf{a}_1^-, \dots, \mathbf{a}_n^+, \mathbf{a}_n^-$ of $\mathcal{F}_{\mathcal{R}}$, so we have $\sigma_{\mathcal{R}}(\hat{\omega}) = 1$ iff $\sigma_i(\hat{\omega}) = 1$ for all i , $1 \leq i \leq n$, and this means

$$\ker \sigma_{\mathcal{R}} = \bigcap_{i=1}^n \ker \sigma_i \quad (3.23)$$

In this way, each conditional in \mathcal{R} contributes to $\ker \sigma_{\mathcal{R}}$. The kernel of $\sigma_{\mathcal{R}}$, however, is not apt to describe uniquely the set \mathcal{R} of conditionals.

Example 3.5.5. Consider the four sets

$$\begin{aligned} \mathcal{R}_0 &= \{(c|a), (c|b)\} \\ \mathcal{R}_1 &= \{(c|a), (c|b), (c|\bar{a})\} \\ \mathcal{R}_2 &= \{(c|a), (c|b), (c|\bar{a}), (c|\bar{b})\} \\ \mathcal{R}_3 &= \{(c|a), (c|b), (c|\bar{b})\} \end{aligned}$$

Some easy calculations show that

$$\begin{aligned} \ker \sigma_{\mathcal{R}_1} &= \left\langle \frac{abc \cdot \bar{a}\bar{b}\bar{c}}{\bar{a}\bar{b}\bar{c} \cdot \bar{a}\bar{b}\bar{c}}, \frac{ab\bar{c} \cdot \bar{a}\bar{b}\bar{c}}{\bar{a}\bar{b}\bar{c} \cdot \bar{a}\bar{b}\bar{c}} \right\rangle \\ &= \ker \sigma_{\mathcal{R}_2} = \ker \sigma_{\mathcal{R}_3} \end{aligned}$$

and

$$\ker \sigma_{\mathcal{R}_0} = \left\langle \bar{a}\bar{b}\bar{c}, \bar{a}\bar{b}\bar{c}, \frac{abc}{\bar{a}\bar{b}\bar{c} \cdot \bar{a}\bar{b}\bar{c}}, \frac{ab\bar{c}}{\bar{a}\bar{b}\bar{c} \cdot \bar{a}\bar{b}\bar{c}} \right\rangle$$

So we have $\ker \sigma_{\mathcal{R}_1} = \ker \sigma_{\mathcal{R}_2} = \ker \sigma_{\mathcal{R}_3}$ and $\mathcal{R}_1 \cap \mathcal{R}_2 \cap \mathcal{R}_3 = \mathcal{R}_0$, but $\ker \sigma_{\mathcal{R}_0} \neq \ker \sigma_{\mathcal{R}_i}$ for $i = 1, 2, 3$.

In general, if $(B|A)$ is a conditional with $\ker \sigma_{\mathcal{R}} \subseteq \ker \sigma_{(B|A)}$, then $\ker \sigma_{\mathcal{R}} = \ker \sigma_{\mathcal{R} \cup \{(B|A)\}}$, as may easily be seen from (3.23). Moreover, Example 3.5.5 above illustrates that it is not always possible to find a minimal (with respect to set inclusion) set of conditionals \mathcal{R}' with $\sigma_{\mathcal{R}} = \sigma_{\mathcal{R}'}$.

In particular, a conditional $(B|A)$ and its negation $(\bar{B}|A)$ give rise to the same kernel:

Lemma 3.5.2. *Let $(B|A) \in (\mathcal{L} \mid \mathcal{L})$. Then $\ker \sigma_{(B|A)} = \ker \sigma_{(\bar{B}|A)}$.*

This is evident by considering (3.21).

The subgroup

$$\hat{\Omega}_0 = \left\langle \frac{\omega_1}{\omega_2} \mid \omega_1, \omega_2 \in \Omega \right\rangle$$

of $\hat{\Omega}$ generated by all quotients $\frac{\omega_1}{\omega_2}$ is of particular interest in connection with conditionals because it only contains ratios and iterated ratios formed by products. It focusses on comparing actions and interactions of conditionals on worlds. Thus by considering $\hat{\Omega}_0$, it is possible to reveal genuinely conditional influences hidden e.g. by normalizing constraints (see Section 3.6; cf. also Lemma 3.5.3). The elements of $\hat{\Omega}_0$ may be described easily:

$$\hat{\Omega}_0 = \{ \hat{\omega} = \omega_1^{r_1} \cdots \omega_m^{r_m} \in \hat{\Omega} \mid \sum_{j=1}^m r_j = 0 \}.$$

Two elements $\hat{\omega}_1 = \omega_1^{r_1} \cdots \omega_m^{r_m}$, $\hat{\omega}_2 = \nu_1^{s_1} \cdots \nu_p^{s_p} \in \hat{\Omega}$ are equivalent modulo $\hat{\Omega}_0$, i.e. $\hat{\omega}_1 \hat{\Omega}_0 = \hat{\omega}_2 \hat{\Omega}_0$, iff $\sum_{1 \leq j \leq m} r_j = \sum_{1 \leq k \leq p} s_k$. This means that $\hat{\omega}_1$ and $\hat{\omega}_2$ are equivalent modulo $\hat{\Omega}_0$ iff they both are a (cancelled) product of the same number of generators, each generator being counted with its corresponding exponent.

Let

$$\ker_0 \sigma_{\mathcal{R}} := \ker \sigma_{\mathcal{R}} \cap \widehat{\Omega}_0$$

be the part of $\ker \sigma_{\mathcal{R}}$ which is included in $\widehat{\Omega}_0$. $\ker_0 \sigma_{\mathcal{R}}$ is less expressive than $\ker \sigma_{\mathcal{R}}$, for instance, it does not contain all $\omega \in \Omega$ with $\sigma_{\mathcal{R}}(\omega) = 1$. But $\ker_0 \sigma_{\mathcal{R}}$ concentrates on considering ratios as essential entities to reveal the influences of conditionals. The following lemma shows that $\ker_0 \sigma_{\mathcal{R}}$ differs from $\ker \sigma_{\mathcal{R}}$ only by taking the conditional tautology $(\top \mid \top)$ into regard:

Lemma 3.5.3. $\widehat{\Omega}_0 = \ker \sigma_{(\top \mid \top)}$, and $\ker_0 \sigma_{\mathcal{R}} = \ker \sigma_{\mathcal{R} \cup \{(\top \mid \top)\}}$.

So by considering $\ker_0 \sigma_{\mathcal{R}}$, implicit normalizing constraints (such as $P(\top) = 1$ for probability functions or $\kappa(\top) = 0$ for ordinal conditional functions) can be taken explicitly into account.

Finally, we will show how to describe the relations \sqsubseteq and $\perp\!\!\!\perp$ between conditionals, introduced in Definitions 3.4.1 and 3.4.3, respectively, by considering the kernels of the corresponding σ -homomorphisms. As a convenient notation, for each proposition $A \in \mathcal{L}$ we define

$$\widehat{A} = \{\widehat{\omega} = \omega_1^{r_1} \dots \omega_m^{r_m} \in \widehat{\Omega} \mid \omega_i \models A \text{ for all } i, 1 \leq i \leq m\}. \quad (3.24)$$

Proposition 3.5.1. Let $(B \mid A), (D \mid C) \in (\mathcal{L} \mid \mathcal{L})$ be conditionals.

1. $(D \mid C)$ is either a subconditional of $(B \mid A)$ or of $(\overline{B} \mid A)$ iff $C \leq A$ and $\ker \sigma_{(D \mid C)} \cap \widehat{C} \subseteq \ker \sigma_{(B \mid A)} \cap \widehat{C}$.
2. $(D \mid C) \perp\!\!\!\perp (B \mid A)$ iff $\widehat{C} \cap \widehat{\Omega}_0 \subseteq \ker \sigma_{(B \mid A)}$.

Within the framework of representing, revising and discovering conditional knowledge, the notion of *conditional indifference of a conditional valuation function*, to be introduced in the next Section 3.6, will play a central part (cf. Sections 4.5, 4.6 and 8.2). For knowledge discovery, it will be crucial to determine a suitable set \mathcal{R} of conditionals such that $\ker \sigma_{\mathcal{R}} \subseteq \widehat{\Omega}_1$ or $\ker_0 \sigma_{\mathcal{R}} \subseteq \widehat{\Omega}_1$, respectively, for some given subgroup $\widehat{\Omega}_1 \subseteq \widehat{\Omega}$ (ideally, equality should hold; see Section 3.6 and Section 8.2 for the details). The set

$$\begin{aligned} \mathcal{R}(\widehat{\Omega}_1) &= \{(B \mid A) \in (\mathcal{L} \mid \mathcal{L}) \mid \sigma_{(B \mid A)}(\widehat{\omega}) = 1 \text{ for all } \widehat{\omega} \in \widehat{\Omega}_1\} \\ &= \{(B \mid A) \in (\mathcal{L} \mid \mathcal{L}) \mid \widehat{\Omega}_1 \subseteq \ker \sigma_{(B \mid A)}\} \end{aligned}$$

is obviously a maximal candidate for such a set \mathcal{R} . But determining $\mathcal{R}(\widehat{\Omega}_1)$ is quite an expensive task. The next lemma and the following corollary provide a first easy criterion for excluding conditionals from being elements of $\mathcal{R}(\widehat{\Omega}_1)$:

Lemma 3.5.4. *Let $(B|A) \in (\mathcal{L} \mid \mathcal{L})$. For any basic subconditional $\psi_{\omega, \omega'} \sqsubseteq (B|A)$, $\sigma_{(B|A)}\left(\frac{\omega}{\omega'}\right) \neq 1$.*

Corollary 3.5.1. *If $\frac{\omega}{\omega'} \in \ker \sigma_{\mathcal{R}}$, then $\psi_{\omega, \omega'} \not\sqsubseteq (B|A)$ for any $(B|A) \in \mathcal{R}$.*

In general, for a given element $\hat{\omega} \in \hat{\Omega}$, it is easy to find conditionals $(B|A)$ with $\hat{\omega} \in \ker \sigma_{(B|A)}$. Suppose, for instance, $\hat{\omega} = \frac{\omega_1}{\nu_1} \cdots \frac{\omega_r}{\nu_r}$, to be an element in $\hat{\Omega}_0$. Set $\Omega_1 := \{\omega_1, \dots, \omega_r, \nu_1, \dots, \nu_r\}$ and $A := \text{form}(\Omega_1)$, and choose $\Omega_2 \subseteq \Omega_1$ suitably to yield $\sigma_{(B|A)}(\hat{\omega}) = 1$ with $B := \text{form}(\Omega_2)$. The following example will illustrate this.

Example 3.5.6. Let the alphabet consist of the three atoms a, b, c . Consider the element $\hat{\omega} = \frac{abc \cdot \bar{a}\bar{b}c}{\bar{a}\bar{b}c \cdot \bar{a}bc} \in \hat{\Omega}$. Set $A := \text{form}(abc, \bar{a}\bar{b}c, \bar{a}bc, \bar{a}bc) = c$, and choose, for instance, $B_1 = \text{form}(abc, \bar{a}\bar{b}c) = ac$ and $B_2 = \text{form}(abc, \bar{a}bc) = bc$. Then for each of $(B_1|A) \equiv (a|c)$ and $(B_2|A) \equiv (b|c)$, we have $\sigma_{(B_i|A)}(\hat{\omega}) = 1$.

In this way, considering the elements of subgroups $\hat{\Omega}_1$ helps us find appropriate sets of conditionals. But interactions of conditionals considerably complicate this task. In Section 8.2, we will present a method that is apt to discovering conditional structures for a special but important type of conditionals.

3.6 Conditional Indifference

To study conditional interactions, we now focus on the behavior of conditional valuation functions $V : \mathcal{L} \rightarrow \mathcal{A}$ with respect to the “multiplication” \odot in \mathcal{A} (see Definition 3.2.1, p. 32). Each such function may be extended to a homomorphism

$$V : \hat{\Omega}_+ \rightarrow (\mathcal{A}, \odot)$$

by setting

$$V(\omega_1^{r_1} \cdots \omega_m^{r_m}) = V(\omega_1)^{r_1} \odot \cdots \odot V(\omega_m)^{r_m},$$

where $\hat{\Omega}_+$ is the subgroup of $\hat{\Omega}$ generated by the set $\Omega_+ := \{\omega \in \Omega \mid V(\omega) \neq 0^{\mathcal{A}}\}$. This allows us to analyze numerical relationships holding between different $V(\omega)$. Thereby, it will be possible to elaborate the conditionals whose structures V follows, that means, to determine sets of conditionals $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ with respect to which V is *indifferent*:

Definition 3.6.1. Suppose $V : \mathcal{L} \rightarrow \mathcal{A}$ is a conditional valuation function and $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ is a set of conditionals such that $V(A) \neq 0^{\mathcal{A}}$ for all $(B \mid A) \in \mathcal{R}$.

1. V is strictly indifferent with respect to \mathcal{R} iff the following two conditions hold:
 - (i) If $V(\omega) = 0^{\mathcal{A}}$ then there is $(B \mid A) \in \mathcal{R}$ such that $\sigma_{(B \mid A)}(\omega) \neq 1$ and $V(\omega') = 0^{\mathcal{A}}$ for all ω' with $\sigma_{(B \mid A)}(\omega') = \sigma_{(B \mid A)}(\omega)$.
 - (ii) $V(\hat{\omega}_1) = V(\hat{\omega}_2)$ whenever $\sigma_{\mathcal{R}}(\hat{\omega}_1) = \sigma_{\mathcal{R}}(\hat{\omega}_2)$ for $\hat{\omega}_1, \hat{\omega}_2 \in \hat{\Omega}_+$.
2. V is (weakly) indifferent with respect to \mathcal{R} iff V is strictly indifferent with respect to $\mathcal{R} \cup \{(\top \mid \top)\}$.

If V is strictly indifferent with respect to $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$, then it does not distinguish between different elements $\hat{\omega}_1, \hat{\omega}_2$ with the same conditional structure with respect to \mathcal{R} . Conversely, any deviation $V(\hat{\omega}) \neq 1^{\mathcal{A}}$ can be explained by the conditionals in \mathcal{R} acting on $\hat{\omega}$ in a non-balanced way. Weak indifference means that tautologies are taken explicitly into account; it concentrates on ratios which conditionals are based upon. So in many applications, it will appear to be the weaker but more adequate notion. Using Lemma 3.5.3, V is (weakly) indifferent with respect to \mathcal{R} iff (i) holds and $V(\hat{\omega}_1) = V(\hat{\omega}_2)$ whenever $\hat{\omega}_1 \hat{\Omega}_0 = \hat{\omega}_2 \hat{\Omega}_0$ and $\sigma_{\mathcal{R}}(\hat{\omega}_1) = \sigma_{\mathcal{R}}(\hat{\omega}_2)$ for $\hat{\omega}_1, \hat{\omega}_2 \in \hat{\Omega}_+$.

Condition (i) in Definition 3.6.1(1) is necessary to deal with worlds $\omega \notin \Omega_+$.

Lemma 3.6.1. If the conditional valuation function V is (strictly or weakly) indifferent with respect to \mathcal{R} , then $\sigma_{\mathcal{R}}(\omega_1) = \sigma_{\mathcal{R}}(\omega_2)$ implies $V(\omega_1) = V(\omega_2)$ for all worlds $\omega_1, \omega_2 \in \Omega$.

The following proposition rephrases conditional indifference by establishing a relationship between the kernels of $\sigma_{\mathcal{R}}$ and V :

Proposition 3.6.1. Let $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ be a set of conditionals, and let $V : \mathcal{L} \rightarrow \mathcal{A}$ be a conditional valuation function with $V(A) \neq 0^{\mathcal{A}}$ for all $(B \mid A) \in \mathcal{R}$.

1. V is strictly indifferent with respect to \mathcal{R} iff condition (i) of Definition 3.6.1(1) holds, and $\ker \sigma_{\mathcal{R}} \cap \hat{\Omega}_+ \subseteq \ker V$.
2. V is (weakly) indifferent with respect to \mathcal{R} iff condition (i) of Definition 3.6.1(1) holds, and $\ker_0 \sigma_{\mathcal{R}} \cap \hat{\Omega}_+ \subseteq \ker_0 V$.

If, in particular, $\ker \sigma_{\mathcal{R}} \cap \hat{\Omega}_+ = \ker V$, or $\ker_0 \sigma_{\mathcal{R}} \cap \hat{\Omega}_+ = \ker_0 V$, respectively, then $V(\hat{\omega}_1) = V(\hat{\omega}_2)$ if and only if $\sigma_{\mathcal{R}}(\hat{\omega}_1) = \sigma_{\mathcal{R}}(\hat{\omega}_2)$ for $\hat{\omega}_1, \hat{\omega}_2 \in \hat{\Omega}_+$

$\widehat{\Omega}_+$ (and $\widehat{\omega}_1\widehat{\Omega}_0 = \widehat{\omega}_2\widehat{\Omega}_0$). In this case, V completely follows the conditional structures imposed by \mathcal{R} – it observes \mathcal{R} *faithfully*:

Definition 3.6.2. Let $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ be a set of conditionals, and let $V : \mathcal{L} \rightarrow \mathcal{A}$ be a conditional valuation function with $V(A) \neq 0^{\mathcal{A}}$ for all $(B|A) \in \mathcal{R}$.

$V : \mathcal{L} \rightarrow \mathcal{A}$ is said to be strictly faithful (or (weakly) faithful) with respect to \mathcal{R} , iff $\ker_{(0)} \sigma_{\mathcal{R}} \cap \widehat{\Omega}_+ = \ker_{(0)} V$.

We will close this section by characterizing probability functions and ordinal conditional functions with indifference properties:

Theorem 3.6.1. A probability function P is strictly indifferent with respect to a set $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\}$ of conditionals iff $P(A_i) \neq 0$ for all $i, 1 \leq i \leq n$, and there are real numbers $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^- \in \mathbb{R}^+$ such that

$$P(\omega) = \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \quad (3.25)$$

for all $\omega \in \Omega$.

This theorem can be rephrased using group-theoretical terms:

Corollary 3.6.1. A probability function P is strictly indifferent with respect to a set $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\}$ iff $P(A_i) \neq 0$ for all $i, 1 \leq i \leq n$, and there is a homomorphism

$$\tilde{P} : \mathcal{F}_{\mathcal{R}} \rightarrow (\mathbb{R}^+, \cdot)$$

such that

$$\tilde{P} \circ \sigma_{\mathcal{R}} = P \quad (3.26)$$

Weak indifference with respect to \mathcal{R} means strict indifference with respect to $\mathcal{R} \cup \{(\top|\top)\}$; so we obtain immediately

Corollary 3.6.2. A probability function P is (weakly) indifferent with respect to a set $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\}$ iff $P(A_i) \neq 0$ for all $i, 1 \leq i \leq n$, and there are real numbers $\alpha_0, \alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^- \in \mathbb{R}^+, \alpha_0 > 0$, such that

$$P(\omega) = \alpha_0 \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \quad (3.27)$$

for all $\omega \in \Omega$.

Each probability function P has to obey the normalization constraint $P(\top) = 1$ which may veil the (inter)actions of conditionals. So weak indifference appears to be particularly appropriate to study conditional structures in a probabilistic framework.

Similar statements and arguments also hold for ordinal conditional functions and OCF-conditionals. We start with reformulating the central Theorem 3.6.1 for the OCF-environment:

Theorem 3.6.2. *An ordinal conditional function κ is strictly indifferent with respect to a set $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\}$ of conditionals iff $\kappa(A_i) \neq \infty$ for all $i, 1 \leq i \leq n$, and there are rational numbers $\kappa_i^+, \kappa_i^- \in \mathbb{Q}, 1 \leq i \leq n$, such that*

$$\kappa(\omega) = \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} \kappa_i^- \quad (3.28)$$

for all $\omega \in \Omega$.

This theorem and the statements to follow are proven in full analogy to Theorem 3.6.1 and the corollaries above; so we omit the proofs.

The numbers $\kappa_i^+, \kappa_i^- \in \mathbb{Q}, 1 \leq i \leq n$, again may serve to define a suitable homomorphism of groups:

Corollary 3.6.3. *An ordinal conditional function κ is strictly indifferent with respect to a set $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\}$ of conditionals iff $\kappa(A_i) \neq \infty$ for all $i, 1 \leq i \leq n$, and there is a homomorphism*

$$\tilde{\kappa} : \mathcal{F}_{\mathcal{R}} \rightarrow (\mathbb{Z}, +)$$

such that

$$\tilde{\kappa} \circ \sigma_{\mathcal{R}} = \kappa \quad (3.29)$$

For weak indifference, we obtain

Corollary 3.6.4. *An ordinal conditional function κ is (weakly) indifferent with respect to a set $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\}$ of conditionals iff $\kappa(A_i) \neq \infty$ for all $i, 1 \leq i \leq n$, and there are rational numbers $\kappa_0, \kappa_i^+, \kappa_i^- \in \mathbb{Q}, 1 \leq i \leq n$, such that*

$$\kappa(\omega) = \kappa_0 + \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} \kappa_i^- \quad (3.30)$$

for all $\omega \in \Omega$.

Theorems 3.6.1 and 3.6.2, as well as Corollaries 3.6.2 and 3.6.4, give simple criteria to check conditional indifference with probability functions and ordinal conditional functions. Moreover, they provide intelligible schemas to construct conditional indifferent functions.

Note that the concept of conditional indifference is a structural notion, using the numerical values of a conditional valuation function, V , as manifestations of conditional structures imposed by a set \mathcal{R} . We do not postulate, however, that the conditionals in \mathcal{R} are satisfied by V . This adoption of \mathcal{R} will be dealt with in the following chapter, *Revising epistemic states by conditional beliefs*.

4. Revising Epistemic States by Conditional Beliefs

Usually, the *belief sets* in AGM theory (cf. Section 2.2) are assumed to be deductively closed sets of propositional formulas, or to be represented by one single propositional formula, respectively, and the revising beliefs are taken to be propositional formulas. So the AGM postulates constrain revisions of the form $\psi * A$, the revision operator $*$ connecting two propositional formulas ψ and A , where ψ represents the initial state of belief and A stands for the new information. A representation theorem (see [KM91a]) establishes a relationship between AGM revision operators and total pre-orders \leq_ψ on the set of possible worlds, proving the revised belief set $\psi * A$ to be satisfied precisely by all minimal A -worlds (see also Section 2.2).

Belief sets represent what is known for certain and are of specific interest. They are, however, only poor reflections of the complex attitudes an individual may hold. The limitation to propositional beliefs severely restricts the frame of AGM theory, in particular, when iterated revisions have to be performed. So belief revision should not only be concerned with the revision of propositional beliefs but also with the modification of *revision strategies* when new information arrives (cf. [DP97a, Bou93, BG93]). These revision strategies may be given implicitly by some kind of preference relation like a plausibility ordering or an epistemic entrenchment (cf. Section 2.4), or may be taken explicitly as conditional beliefs. Revisions of the complex structure of an epistemic state so as to allow iterated revisions are denoted as *transmutations of knowledge systems* in [Wil94]. As a counterpart to the paradigm of *minimal propositional change* guiding the AGM postulates, the new paradigm of *preserving conditional beliefs*, shortly referred to as *conditional preservation*, arises in the framework of revising epistemic states.

Darwiche and Pearl [DP97a] explicitly took conditional beliefs into account by revising epistemic states instead of belief sets, and they advanced four postulates in addition to the AGM axioms as an approach to describe conditional preservation under revision by propositional beliefs (cf. the *DP-postulates* on page 22 in Section 2.4).

In the sequel, we broaden the framework for revising epistemic states (as presented, for instance, in [DP97a, Bou94, Wil94]) so as to include also the

revision by conditional beliefs. Thus belief revision is considered here in quite a general framework, exceeding the AGM-theory in two respects:

- We revise epistemic states; this makes it necessary to allow for the changes in conditional beliefs caused by new information.
- The new belief A may be of a conditional nature, thus reflecting a changed or newly acquired revision policy that has to be incorporated adequately.

First, we present a scheme of eight postulates appropriate to guide the revision of epistemic states by conditional beliefs (cf. Section 4.1). These postulates are supported mainly by following the specific, non-classical nature of conditionals. The aim of preserving conditional beliefs is achieved by studying specific interactions between conditionals, represented properly by two relations. Because one of the postulates claims propositional belief revision to be a special case of conditional belief revision, our framework also covers the topic of Darwiche and Pearl's work [DP97a], and we show that all four postulates presented there may be derived from our postulates. We state representation theorems for the principal postulates, and to exemplify our ideas, we present a conditional belief operator obeying all of the postulates by using ordinal conditional functions as representations of epistemic states.

Like the postulates of Darwiche and Pearl, our postulates aim at *describing* an appropriate *principle of conditional preservation* to be obeyed when revising epistemic states. The main result of this chapter, however, is a complete formalization of this principle not only with respect to one revising conditional, but to a (finite) set of conditionals to be simultaneously incorporated into the knowledge system (see Section 4.5). We base this principle on representations of epistemic states via conditional valuation functions (cf. Section 3.2) by making use of the notion of conditional indifference (see Section 3.6). As a particular application of these ideas, we deal with the representation of incompletely specified (conditional and factual) knowledge by epistemic states in an appropriate way.

4.1 Postulates for Revising by Conditional Beliefs

Revising an epistemic state Ψ by a conditional $(B|A)$ becomes necessary if a new conditional belief, or a new revision policy, respectively, is to be included in Ψ , yielding a changed epistemic state $\Psi' = \Psi * (B|A)$ such that $\Psi' \models (B|A)$, i.e. $\Psi' * A \models B$. We will use the same operator $*$ for propositional as well as for conditional revision, thus expressing that conditional revision should extend propositional revision in accordance with the Ramsey test (RT) (see page 18).

In this section, we propose a catalogue of postulates a revision of an epistemic state by a conditional should satisfy. The rationale behind our postulates is not to minimize conditional change, as in Boutilier and Goldszmidt's work [BG93], but to preserve the *conditional structure* of the knowledge, as far as possible. Here the key idea is to follow the conditionals in Ψ as long as there is no conflict between them and the new conditional belief. We will make use of the relations \sqsubseteq (subconditional) and $\perp\!\!\!\perp$ (perpendicularity), introduced in 3.4, Definitions 3.4.1 and 3.4.3, to relate conditionals appropriately.

Postulates for conditional revision:

Suppose Ψ is an epistemic state and $(B|A), (D|C)$ are conditionals. Let $\Psi * (B|A)$ denote the result of revising Ψ by $(B|A)$.

- (CR0) $\Psi * (B|A)$ is an epistemic state.
- (CR1) $\Psi * (B|A) \models (B|A)$ (*success*).
- (CR2) $\Psi * (B|A) = \Psi$ iff $\Psi \models (B|A)$ (*stability*).
- (CR3) $\Psi * B := \Psi * (B|\top)$ induces a propositional AGM-revision operator.
- (CR4) $\Psi * (B|A) = \Psi * (D|C)$ whenever $(B|A) \equiv (D|C)$.
- (CR5) If $(D|C) \perp\!\!\!\perp (B|A)$ then $\Psi \models (D|C)$ iff $\Psi * (B|A) \models (D|C)$.
- (CR6) If $(D|C) \sqsubseteq (B|A)$ and $\Psi \models (D|C)$ then $\Psi * (B|A) \models (D|C)$.
- (CR7) If $(D|C) \sqsubseteq (\overline{B}|A)$ and $\Psi * (B|A) \models (D|C)$ then $\Psi \models (D|C)$.

Postulates (CR0) and (CR1) are self-evident. (CR2) postulates that Ψ should be left unchanged precisely if it already entails the conditional. (CR3) says that the induced propositional revision operator should be in accordance with the AGM postulates. (CR4) requires the result of the revision process to be independent of the syntactical representation of conditionals.

The next three postulates aim at preserving the conditional structure of knowledge:

(CR5) claims that revising by a conditional should preserve all conditionals to which that conditional is irrelevant, in the sense described by the relation $\perp\!\!\!\perp$. The rationale behind this postulate is the following: The validity of a conditional $(B|A)$ in an epistemic state Ψ depends on the relation between (some) worlds in $Mod(AB)$ and (some) worlds in $Mod(A\overline{B})$ (see Lemmata 2.4.1, 2.4.2). So incorporating $(B|A)$ into Ψ may require a shift between $Mod(AB)$ on one side and $Mod(A\overline{B})$ on the other side, but should leave intact any relations between worlds within $Mod(AB)$, $Mod(A\overline{B})$, or $Mod(\overline{A})$. These relations may be captured by conditionals $(D|C)$ not affected by $(B|A)$, that is, by conditionals $(D|C) \perp\!\!\!\perp (B|A)$.

(CR6) states that conditional revision should bring about no change for conditionals that are already in line with the revising conditional, and (CR7) guarantees that no conditional change contrary to the revising conditional is caused by conditional revision.

An idea of *conditional preservation* is also inherent to the postulates (C1)-(C4) of Darwiche and Pearl ([DP97a], see also page 22) which we will show to be covered by our postulates.

Theorem 4.1.1. *Suppose $*$ is a conditional revision operator obeying the postulates (CR0)-(CR7). Then for the induced propositional revision operator, postulates (C1)-(C4) are satisfied, too.*

This theorem provides further justifications for the postulates of Darwiche and Pearl from within the framework of conditionals.

4.2 Representation Theorems

To formulate and prove the representation theorems of this section, we assume that each epistemic state is equipped with a plausibility pre-ordering \leq_Ψ underlying propositional revision. Thus we actually presuppose Postulate (CR3), in observing Theorem 2.4.1, page 21.

Postulates (CR5)-(CR7) claim specific connections to hold between Ψ and the revised $\Psi * (B|A)$, thus relating \leq_Ψ and $\leq_{\Psi * (B|A)}$. We will elaborate this relationship in order to characterize those postulates by properties of the pre-orders associated with Ψ and $\Psi * (B|A)$.

Postulate (CR5) proves to be of particular importance because it guarantees the ordering within $Mod(AB)$, $Mod(A\bar{B})$, $Mod(\bar{A})$, respectively, to be preserved:

Theorem 4.2.1. *The conditional revision operator $*$ satisfies (CR5) iff for each epistemic state (Ψ, \leq_Ψ) and for each conditional $(B|A)$ it holds that:*

$$\omega \leq_\Psi \omega' \quad \text{iff} \quad \omega \leq_{\Psi * (B|A)} \omega' \quad (4.1)$$

for all worlds ω, ω' both being elements of $Mod(AB)$ (or of $Mod(A\bar{B})$, or of $Mod(\bar{A})$, respectively).

As an immediate consequence, equation (4.1) yields

Lemma 4.2.1. *Suppose (4.1) holds for all worlds ω, ω' both being elements of $Mod(AB)$ (or of $Mod(A\bar{B})$, or of $Mod(\bar{A})$, respectively). Let $E \in \mathcal{L}$ be a proposition such that either $E \leq AB$, or $E \leq A\bar{B}$, or $E \leq \bar{A}$. Then*

$$\min(E; \Psi) = \min(E; \Psi * (B | A))$$

Together with the Ramsey test (RT) (see page 18), (CR5) yields equalities of belief sets as are stated in the following proposition:

Proposition 4.2.1. *If the conditional revision operator $*$ satisfies Postulate (CR5), then*

$$\begin{aligned} \text{Bel}((\Psi * (B|A)) * AB) &\equiv \text{Bel}(\Psi * AB) \\ \text{Bel}((\Psi * (B|A)) * A\bar{B}) &\equiv \text{Bel}(\Psi * A\bar{B}) \\ \text{Bel}((\Psi * (B|A)) * \bar{A}) &\equiv \text{Bel}(\Psi * \bar{A}) \end{aligned}$$

For the representation theorems of Postulates (C6) and (C7), we need Postulate (CR5), respectively equation (4.1) and its consequence, Lemma 4.2.1: We have to ensure that the property of being a minimal world in the affirmative or in the contradictory set associated with some conditionals is not touched under revision.

Theorem 4.2.2. *Suppose $*$ is a conditional revision operator satisfying (CR5). Let Ψ be an epistemic state, and let $(B|A)$ be a conditional.*

1. *$*$ satisfies (CR6) iff for all $\omega \in \text{Mod}(AB)$, $\omega' \in \text{Mod}(A\bar{B})$, $\omega <_{\Psi} \omega'$ implies $\omega <_{\Psi * (B|A)} \omega'$.*
2. *$*$ satisfies (CR7) iff for all $\omega \in \text{Mod}(AB)$, $\omega' \in \text{Mod}(A\bar{B})$, $\omega' <_{\Psi * (B|A)} \omega$ implies $\omega' <_{\Psi} \omega$.*

4.3 Conditional Valuation Functions and Revision

The theory of revising epistemic states is mainly devised and developed for qualitative representations, such as ordinal conditional functions and the like (see, for instance, [Bou94, DP97a]). The theorems of the previous section are only to be used in such a qualitative framework. Furthermore, revision, propositional beliefs and conditional beliefs may be tightly connected via the Ramsey test (see Section 2.4, equation (2.10), page 18).

Conditional valuation functions $V : \mathcal{L} \rightarrow \mathcal{A}$ were introduced in Section 3.2 as abstract representations of epistemic states, with V providing a measure of plausibility adequate to model belief revisions (see Section 3.3). Conditional valuation function subsume in particular probability functions as well as ordinal conditional functions. Therefore, they allow us to consider revisions in a quantitative as well as in a qualitative framework. It is worth-while studying how conditional valuation functions can be revised in accordance with the Ramsey test within both frameworks.

An ordinal conditional function $\kappa : \Omega \rightarrow \mathbb{N} \cup \{0, \infty\}$ induces a (propositional) AGM-revision operator $*$ by setting

$$Mod(\kappa * A) = \min_{\kappa}(Mod(A)) \quad (4.2)$$

(see Theorem 2.2.2, page 16). The Ramsey test then reads

$$\kappa \models (B|A) \quad \text{iff} \quad \kappa * A \models B. \quad (4.3)$$

This is in accordance with the plausibility relation imposed by κ , as the following lemma shows:

Lemma 4.3.1. *Let $(B|A)$ be a conditional in $(\mathcal{L} | \mathcal{L})$, let κ be an ordinal conditional function. The following three statements are equivalent:*

- (i) $\kappa \models (B|A)$.
- (ii) $\kappa(AB) < \kappa(A\bar{B})$.
- (iii) $\kappa(\bar{B}|A) > 0$.

So κ accepts a conditional (via the Ramsey test) iff AB is more plausible than $A\bar{B}$. The proof of this lemma is immediate by using Lemma 2.4.2, page 22.

In general, we will say that a conditional valuation function V *qualitatively accepts* a conditional, written as $V \models (B|A)$, iff $V(A\bar{B}) < V(AB)$ is satisfied¹. Making this compatible with the Ramsey test means to postulate

$$(V * A)(\bar{B}) < (V * A)(B) \quad \text{iff} \quad V(A\bar{B}) < V(AB).$$

This suggests that an adequate revision operator should fulfill

$$(V * A)(B) = \alpha \odot V(AB)$$

for some $\alpha \in \mathcal{A}, \alpha \neq 0^{\mathcal{A}}$, and for all $B \in \mathcal{L}$ with $AB \not\equiv \perp$. Observing that $(V * A)(\top) = 1$, by definition of a conditional valuation function, this would imply $\alpha = V(A)^{-1}$, i.e.

$$V * A(B) = V(B|A) \quad (4.4)$$

for all $B \in \mathcal{L}$ with $AB \not\equiv \perp$. So, following a qualitative approach and using the Ramsey test, an important quantitative idea arises how to realize the (propositional) revision of conditional valuation functions appropriately. Note that,

¹ Note that the degrees of plausibility are measured differently for OCF's and for general conditional valuation functions. Those of conditional valuation functions follow the intuitive ordering of real numbers (as is suitable for possibility and probability distributions), while those of OCF's traditionally reverse this ordering.

in spite of the obvious similarity of (4.4) with Bayesian conditionalization (see, for instance, (4.7) below), this condition does not force $(V * A)(B) = 0^A$ for $AB \equiv \top$, because (4.4) is only assumed to apply for $B \in \mathcal{L}$ with $AB \neq \perp$.

Things are more complicated for probability functions, P . Here the notion of belief may be dealt with in two different, but nevertheless intuitive, ways, one by setting

$$P \models A \quad \text{iff} \quad P(A) = 1, \quad (4.5)$$

the other one by setting

$$P \models A \quad \text{iff} \quad P(\bar{A}) < P(A) \quad (4.6)$$

(see also Section 3.3). While the first approach seems to be most appropriate for establishing propositional beliefs, the second one fits better the intended meaning of conditional beliefs. Combining both methods with a revision operator and applying the Ramsey test, respectively, gives rise to two different criteria for accepting conditional beliefs. Although (4.6) appears to be a proper candidate for realizing qualitative belief revision within a probabilistic framework, the problem here is that, due to the non-discreteness of real numbers, no generally accepted “best” revision operator to establish $P(\bar{A}) < P(A)$ exists (though there is a vast amount of such operators; in Section 4.5, we will at least identify those revision operators being appropriate from a conditional logical point of view).

On the contrary, criterion (4.5) gives rise to one of the oldest and most popular revision operators, namely (*Bayesian*) *conditionalization*:

$$(P * A)(\omega) = P(\omega|A) = \begin{cases} \frac{P(\omega)}{P(A)} & \text{if } \omega \models A \\ 0 & \text{if } \omega \not\models A \end{cases} \quad (4.7)$$

But conditionalization is not a full revision operator in the AGM sense, it is only an expansion (see Section 2.2, page 14) because it does not allow us to establish beliefs which are contrary to the certain beliefs held in P .

So in a probabilistic framework, taking the *quantitative* degrees of probability into account in a belief revision process actually seems to be more appropriate than a purely qualitative point of view. That is, instead of considering belief revision as restricted to establishing beliefs for certain, we should be concerned with revisions of the form $P * A[x]$, yielding a probability function P^* such that $P^*(A) = x$. Moreover, a quantitative version of the Ramsey test can be formalized as

$$P \models (B|A)[x] \quad \text{iff} \quad P * A[1] \models B[x] \quad (4.8)$$

By taking $*$ as conditionalization, we obtain

$$P \models (B|A)[x] \quad \text{iff} \quad P(B|A) = x$$

for the acceptance of probabilistic conditionals (confer (3.8), page 34). Further generalizations are possible, for instance

$$P \models (B[x]|A[y]) \quad \text{iff} \quad P * A[y] \models B[x], \quad (4.9)$$

permitting us to deal with generalized conditionals. One method for achieving revisions of the form $P * A[x]$ is *Jeffrey's rule* (see (2.12), page 24), which coincides in this case with the more powerful principle of minimum cross-entropy (see (2.13), page 25). The latter technique also allows of realizing revisions by (sets of) probabilistic conditionals and thereby constitutes an interesting and important example for revision operators of epistemic states.

One of the aims of this chapter is to develop a formal *principle of conditional preservation* which underlies both qualitative and quantitative revisions by conditional beliefs, in that it implies the qualitative axioms (CR5) - (CR7) and can also be applied in a purely quantitative framework.

First, however, we will exemplify the eight postulates (CR0)-(CR7) of Section 4.1 by presenting a revision operator for ordinal conditional functions which satisfies all of these postulates.

4.4 A Revision Operator for Ordinal Conditional Functions

Propositional revisions of ordinal conditional functions $\kappa : \Omega \rightarrow \mathbb{N} \cup \{0, \infty\}$ have been investigated by several authors. For instance, Spohn [Spo88] proposed the following revision operator $*_S$ to establish belief in A *with firmness* n :

$$(\kappa *_S A[m])(\omega) = \begin{cases} \kappa(\omega) - \kappa(A) & \text{if } \omega \models A, \\ \kappa(\omega) - \kappa(\overline{A}) + m & \text{if } \omega \models \overline{A}. \end{cases} \quad (4.10)$$

Spohn calls $\kappa *_S A[m]$ the (A, m) -*conditionalization* of κ .

Darwiche and Pearl [DP97a] modified this approach to define a revision operator $*_{DP}$ which always strengthens the belief in A :

$$(\kappa *_{DP} A)(\omega) = \begin{cases} \kappa(\omega) - \kappa(A) & \text{if } \omega \models A, \\ \kappa(\omega) + 1 & \text{if } \omega \models \overline{A}. \end{cases} \quad (4.11)$$

That is, $\kappa *_{DP} A = \kappa *_S A[\kappa(\overline{A}) + 1]$.

Goldszmidt and Pearl [GP96] considered two different types of propositional revision, namely *conditionalization of Type-J* and *conditionalization of*

Type-L. Their conditionalization of Type-J corresponds exactly to Spohn's (A, m) -conditionalization.

In the following, we will define a *conditional revision operator* $*_0$ for ordinal conditional functions that satisfy all of the postulates (CR0)-(CR7). In particular, $*_0$ realizes the idea of conditional preservation developed so far:

For an ordinal conditional function $\kappa : \Omega \rightarrow \mathbb{N} \cup \{0, \infty\}$ and a conditional $(B|A)$, we define $\kappa *_0 (B|A)$ by setting

$$\kappa *_0 (B|A) (\omega) = \begin{cases} \kappa(\omega) - \kappa(B|A) & \text{if } \omega \models AB \\ \kappa(\omega) + \alpha + 1 & \text{if } \omega \models A\bar{B} \\ \kappa(\omega) & \text{if } \omega \models \bar{A} \end{cases} \quad (4.12)$$

where

$$\alpha = \begin{cases} -1, & \text{if } \kappa(AB) < \kappa(A\bar{B}), \\ 0, & \text{else} \end{cases}$$

This revision operator $*_0$ generalizes the propositional revision operator $*_{DP}$ (see (4.11) above), except for one crucial difference: $*_{DP}$ always strenghtens the belief in the revising proposition A , even if A is already believed with firmness > 0 . So Darwiche and Pearl's revision operator violates Postulate (CR2), while $*_0$ complies with it: If already $\kappa \models (B|A)$, then $\kappa *_0 (B|A) = \kappa$, due to $\alpha = -1$ in this case.

The check of the postulates (CR0) - (CR7) is straightforward, due to the representation Theorems 4.2.1 and 4.2.2. So we have

Proposition 4.4.1. *The conditional revision operator $*_0$ defined by (4.12) satisfies all of the postulates (CR0) - (CR7).*

Example 4.4.1 (Diagnosis). A physician has to make a diagnosis. The patient he is facing obviously feels ill, and at first glance, the physician supposes that the patient is suffering from disease D causing mainly two symptoms, S_1 (major symptom) and S_2 (minor symptom). To obtain certainty about these symptoms, further examinations will be necessary the next days. The following table shows an ordinal conditional function κ representing the epistemic state of the physician under these conditions:

D	S_1	S_2	κ	D	S_1	S_2	κ
0	0	0	5	1	0	0	4
0	0	1	2	1	0	1	2
0	1	0	1	1	1	0	0
0	1	1	4	1	1	1	3

So in this epistemic state, $DS_1\bar{S}_2$ is considered to be the most plausible world, i.e. $Bel(\kappa) \models D, S_1, \bar{S}_2$. Moreover, we have $\kappa \models (S_1|D), (D|S_1), (\bar{S}_2|D)$, as well as $\kappa \models (S_1|\bar{D})$.

In the evening, the physician finds an article in a medical journal, pointing out that symptom D_2 has recently proved to be of major importance for disease D . So the physician wants to revise κ to incorporate the new conditional belief $(S_2|D)$.

We make use of the revision operator $*_0$ introduced by (4.12) to calculate a revised ordinal conditional function $\kappa *_0 (S_2|D)$. Here we have $\kappa(DS_2) = 2, \kappa(D\bar{S}_2) = 0$, so $\kappa \not\models (S_2|D)$ and $\kappa(S_2|D) = 2$. By (4.12), we obtain

D	S_1	S_2	κ	$\kappa *_0 (S_2 D)$
0	0	0	5	5
0	0	1	2	2
0	1	0	1	1
0	1	1	4	4
1	0	0	4	5
1	0	1	2	0
1	1	0	0	1
1	1	1	3	1

Comparing κ and $\kappa *_0 (S_2|D)$, we see that the physician still believes in D , but that his belief in S_1 is given up in favor of now believing S_2 . Moreover, the conditional relationship between D and S_1 is weakened, $(D|S_1)$ can no longer be entailed from $\kappa *_0 (S_2|D)$, and $\kappa *_0 (S_2|D) \models (\bar{S}_1|D)$. The plausibility of DS_1 , however, is still quite high ($\kappa *_0 (S_2|D)(DS_1) = 1$). Although a common occurrence of both symptoms S_1 and S_2 is still excluded ($\kappa *_0 (S_2|D)(S_1S_2) \neq 0$), its plausibility is raised from 3 to 1.

The operator $*_0$ defined by (4.12) can be extended to adopt conditionals $(B|A)[m]$ quantified by a degree of firmness $m > 0$:

$$\kappa *_0 (B|A)[m](\omega) = \begin{cases} \kappa(\omega) - \kappa(B|A) & \text{if } \omega \models AB \\ \kappa(\omega) + m - \kappa(\bar{B}|A) & \text{if } \omega \models A\bar{B} \\ \kappa(\omega) & \text{if } \omega \models \bar{A} \end{cases} \quad (4.13)$$

It is easy to check that also this extended operator $*_0$ satisfies (a quantitative version of) Postulate (CR2).

4.5 The Principle of Conditional Preservation

Minimality of change is a crucial paradigm for belief revision, and a “principle of conditional preservation” is to realize this idea of minimality when conditionals are involved in change. Minimizing absolutely the changes in conditional beliefs, as in [BG93], is an important proposal to this aim, but

it does not always lead to intuitive results (cf. [DP97a]). The idea we will develop here rather aims at *preserving the conditional structure of knowledge* within an epistemic state which we assume to be represented by a conditional valuation function $V : \mathcal{L} \rightarrow \mathcal{A}$ (cf. Section 3.2). The propositions and theorems to be presented in this section extend the results of [KI99a].

The notion of a conditional structure with respect to a set \mathcal{R} of conditionals was defined in Section 3.5, and in Section 3.6, we explained what it means for V to follow the structure imposed by \mathcal{R} on the set of worlds by introducing the notion of conditional indifference (cf. Definition 3.6.1).

Pursuing this approach further in the framework of belief revision, a revision of V by simultaneously incorporating the conditionals in \mathcal{R} , $V^* = V * \mathcal{R}$, can be said to preserve the conditional structure of V with respect to \mathcal{R} if the *relative change function* $V^* \odot V^{-1}$ is indifferent with respect to \mathcal{R} . Taking into regard prior knowledge V and the worlds ω with $V(\omega) = 0^{\mathcal{A}}$ appropriately, this gives rise to the following definitions:

Definition 4.5.1. *Let $V : \mathcal{L} \rightarrow \mathcal{A}$ be a conditional valuation function, and let \mathcal{R} be a finite set of (quantified) conditionals. Let $V^* = V * \mathcal{R}$ denote the result of revising V by \mathcal{R} ; in particular, suppose that $V^*(A) \neq 0^{\mathcal{A}}$ for all $(B|A) \in \mathcal{R}$.*

1. V^* is called V -consistent iff $V(\omega) = 0^{\mathcal{A}}$ implies $V^*(\omega) = 0^{\mathcal{A}}$; V^* is called strictly V -consistent iff $V(\omega) = 0^{\mathcal{A}} \Leftrightarrow V^*(\omega) = 0^{\mathcal{A}}$;
2. If V^* is V -consistent, then the relative change function $(V^*/V) : \Omega \rightarrow \mathcal{A}$ is defined by

$$(V^*/V)(\omega) = \begin{cases} V^*(\omega) \odot V(\omega)^{-1} & \text{if } V(\omega) \neq 0^{\mathcal{A}} \\ 0^{\mathcal{A}} & \text{if } V(\omega) = 0^{\mathcal{A}} \end{cases}$$

3. V^* is strictly indifferent (indifferent) with respect to \mathcal{R} and V iff V^* is V -consistent and the following two conditions hold:
 - (i) If $V^*(\omega) = 0^{\mathcal{A}}$ then $V(\omega) = 0^{\mathcal{A}}$, or there is $(B|A) \in \mathcal{R}$ such that $\sigma_{(B|A)}(\omega) \neq 1$ and $V^*(\omega') = 0^{\mathcal{A}}$ for all ω' with $\sigma_{(B|A)}(\omega') = \sigma_{(B|A)}(\omega)$.
 - (ii) $(V^*/V)(\hat{\omega}_1) = (V^*/V)(\hat{\omega}_2)$ whenever $\sigma_{\mathcal{R}}(\hat{\omega}_1) = \sigma_{\mathcal{R}}(\hat{\omega}_2)$ (and $\hat{\omega}_1 \hat{\omega}_0 = \hat{\omega}_2 \hat{\omega}_0$) for $\hat{\omega}_1, \hat{\omega}_2 \in \hat{\Omega}_+^*$, where $\hat{\Omega}_+^* = \langle \omega \in \Omega \mid V^*(\omega) \neq 0^{\mathcal{A}} \rangle$.
4. A revision V^* is said to satisfy the (strict) principle of conditional preservation with respect to \mathcal{R} and V iff V^* is (strictly) indifferent with respect to \mathcal{R} and V .

Thus in a numerical framework, the principle of conditional preservation is realized as an indifference property.

Remark 4.5.1. Though the relative change function (V^*/V) is not a conditional valuation function, it may nevertheless be extended to a homomorphism $(V^*/V) : \widehat{\Omega}_+^* \rightarrow (\mathcal{A}, \odot)$ (see Section 3.6). Therefore, Definition 4.5.1(3) is an appropriate modification of Definition 3.6.1 for revisions.

Note that the principle of conditional preservation is based only on observing conditional structures, without using any acceptance conditions or taking quantifications of conditionals into account.

To illustrate the postulate for conditional preservation once again in a probabilistic setting, we will return to the Florida murderers-example (see Example 3.5.1 in Section 3.5, page 40).

Example 4.5.1 (Florida murderers, continued). The propositional variables involved here are $V = \underline{Victim}$ (of the murder) is black or white, respectively, $v \in \{v_b, v_w\}$, $M = \underline{Murderer}$ is black or white, respectively, $m \in \{m_b, m_w\}$, and $D = \underline{Murderer}$ is sentenced to \underline{Death} , $d \in \{d, \bar{d}\}$. The following probability distribution P mirrors the sentencing policy in the US state of Florida during a six years period:

ω	$P(\omega)$	ω	$P(\omega)$
$v_w m_w d$	0.0151	$v_w m_w \bar{d}$	0.4353
$v_w m_b d$	0.0101	$v_w m_b \bar{d}$	0.0502
$v_b m_w d$	0	$v_b m_w \bar{d}$	0.0233
$v_b m_b d$	0.0023	$v_b m_b \bar{d}$	0.4637

Assume that in a following year, we observe a slightly changed relationship between m_b and d , say $(d|m_b)[0.03]$ instead of $(d|m_b)[0.0236]$, and we want P to be adjusted to this new information. So we have $\mathcal{R} = \{(d|m_b)[0.03]\}$, and let two symbols $\mathbf{a}^+, \mathbf{a}^-$ be associated with \mathcal{R} . The conditional structures with respect to \mathcal{R} are calculated easily as follows:

$$\begin{aligned}
 \sigma_{\mathcal{R}}(v_w m_w d) = \sigma_{\mathcal{R}}(v_w m_w \bar{d}) &= \sigma_{\mathcal{R}}(v_b m_w d) = \sigma_{\mathcal{R}}(v_b m_w \bar{d}) = 1, \\
 \sigma_{\mathcal{R}}(v_w m_b d) = \sigma_{\mathcal{R}}(v_b m_b d) &= \mathbf{a}^+, \\
 \sigma_{\mathcal{R}}(v_w m_b \bar{d}) = \sigma_{\mathcal{R}}(v_b m_b \bar{d}) &= \mathbf{a}^-.
 \end{aligned}$$

Consider the elements $\widehat{\omega}_1 = v_w m_b d \cdot v_b m_b \bar{d}$ and $\widehat{\omega}_2 = v_b m_b d \cdot v_w m_b \bar{d}$ with equal conditional structures $\sigma_{\mathcal{R}}(\widehat{\omega}_1) = \mathbf{a}^+ \mathbf{a}^- = \sigma_{\mathcal{R}}(\widehat{\omega}_2)$. Therefore for P^* to be indifferent with respect to \mathcal{R} and P , it has to satisfy

$$\frac{P^*(v_w m_b d) P^*(v_b m_b \bar{d})}{P^*(v_b m_b d) P^*(v_w m_b \bar{d})} = \frac{P(v_w m_b d) P(v_b m_b \bar{d})}{P(v_b m_b d) P(v_w m_b \bar{d})},$$

which corresponds to equation (3.17), page 41.

Thus the concept of conditional structures helps us to get a technically clear and precise formalization of the intuitive idea of conditional preservation.

Thus we have developed two formal approaches to realize the idea of preserving conditional beliefs, one via the postulates (CR5)-(CR7) presented in Section 4.1, the other by applying the concept of conditional indifference appropriately in Definition 4.5.1(4). Though Proposition 3.5.1 reveals a connection between conditional structures, as introduced in Section 3.5, and the relations \sqsubseteq and \sqcup used to formalize (CR5)-(CR7), it still remains to make clear the compatibility of both approaches. Theorem 4.5.1 below will show that, within a qualitative setting, in the case that \mathcal{R} consists of only one conditional $(B|A)$, any strictly V -consistent revision $V * \mathcal{R}$ satisfying the principle of conditional preservation with respect to \mathcal{R} and V also obeys the postulates (CR5)-(CR7).

We begin by characterizing revisions $V^* = V * \mathcal{R} = V * (B|A)$ satisfying the principle of conditional preservation with respect to $\mathcal{R} = \{(B|A)\}$ and V . As a basic requirement for such revisions, we will only presuppose that $V^*(A) \neq 0^A$, instead of the (stronger) success postulate $V^* \models (B|A)$. This makes the results to be presented independent of acceptance conditions and helps concentrating on conditional structures; in particular, it will be possible to make use of these results even when conditionals are assigned numerical degrees of acceptance. Note that the principle of conditional preservation with respect to \mathcal{R} does not imply the acceptance of \mathcal{R} in general.

Proposition 4.5.1. *Let $V : \mathcal{L} \rightarrow \mathcal{A}$ be a conditional valuation function, and let $\mathcal{R} = \{(B|A)\}$, $(B|A) \in (\mathcal{L} \mid \mathcal{L})$ consist of only one conditional. Let $V^* = V * \mathcal{R} = V * (B|A)$ denote a revision of V by $(B|A)$ such that $V^*(A) \neq 0^A$.*

1. *V^* satisfies the strict principle of conditional preservation with respect to V and \mathcal{R} iff there are constants $\alpha^+, \alpha^- \in \mathcal{A}$ such that*

$$V^*(\omega) = \begin{cases} \alpha^+ \odot V(\omega) & \text{if } \omega \models AB \\ \alpha^- \odot V(\omega) & \text{if } \omega \models A\bar{B} \\ V(\omega) & \text{if } \omega \models \bar{A} \end{cases} \quad (4.14)$$

2. *V^* satisfies the principle of conditional preservation with respect to V and \mathcal{R} iff there are constants $\alpha_0, \alpha^+, \alpha^- \in \mathcal{A}$ such that*

$$V^*(\omega) = \begin{cases} \alpha^+ \odot V(\omega) & \text{if } \omega \models AB \\ \alpha^- \odot V(\omega) & \text{if } \omega \models A\bar{B} \\ \alpha_0 \odot V(\omega) & \text{if } \omega \models \bar{A} \end{cases} \quad (4.15)$$

3. *If V^* is strictly V -consistent, then all constants $\alpha_0, \alpha^+, \alpha^- \in \mathcal{A}$ in parts 1. and 2. may be chosen $\neq 0^A$.*

Before focussing on qualitative acceptance of conditionals (see Section 4.3), we will formulate a quantitative counterpart of postulate (CR5) for conditional valuation functions V :

(CR5^{quant}) If $(D|C) \perp\!\!\!\perp (B|A)$ and $V(CD), (V * (B|A))(CD) \neq 0^A$, then

$$V(C\overline{D}) \odot V(CD)^{-1} = (V * (B|A))(C\overline{D}) \odot (V * (B|A))(CD)^{-1}.$$

(CR5^{quant}) ensures that essentially, the values assigned to conditionals which are perpendicular to the revising conditional are not changed under revision:

Lemma 4.5.1. *Suppose the revision $V * (B|A)$ is strictly V -consistent and satisfies (CR5^{quant}). Then for any conditional $(D|C) \perp\!\!\!\perp (B|A)$ with $V(C) \neq 0^A$, it holds that*

$$V(D|C) = (V * (B|A))(D|C)$$

The following proposition shows that substantially, (CR5^{quant}) is stronger than its qualitative counterpart (CR5):

Proposition 4.5.2. *Let $V^* = V * \mathcal{R} = V * \{(B|A)\}$ denote a strictly V -consistent revision of V by $(B|A)$ such that $V^*(A) \neq 0^A$. If V^* fulfills (CR5^{quant}), then it also satisfies (CR5).*

The following theorem identifies the principle of conditional preservation (or conditional indifference, respectively) as a fundamental device to guide reasonable changes in the conditional structure of knowledge:

Theorem 4.5.1. *Let $V : \mathcal{L} \rightarrow \mathcal{A}$ be a conditional valuation function, and let $\mathcal{R} = \{(B|A)\}$, $(B|A) \in (\mathcal{L} | \mathcal{L})$, consist of only one conditional. Let $V^* = V * \mathcal{R}$ denote a strictly V -consistent revision of V by \mathcal{R} fulfilling the postulates (CR1) (success) and (CR2) (stability).*

If V^ satisfies the principle of conditional preservation, then the revision also satisfies postulate (CR5^{quant}) and the postulates (CR6) and (CR7); in particular, it satisfies all of the postulates (CR5)-(CR7) (see Section 4.1).*

4.6 C-Revisions and C-Representations

Essentially, the principle of conditional preservation means the strict indifference of the relative change function (V^*/V) with respect to the revising set \mathcal{R} of conditionals. In the preceding section, we showed that in case that

\mathcal{R} consists of only one conditional, this is compatible with the qualitative notion of conditional preservation described by the postulates (CR5)-(CR7).

In this section, we will return to the more general framework of considering the simultaneous revision by a *set* of (quantified) conditionals, and we will take the success condition (CR1), $V^* \models \mathcal{R}$, more explicitly into account. In particular, for probability functions and for ordinal conditional functions, we will present practical schemes to obtain revisions satisfying the principle of conditional preservation by making use of the results in Section 3.6.

Definition 4.6.1. *Let $V, V^* : \mathcal{L} \rightarrow \mathcal{A}$ be conditional valuation functions, and let \mathcal{R} be a set of (quantified) conditionals. A conditional valuation function $V^* : \mathcal{L} \rightarrow \mathcal{A}$ is called a (strict) c-revision of V by \mathcal{R} iff V^* satisfies the (strict) principle of conditional preservation with respect to V and \mathcal{R} , and $V^* \models \mathcal{R}$.*

A c-revision is based both on V and \mathcal{R} , using V as a reference point and \mathcal{R} as a guideline for changes. The prefix “c” marks the conditional well-behavedness of such revisions. C-revisions will sometimes also be called *c-adaptations*, in particular within a probabilistic framework (cf. [KI98a]).

A special case arises if no prior knowledge to be revised is at hand, and the set \mathcal{R} of (quantified) conditionals constitutes the only knowledge available. Then the problem is to find a conditional valuation function, i.e. an epistemic state, that represents \mathcal{R} most adequately. This *representation problem* may be considered as a particular revision problem, in taking the uniform conditional valuation function (see Definition 3.2.2, page 33), V_0 , as an appropriate conditional valuation function to start revision with. By assigning the same degree of probability, plausibility, etc. to each world in Ω , V_0 represents a state of complete ignorance in this framework.

Definition 4.6.2. *A conditional valuation function $V^* : \mathcal{L} \rightarrow \mathcal{A}$ is called a (strict) c-representation of a set \mathcal{R} of (quantified) conditionals, iff V^* satisfies the (strict) principle of conditional preservation with respect to V_0 and \mathcal{R} , and $V^* \models \mathcal{R}$.*

The following proposition shows that the uniform distribution actually plays a neutral part for c-representations, having only a normalizing effect on the representing conditional valuation function. Furthermore, the compatibility of the concepts of conditional indifference, developed for revisions in Section 4.5 and for conditional valuation functions in Section 3.6, is stated (see, in particular, Definitions 3.6.1 and 4.5.1).

Proposition 4.6.1. *A conditional valuation function V^* is indifferent with respect to \mathcal{R} and V_0 iff V^* is indifferent with respect to \mathcal{R} .*

We will now transfer the notions and results of Section 3.6 into the framework of revision (see pages 48 ff.).

Definition 4.6.3. *Let $V, V^* : \mathcal{L} \rightarrow \mathcal{A}$ be conditional valuation functions, let \mathcal{R} be a set of (quantified) conditionals.*

1. *V^* is called a faithful c-revision of V by \mathcal{R} iff V^* is a c-revision of V by \mathcal{R} satisfying $\ker_0 \sigma_{\mathcal{R}} \cap \widehat{\Omega}_+^* = \ker_0 (V^*/V)$.*
2. *V^* is called a faithful c-representation of \mathcal{R} iff V^* is a c-representation of \mathcal{R} satisfying $\ker_0 \sigma_{\mathcal{R}} \cap \widehat{\Omega}_+^* = \ker_0 V^*$.*

Due to their strict obeying of prior knowledge and new information, faithful c-revisions and c-representations will prove to be especially useful in the context of knowledge discovery and data mining (see Chapter 8).

The next theorem characterizes revisions of ordinal conditional functions that satisfy the principle of conditional preservation. The theorem is obvious by observing Corollary 3.6.4, page 51.

Theorem 4.6.1. *Let κ, κ^* be ordinal conditional functions, and let $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\}$ be a (finite) set of conditionals in $(\mathcal{L} | \mathcal{L})$.*

A revision $\kappa^ = \kappa * \mathcal{R}$ satisfies the principle of conditional preservation iff $\kappa^*(A_i) \neq \infty$ for all $i, 1 \leq i \leq n$, and there are numbers $\kappa_0, \kappa_i^+, \kappa_i^- \in \mathbb{Q}, 1 \leq i \leq n$, such that*

$$\kappa^*(\omega) = \kappa(\omega) + \kappa_0 + \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i \overline{B_i}}} \kappa_i^- \quad (4.16)$$

for all $\omega \in \Omega$.

Combining Theorem 4.6.1 with Lemma 4.3.1, we obtain

Corollary 4.6.1. *Let $\mathcal{R} = \{(B_1|A_1), \dots, (B_n|A_n)\}$ be a (finite) set of conditionals in $(\mathcal{L} | \mathcal{L})$, and let κ be an ordinal conditional function.*

κ is a c-representation of \mathcal{R} iff $\kappa(A_i) \neq \infty$ for all $i, 1 \leq i \leq n$, and there are numbers $\kappa_0, \kappa_i^+, \kappa_i^- \in \mathbb{Q}, 1 \leq i \leq n$, such that

$$\kappa(\omega) = \kappa_0 + \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i \overline{B_i}}} \kappa_i^-, \quad \omega \in \Omega, \quad (4.17)$$

and

$$\begin{aligned} \kappa_i^+ - \kappa_i^- &< \min_{\omega \models A_i \bar{B}_i} \left(\sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \bar{B}_j}} \kappa_j^- \right) \\ &\quad - \min_{\omega \models A_i B_i} \left(\sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \bar{B}_j}} \kappa_j^- \right) \end{aligned}$$

A well-known method to calculate an ordinal conditional function apt to represent a (finite) set $\mathcal{R} = \{r_i = (B_i|A_i) \mid 1 \leq i \leq n\}$ of conditionals is the *system-Z* of Goldszmidt and Pearl ([GP92, GP96]). The basic idea of system-Z is to observe the (logical) interactions of the conditionals in \mathcal{R} which are described by the notion of *tolerance*. A conditional $(B|A)$ is *tolerated* by a set of conditionals \mathcal{S} iff there is a world ω such that ω confirms $(B|A)$ and ω does not refute any of the conditionals in \mathcal{S} . If \mathcal{R} is consistent, then there is an ordered partition $\mathcal{R}_0, \mathcal{R}_1, \dots, \mathcal{R}_k$ of \mathcal{R} such that each conditional in \mathcal{R}_m is tolerated by $\bigcup_{j=m}^k \mathcal{R}_j$, $0 \leq m \leq k$ (for more details see, for instance, [GP96]).

The system-Z ranking function, κ^z , representing \mathcal{R} is given by

$$\kappa^z(\omega) = \begin{cases} 0, & \text{if } \omega \text{ does not falsify any } r_i, \\ 1 + \max_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} Z(r_i), & \text{otherwise} \end{cases} \quad (4.18)$$

where $Z(r_i) = j$ iff $r_i \in \mathcal{R}_j$. κ^z assigns to each world ω the lowest possible rank admissible with respect to the constraints in \mathcal{R} .

Comparing (4.18) with (4.17), we see that in general, κ^z is *not* a c-representation of \mathcal{R} , since in its definition (4.18), *maximum* is used instead of *summation* (see Example 4.6.1 below). The partition $\mathcal{R}_0, \mathcal{R}_1, \dots, \mathcal{R}_k$ of \mathcal{R} , however, may well serve to define appropriate constants κ_i^- in (4.17). Setting $\kappa_0 := \kappa_i^+ := 0$, and

$$\kappa_i^- := \left(\min_{\omega \models A_i B_i} \sum_{\substack{r_j \in \bigcup_{l=0}^{m-1} \mathcal{R}_l \\ \omega \models A_j \bar{B}_j}} \kappa_j^- \right) + 1 \quad (4.19)$$

for each conditional $r_i \in \mathcal{R}_m$, $m = 0, \dots, k$ successively, we obtain a c-representation κ_c^z of \mathcal{R} via

$$\kappa_c^z(\omega) := \sum_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} \kappa_i^- \quad (4.20)$$

Example 4.6.1 (System Z). Consider the set \mathcal{R} consisting of the following conditionals:

- $r_1 : (f|b) \quad \text{Birds fly.}$
- $r_2 : (b|p) \quad \text{Penguins are birds.}$
- $r_3 : (\bar{f}|p) \quad \text{Penguins do not fly.}$
- $r_4 : (w|b) \quad \text{Birds have wings.}$
- $r_5 : (a|f) \quad \text{Animals that fly are airborne.}$

(see Example 17 in [GP96, p. 68f]). Here \mathcal{R} is partitioned by $\mathcal{R}_0 = \{r_1, r_4, r_5\}$ and $\mathcal{R}_1 = \{r_2, r_3\}$ (for the details, see [GP96, p. 69]). By applying (4.19), we calculate $\kappa_1^- = \kappa_4^- = \kappa_5^- = 1$ and $\kappa_2^- = \kappa_3^- = 2$. Actually, in this example, κ_c^z from (4.20) coincides with the system- Z^* ranking function (cf. [GMP93, BP99]). We obtain, for example,

$$\begin{aligned} \kappa_c^z(\bar{p}\bar{b}\bar{f}wa) &= 0, \\ \kappa_c^z(pb\bar{f}wa) &= \kappa_1^- = 1, \\ \kappa_c^z(p\bar{b}f\bar{w}\bar{a}) &= \kappa_2^- + \kappa_3^- + \kappa_5^- = 5. \end{aligned}$$

In Table 4.1, we list the ranks of all possible worlds, first computed by system- Z , according to (4.18), and then computed as a c-representation of \mathcal{R} , according to (4.20). Comparing $\kappa^z(\omega)$ to $\kappa_c^z(\omega)$, we see that κ_c^z is more fine-grained. Table 4.1 also reveals that κ^z is not a c-representation of \mathcal{R} : Assigning the

Table 4.1. Rankings for Example 4.6.1

ω	$\kappa^z(\omega)$	$\kappa_c^z(\omega)$	ω	$\kappa^z(\omega)$	$\kappa_c^z(\omega)$
$pbfwa$	2	2	$\bar{p}\bar{b}fwa$	0	0
$pbfw\bar{a}$	2	3	$\bar{p}\bar{b}fw\bar{a}$	1	1
$pbf\bar{w}a$	2	3	$\bar{p}\bar{b}f\bar{w}a$	1	1
$pbf\bar{w}\bar{a}$	2	4	$\bar{p}\bar{b}f\bar{w}\bar{a}$	1	2
$p\bar{b}\bar{f}wa$	1	1	$\bar{p}\bar{b}\bar{f}wa$	1	1
$p\bar{b}\bar{f}w\bar{a}$	1	1	$\bar{p}\bar{b}\bar{f}w\bar{a}$	1	1
$p\bar{b}\bar{f}\bar{w}a$	1	2	$\bar{p}\bar{b}\bar{f}\bar{w}a$	1	2
$p\bar{b}\bar{f}\bar{w}\bar{a}$	1	2	$\bar{p}\bar{b}\bar{f}\bar{w}\bar{a}$	1	2
$p\bar{b}fwa$	2	4	$\bar{p}\bar{b}fwa$	0	0
$p\bar{b}fw\bar{a}$	2	5	$\bar{p}\bar{b}fw\bar{a}$	1	1
$p\bar{b}f\bar{w}a$	2	4	$\bar{p}\bar{b}f\bar{w}a$	0	0
$p\bar{b}f\bar{w}\bar{a}$	2	5	$\bar{p}\bar{b}f\bar{w}\bar{a}$	1	1
$p\bar{b}\bar{f}wa$	2	2	$\bar{p}\bar{b}\bar{f}wa$	0	0
$p\bar{b}\bar{f}w\bar{a}$	2	2	$\bar{p}\bar{b}\bar{f}w\bar{a}$	0	0
$p\bar{b}\bar{f}\bar{w}a$	2	2	$\bar{p}\bar{b}\bar{f}\bar{w}a$	0	0
$p\bar{b}\bar{f}\bar{w}\bar{a}$	2	2	$\bar{p}\bar{b}\bar{f}\bar{w}\bar{a}$	0	0

symbols \mathbf{a}_i^+ , \mathbf{a}_i^- to the conditionals r_i in \mathcal{R} , $1 \leq i \leq 5$, respectively, we obtain

$$\sigma_{\mathcal{R}} \left(\frac{p\bar{b}fwa \cdot \bar{p}b\bar{f}w\bar{a}}{p\bar{b}f\bar{w}\bar{a} \cdot \bar{p}b\bar{f}wa} \right) = \frac{\mathbf{a}_2^- \mathbf{a}_3^- \mathbf{a}_5^+ \cdot \mathbf{a}_1^+ \mathbf{a}_4^+ \mathbf{a}_5^-}{\mathbf{a}_2^- \mathbf{a}_3^- \mathbf{a}_5^- \cdot \mathbf{a}_1^+ \mathbf{a}_4^+ \mathbf{a}_5^+} = 1,$$

but

$$\begin{aligned} \kappa^z \left(\frac{p\bar{b}fwa \cdot \bar{p}b\bar{f}w\bar{a}}{p\bar{b}f\bar{w}\bar{a} \cdot \bar{p}b\bar{f}wa} \right) &= \kappa^z(p\bar{b}fwa) + \kappa^z(\bar{p}b\bar{f}w\bar{a}) \\ &\quad - \kappa^z(p\bar{b}f\bar{w}\bar{a}) - \kappa^z(\bar{p}b\bar{f}wa) \\ &= 3 - 2 = 1 \neq 0. \end{aligned}$$

In a probabilistic framework, we obtain the following theorem characterizing c-revisions of probability distributions, immediately by making use of Corollary 3.6.2, page 50 and observing the success condition (see also [KI98a]).

Theorem 4.6.2. *Suppose P is a probability distribution and $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$ is a P -consistent set of probabilistic conditionals.*

A probability distribution P^ is a c-revision of P with respect to \mathcal{R} if and only if there are real numbers $\alpha_0, \alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ with $\alpha_0 > 0$ and $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ satisfying the positivity condition*

$$\alpha_i^+, \alpha_i^- \geq 0, \quad \alpha_i^+ = 0 \text{ iff } x_i = 0, \quad \alpha_i^- = 0 \text{ iff } x_i = 1, \quad 1 \leq i \leq n, \quad (4.21)$$

and the adjustment condition

$$\begin{aligned} (1 - x_i) \alpha_i^+ \sum_{\omega \models A_i B_i} P(\omega) \prod_{\substack{j \neq i \\ \omega \models A_j B_j}} \alpha_j^+ \prod_{\substack{j \neq i \\ \omega \models A_j \bar{B}_j}} \alpha_j^- \\ = x_i \alpha_i^- \sum_{\omega \models A_i \bar{B}_i} P(\omega) \prod_{\substack{j \neq i \\ \omega \models A_j B_j}} \alpha_j^+ \prod_{\substack{j \neq i \\ \omega \models A_j \bar{B}_j}} \alpha_j^-, \quad 1 \leq i \leq n, \end{aligned} \quad (4.22)$$

such that

$$P^*(\omega) = \alpha_0 P(\omega) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} \alpha_i^- \quad (4.23)$$

for all worlds $\omega \in \Omega$.

Results for probabilistic and OCF-c-representations are easily obtained by using the corresponding uniform function as prior knowledge P_0 or κ_0 , respectively. The normalizing effect of the uniform prior function may then be

subsumed by the constant α_0 so that no explicit occurrence of prior knowledge will be necessary in the formulas (4.23) and (4.16), respectively.

Theorems 4.6.1 and 4.6.2 above provide a practical method to realize revisions following the principle of conditional preservation in an OCF- or probabilistic framework, respectively. Note that they generalize Proposition 4.5.1 appropriately in so far as now sets of conditionals are considered. Intuitively, the involved constants $\alpha_0, \alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ and $\kappa_0, \kappa_1^+, \kappa_1^-, \dots, \kappa_n^+, \kappa_n^-$, respectively, should be chosen appropriately to keep the amount of change “minimal”. The revision operator $*_0$ presented in Section 4.4 (see (4.12), page 61) obviously intends to meet this requirement.

In the next chapter, we will develop the ME-revision of a probability distribution P by a set of quantified conditionals \mathcal{R} from the approach (4.23) by imposing suitable constraints on the constants $\alpha_0, \alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$.

5. Characterizing the Principle of Minimum Cross-Entropy

Probability theory provides a sound and convenient machinery to be used for knowledge representation and automated reasoning (see, for instance, [Cox46, DPT90, DP91b, LS88, Pea88, TKG92]). In many cases, only relatively few relationships between relevant variables are known, due to incomplete information. Or maybe, an abstractional representation is intended, incorporating only fundamental relationships. In both cases, the knowledge explicitly stated is not sufficient to determine uniquely a probability distribution. One way to cope with this indetermination is to calculate upper and lower bounds for probabilities (cf. [Nil86, TKG92, DPT90]). This method, however, brings about two problems: Sometimes the inferred bounds are quite bad, and one has to handle intervals instead of single values.

An alternative way that provides best expectation values for the unknown probabilities and guarantees a logically sound reasoning is to use the *principle of maximum entropy* resp. the *principle of minimum cross entropy* to represent all available probabilistic knowledge by a unique distribution (see Section 2.5; cf. [Sho86, Kul68, Jay83a, GHK94]). Here we assume the available knowledge to constitute of a (consistent) set \mathcal{R} of conditionals, each equipped with a probability, usually providing only incomplete probabilistic knowledge.

The aim of this chapter is to establish a direct and constructive link between probabilistic conditionals and their suitable representation via distributions, taking prior knowledge into account if necessary. We develop the following four principles which mark the corner-stones for using quantified conditionals consistently for probabilistic knowledge representation and updating:

- (P1) The *principle of conditional preservation*: this is to express that prior conditional dependencies shall be preserved “as far as possible” under adaptation;
- (P2) the idea of a *functional concept* which underlies the adaptation and which allows us to calculate a posterior distribution from prior and new knowledge;

- (P3) the *principle of logical coherence*¹: posterior distributions shall be used coherently as priors for further inferences; and
- (P4) the *principle of representation invariance*: the resulting distribution shall not depend upon the actual probabilistic representation of the new information.

(P1) links numerical changes to the conditional structure of the new information. (P2) realizes a computable relationship between prior and posterior knowledge by means of appropriate real functions. (P3) forestalls ambivalent results of updating procedures, and (P4) should be self-evident within a probabilistic framework.

As we will show, the only method that solves the probabilistic revision problem

- (*_{prob}) Given a (prior) distribution P and some finite set of probabilistic conditionals $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\} \subseteq (\mathcal{L} | \mathcal{L})^{prob}$, how should P be modified to yield a (posterior) distribution P^* with $P^* \models \mathcal{R}$?

while obeying all of the principles (P1) to (P4) is provided by the principle of minimum cross-entropy. The first two axioms (P1) and (P2) will lead to a scheme for adjusting a prior distribution to new conditional information, and the principles of logical coherence and of representation invariance will be applied to this scheme, yielding the desired result. Thus a new characterization of the ME-principles is obtained, completely based on probabilistic conditionals and establishing reasoning at optimum entropy as a most fundamental inference method in the area of quantified uncertain reasoning.

Compared to the earlier papers [PV90, SJ80], the characterization presented here points out a more constructive approach to the ME-principles. We will show that ME-inference not only respects (*conditional*) *independencies* but that it is basically determined by *conditional dependencies* (obeying independence properties where no dependency exists), recommending the ME-principles as most adequate methods for reasoning with probabilistic conditionals. So in contrast to Bayesian networks (cf. e.g. [LS88]), probabilistic networks based on ME-techniques (cf. [RKI97b, RM96]) do not require lots of probabilities and independence assumptions to process quantified conditional knowledge properly. Moreover, the methods used in this book are quite different from those in [SJ80] and in [PV90]. In particular, there will be no need to make use of optimization theory, as in [SJ80], or to transfer the problem into the context of linear algebra, as in [PV90]. Our development explains clearly how the ME-principles may be completely based on probabi-

¹ this principle is called *principle of logical consistency* in [KI98a]

listic conditionals. This may improve significantly the explanatory features of computational systems that use these principles for knowledge representation and processing (as, e.g. SPIRIT, cf. [RKI97b, RM96]).

The results presented in this chapter were published in [KI96a] and in [KI98a].

5.1 Conditional Preservation

The principle of conditional preservation (P1) has already been formalized and explained in Section 4.5, using the notions of conditional structures of worlds and that of indifference with respect to (sets of) conditionals. In Section 4.6, Definition 4.6.1, revisions satisfying this principle of conditional preservation have been introduced as *c-revisions*, realizing perfectly a conditional-logical approach to the adaptation problem ($*_{prob}$):

Postulate (P1): conditional preservation

The solution P^* of ($*_{prob}$) is a c-revision.

Theorem 4.6.2 characterizes c-revisions providing a solution to problem ($*_{prob}$) as distributions of the form

$$P^*(\omega) = \alpha_0 P(\omega) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \quad (5.1)$$

where $\alpha_0, \alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ are real numbers with $\alpha_0 > 0$ and $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ satisfying the *positivity condition*, (4.21),

$$\alpha_i^+, \alpha_i^- \geq 0, \quad \alpha_i^+ = 0 \text{ iff } x_i = 0, \quad \alpha_i^- = 0 \text{ iff } x_i = 1 \quad (5.2)$$

and the *adjustment condition*, (4.23),

$$\begin{aligned} & (1 - x_i) \alpha_i^+ \sum_{\omega \models A_i B_i} P(\omega) \prod_{\substack{j \neq i \\ \omega \models A_j B_j}} \alpha_j^+ \prod_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} \alpha_j^- \\ &= x_i \alpha_i^- \sum_{\omega \models A_i \overline{B_i}} P(\omega) \prod_{\substack{j \neq i \\ \omega \models A_j B_j}} \alpha_j^+ \prod_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} \alpha_j^-, \end{aligned} \quad (5.3)$$

for $1 \leq i \leq n$. Due to their simple structure, c-revisions were taken as an intuitively appealing approach to realize adaptations to conditional constraints in [KI97a].

To maintain compatibility between prior and posterior distributions, P^* has to be P -consistent (which coincides with the notion of P -continuity, denoted as $P^* \ll P$; cf. Definition 4.5.1, p. 63), i.e. $P(\omega) = 0$ implies $P^*(\omega) = 0$. Thus to avoid obvious inconsistencies, the set \mathcal{R} is supposed to be P -consistent:

Definition 5.1.1. A set $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$ of probabilistic conditionals is said to be P -consistent iff there is some distribution Q with $Q \ll P$ and $Q \models \mathcal{R}$.

Definition 5.1.2. For a prior distribution P and some P -consistent set \mathcal{R} of probabilistic conditionals, let $C(P, \mathcal{R})$ denote the set of all c -revisions of P by \mathcal{R} :

$$C(P, \mathcal{R}) := \{P_c^* \mid P_c^* \text{ is a } c\text{-revision of } P \text{ by } \mathcal{R}\}.$$

Remark 5.1.1. Throughout this chapter, we will assume without further mentioning that the necessity of zero posterior probabilities is stated explicitly in \mathcal{R} , i.e. if for any $Q \ll P$, $Q \models \mathcal{R}$ implies $Q(\omega) = 0$, then $P(\omega) = 0$, or there is a conditional $(B|A)[x] \in \mathcal{R}$ such that either $x = 1$ and $\omega \models A\bar{B}$ or $x = 0$ and $\omega \models AB$.

The ME-solution to $(*_{prob})$ is the one distribution P_{ME} that satisfies all constraints in \mathcal{R} and has minimal cross-entropy with respect to P , i.e. P_{ME} solves the minimization problem

$$\min R(Q, P) = \sum_{\omega \in \Omega} Q(\omega) \log \frac{Q(\omega)}{P(\omega)} \quad (5.4)$$

s.t. Q is a probability distribution with $Q \models \mathcal{R}$

(see Section 2.5, p. 25).

If \mathcal{R} is a P -consistent (finite) set of conditionals, then the ME-solution P_{ME} of $(*_{prob})$ is guaranteed to exist (cf. [Csi75]).

The condition $Q \models \mathcal{R}$ imposed on a distribution Q can be transformed equivalently into a system of linear equality constraints for the probabilities $Q(\omega)$. Using the well-known Lagrange techniques (see, for instance, [Jay83b]), we may represent P_{ME} in the form

$$P_{ME}(\omega) = \alpha_0 P(\omega) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \alpha_i^{1-x_i} \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} \alpha_i^{-x_i}, \quad (5.5)$$

with the α_i 's being exponentials of the Lagrange multipliers, one for each conditional in \mathcal{R} , and $\alpha_0 = \exp(\lambda_0 - 1)$, where λ_0 is the Lagrange multiplier of the constraint $\sum_{\omega} Q(\omega) = 1$.

By construction, P_{ME} satisfies all conditionals in \mathcal{R} : $P_{ME}(B_i|A_i) = x_i$, which is equivalent to $(1-x)P_{ME}(A_i B_i) = xP_{ME}(A_i \overline{B}_i)$ for all i , $1 \leq i \leq n$. So $\alpha_1, \dots, \alpha_n$ are solutions of the nonlinear equations

$$\alpha_i = \frac{x_i}{1-x_i} \frac{\sum_{\omega \models A_i \overline{B}_i} P(\omega) \prod_{\substack{j \neq i \\ \omega \models A_j B_j}} \alpha_j^{1-x_j} \prod_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \alpha_j^{-x_j}}{\sum_{\omega \models A_i B_i} P(\omega) \prod_{\substack{j \neq i \\ \omega \models A_j B_j}} \alpha_j^{1-x_j} \prod_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \alpha_j^{-x_j}}, \quad (5.6)$$

$$\text{with } \alpha_i \begin{cases} > 0 & : x_i \in (0, 1) \\ = \infty & : x_i = 1 \\ = 0 & : x_i = 0 \end{cases}, \quad 1 \leq i \leq n, \quad (5.7)$$

using the conventions $\infty^0 = 1$, $\infty^{-1} = 0$ and $0^0 = 1$. α_0 arises simply as a normalizing factor. Each α_i symbolizes the impact of the corresponding rule when P is modified. It depends upon the prior distribution P , the other rules and probabilities in \mathcal{R} and - in a distinguished way - on the probability of the conditional it corresponds to. Using the representation formula (5.5) above, it is possible to indicate which of the conditionals in \mathcal{R} actually makes a contribution to a conditional information derived from the posterior ME-distribution (similar to listing active rules in rule based systems).

Comparing (5.5), (5.6) and (5.7) to (5.1), (5.3) and (5.2) above, we see that the ME-distribution P_{ME} is in particular a c-revision of P by \mathcal{R} , with $\alpha_i^+ = \alpha_i^{1-x_i}$ and $\alpha_i^- = \alpha_i^{-x_i}$.

Because ME-revisions exist for priors P and P -consistent sets \mathcal{R} , we have

Corollary 5.1.1. *For any prior distribution P and any P -consistent set \mathcal{R} of probabilistic conditionals, $C(P, \mathcal{R}) \neq \emptyset$.*

Note that we presupposed zero probabilities to be represented explicitly (cf. remark after Definition 5.1.2).

So c-revisions generalize the concept of ME-revisions and embed it into a conditional-logical environment. We will make use of probabilistic c-revisions in the form (5.1). Distributions of this type will play a major part in the following.

Definition 5.1.3. *Let P be a distribution, and let $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ be non-negative real numbers such that $\sum_{\omega} P(\omega) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \overline{B}_i}} \alpha_i^- > 0$.*

Then $P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-]$ denotes the distribution

$$P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-](\omega) := \alpha_0 P(\omega) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \overline{B}_i}} \alpha_i^-$$

$$\text{where } \alpha_0 = \left(\sum_{\omega} P(\omega) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} \alpha_i^- \right)^{-1}.$$

The normalizing factor α_0 is completely determined by P and $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$. Note that $P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-]$ is P -consistent.

According to Theorem 4.6.2, for any c-revision P^* of P by \mathcal{R} , there are non-negative real *weight factors* $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ satisfying (5.2) and (5.3) such that $P^* = P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-]$. Define

$$\begin{aligned} wf(P^*) &:= \{(\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-) \in \mathbb{R}^{2n} \mid \\ &\quad (\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-) \text{ satisfies (5.2), (5.3)} \\ &\quad \text{and } P^* = P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-]\} \end{aligned}$$

for any $P^* \in C(P, \mathcal{R})$. In general, weight factors of c-revisions are not uniquely determined, so that $\text{card}(wf(P^*)) \geq 1$.

As the proof of Theorem 4.6.2 shows, (5.2) ensures that all premises A_i occurring in \mathcal{R} have positive probabilities in $P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-]$, and (5.3) then is equivalent to $P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-] \models \mathcal{R}$.

Corollary 5.1.2. *Let P be a distribution, and suppose \mathcal{R} is a P -consistent set of probabilistic rules.*

If $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ are reals satisfying (5.2) then $P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-] \models \mathcal{R}$ iff $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ fulfill (5.3).

Therefore, we define

$$WF(P, \mathcal{R}) := \bigcup_{P^* \in C(P, \mathcal{R})} wf(P^*) \quad (5.8)$$

$$= \{(\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-) \in \mathbb{R}^{2n} \mid (\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-) \text{ satisfies (5.2), (5.3)}\} \quad (5.9)$$

So, c-revisions actually realize quite a simple idea of adaptation to new conditional information:

When calculating the posterior probability function P^* , one only has to check the conditional structure of each elementary event ω with respect to $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$, set up $P^*(\omega)$ according to (5.1) with unknown quantities $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ and then determine appropriate values for these $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ using (5.2) and (5.3), so that \mathcal{R} is satisfied. Finally, α_0 is computed as a normalizing factor to make P^* a probability distribution.

The following example briefly illustrates this adaptation scheme.

Example 5.1.1. Let P be a positive distribution over two atoms a, b , and suppose $\mathcal{R} = \{(b|a)[x]\}$ with $x \in (0, 1)$. Applying the formulas above, any c-revision P_c^* of P by \mathcal{R} may be written as

$$\begin{aligned} P_c^*(ab) &= \alpha_0 P(ab) \alpha_1^+ & P_c^*(a\bar{b}) &= \alpha_0 P(a\bar{b}) \alpha_1^- \\ P_c^*(\bar{a}b) &= \alpha_0 P(\bar{a}b) & P_c^*(\bar{a}\bar{b}) &= \alpha_0 P(\bar{a}\bar{b}) \end{aligned}$$

with

$$\alpha_1^+, \alpha_1^- > 0, (1-x)\alpha_1^+ P(ab) = x\alpha_1^- P(a\bar{b})$$

and

$$\alpha_0 = \left(\frac{1}{x} P(ab) \alpha_1^+ + P(a\bar{b}) + P(\bar{a}\bar{b}) \right)^{-1}$$

So, for instance, we obtain an infinite set of c-revisions of P by \mathcal{R} by setting simply

$$\alpha_1^+ = xP(a\bar{b})m, \alpha_1^- = (1-x)P(ab)m \text{ with } m \in \mathbb{N},$$

and choosing α_0 appropriately.

C-revisions provide a straightforward scheme to calculate solutions to the adjustment problem $(*_\text{prob})$. ME-revisions are a special instance of this scheme, and it is of interest to investigate which of the characteristics of ME-distributions also hold for c-revisions in general.

The author proved in [KI96b, KI98c] that c-revisions possess the properties of *system independence* and of *subset independence* which both played an outstanding part in Shore and Johnson's [SJ80] characterization of the ME-principle. They also cope in an elegant manner with irrelevant information in that posterior marginals are determined only by conditionals involving the respective variables (cf. [KI96b, KI98c]; see also [PV90]). All this is due to their modular, conditional-logical structure.

There is another principle that ME-adaptations actually seem to fail at first sight and that can now be formulated adequately and proved in terms of c-revisions: it is the *Atomicity Principle* stating that substituting formulas for variables shall not affect the adjustment process (cf. [PV97]):

Theorem 5.1.1 (Atomicity principle). *Let $\mathcal{V} = \{V_1, V_2, \dots\}$ and $\mathcal{V}' = \{V'_1, V'_2, \dots\}$ be two finite disjoint sets of binary propositional variables with corresponding sets of elementary events Ω resp. Ω' , and let V be another binary variable not contained in either of them. Suppose $\Delta \in \mathcal{L}(\mathcal{V}')$ is a propositional formula that is neither a tautology nor a contradiction, using only variables in \mathcal{V}' . Let $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$ be a set of probabilistic conditionals with antecedents A_i and consequences B_i in $\mathcal{L}(\mathcal{V} \cup \{V\})$. Let A_i^Δ resp. B_i^Δ denote the formulas that arise when each*

occurrence of V in A_i resp. B_i is replaced by Δ , $1 \leq i \leq n$, and so $\mathcal{R}^\Delta = \{(B_1^\Delta | A_1^\Delta)[x_1], \dots, (B_n^\Delta | A_n^\Delta)[x_n]\} \subseteq (\mathcal{L} | \mathcal{L})^{prob}(\mathcal{V} \cup \mathcal{V}')$.

Consider the two distributions P' over $\mathcal{V} \cup \mathcal{V}'$ and P over $\mathcal{V} \cup \{V\}$, respectively, that are related via $P(\dot{v}\omega) = \sum_{\substack{\omega' \in \Omega' \\ \omega' \models \Delta}} P'(\omega'\omega)$, and suppose \mathcal{R} to be

P -consistent. Then \mathcal{R}^Δ is P' -consistent, and $WF(P, \mathcal{R}) = WF(P', \mathcal{R}^\Delta)$.

We omit the straightforward but technical proof. This result emphasizes the importance of the weight factors $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ as logical representatives of an adaptation scheme.

In the rest of this chapter, we will investigate which postulates are to be imposed on a probabilistic c-revision so as to force the constants $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ to take on the ME-form, where the α_i 's are given by (5.6).

5.2 The Functional Concept

The concept of c-revisions is not perfect - it fails to satisfy uniqueness: Example 5.1.1 above shows that, even in the simple case when dealing with two variables and one conditional to be adjusted to, the resulting c-revision is not uniquely determined. In general, $WF(P, \mathcal{R})$ will contain lots of elements, and there will be many different posterior c-revisions. Demanding uniqueness means to assume a *functional concept* that guides the finding of a “best solution” so that a unique distribution of type (5.1) arises in dependence of the prior knowledge P and the new conditional information \mathcal{R} .

It is not only the abstract property of uniqueness, however, that makes a functional concept desirable. In a fundamental sense, there should be a clear and understandable dependence between prior distribution, new (conditional) information and resulting posterior distribution, i.e. a - somehow well-behaved - function $F : (P, \mathcal{R}) \mapsto P^*$ that works for all distributions P and all P -consistent sets \mathcal{R} (cf. [Gär88]). These arguments P and \mathcal{R} , however, are quite monstrous. The knowledge represented by them is usually huge and hard to grasp, let alone introducing such concepts as continuity or even differentiability to describe a functional well-behavedness.

Moreover, P^* should depend *significantly* only on the *relevant* parts of the prior P , i.e. relevant with respect to the new information \mathcal{R} . Treating this problem requires making clear what relevant information is, and how irrelevant information should be handled.

Let $\mathcal{R} = \{(B_1 | A_1)[x_1], \dots, (B_n | A_n)[x_n]\}$, and suppose P_1, P_2 are two distributions with $P_1(\omega | A_i) = P_2(\omega | A_i)$ for all $\omega \in \Omega$ and for all $i = 1, \dots, n$.

Then P_1, P_2 match on all parts which are *relevant* with respect to \mathcal{R} , so the difference in their posterior relative changes should be *insignificant*, namely a constant (due to possible differences in irrelevant parts):

Definition 5.2.1. Suppose $F : (P, \mathcal{R}) \mapsto P^*$ is a function that assigns a P -consistent distribution P^* satisfying \mathcal{R} to any pair (P, \mathcal{R}) with P being a distribution and \mathcal{R} representing a P -consistent set of probabilistic rules. Then F fulfills the relevance condition iff the following holds:

Let $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$, and suppose P_1, P_2 are two distributions with $P_1(\omega|A_i) = P_2(\omega|A_i)$ for all $\omega \in \Omega$ and for all $i = 1, \dots, n$, $P_1(\omega) = 0$ iff $P_2(\omega) = 0$ and such that \mathcal{R} is P_1 - and P_2 -consistent. Let $P_k^* := F(P_k, \mathcal{R}), k = 1, 2$. Then $P_1^*(\omega) = 0$ iff $P_2^*(\omega) = 0$, and there is a constant $const$ such that

$$\frac{P_1^*(\omega)}{P_1(\omega)} : \frac{P_2^*(\omega)}{P_2(\omega)} = const$$

for all $\omega \in \Omega$ with $P_k^*(\omega) \neq 0, k = 1, 2$.

Let \mathcal{AP} denote the set of all pairs (P, \mathcal{R}) representing a solvable adjustment problem $(*_prob)$:

$$\mathcal{AP} = \left\{ (P, \mathcal{R}) \mid P \text{ distribution, } \mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}, \mathcal{R} \text{ } P\text{-consistent} \right\}$$

According to postulate (P1), the solution to the adaptation problem $(*_prob)$ should be a c-revision. The following proposition shows that the prerequisites formulated in Definition 5.2.1 in fact are able to capture the idea of relevant information for c-revisions:

Proposition 5.2.1. Let $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$, and suppose P_1, P_2 are two distributions with $P_1(\omega|A_i) = P_2(\omega|A_i)$ for all $\omega \in \Omega$ and for all $i = 1, \dots, n$, $P_1(\omega) = 0$ iff $P_2(\omega) = 0$ and such that \mathcal{R} is P_1 - and P_2 -consistent. Then $WF(P_1, \mathcal{R}) = WF(P_2, \mathcal{R})$.

Therefore distributions incorporating the same relevant conditional knowledge have the same sets of weight factors occurring in the corresponding c-adaptations.

Let us henceforth assume that there is a function

$$F_c : \mathcal{AP} \ni (P, \mathcal{R}) \mapsto P_c^* \in C(P, \mathcal{R}) \quad (5.10)$$

that assigns to each pair $(P, \mathcal{R}) \in \mathcal{AP}$ a particular c-revision. We will describe F_c by specific properties of the weight factors involved.

Proposition 5.2.2. Assume $(P, \mathcal{R}) \in \mathcal{AP}$, $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$, and let $P_c^* = P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-] \in C(P, \mathcal{R})$ be a c-revision of P by \mathcal{R} with weight factors $(\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-) \in wf(P_c^*)$. Suppose $I \cup J$ is a partition of $\{1, \dots, n\}$, and set $P_I := P[\alpha_i^+, \alpha_i^-]_{i \in I}$, $\mathcal{R}_J := \{(B_j|A_j)[x_j]\}_{j \in J}$.

Then $(\alpha_j^+, \alpha_j^-)_{j \in J} \in WF(P_I, \mathcal{R}_J)$ and

$$P_I[\alpha_j^+, \alpha_j^-]_{j \in J} = P[\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-].$$

So, once weight factors $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^-$ are chosen to yield a “best” solution $P_c^* \in C(P, \mathcal{R})$, they should yield “best” solutions in $C(P_I, \mathcal{R}_J)$ (with all notations as stated in the text of Proposition 5.2.2). We name this property *continuity (of solutions)*:

Definition 5.2.2. Let F_c be as described in (5.10). F_c satisfies the continuity condition if the following holds:

Suppose $(P, \mathcal{R}) \in \mathcal{AP}$ with $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$, and assume $\alpha_1^+, \alpha_1^-, \dots, \alpha_n^+, \alpha_n^- \in wf(F_c(P, \mathcal{R}))$. Let $I \cup J$ be a partition of $\{1, \dots, n\}$, and set $P_I := P[\alpha_i^+, \alpha_i^-]_{i \in I}$, $\mathcal{R}_J := \{(B_j|A_j)[x_j]\}_{j \in J}$. Then $(\alpha_j^+, \alpha_j^-)_{j \in J} \in wf(F_c(P_I, \mathcal{R}_J))$, i.e. $F_c(P_I, \mathcal{R}_J) = P_I[\alpha_j^+, \alpha_j^-]_{j \in J} = F_c(P, \mathcal{R})$.

Finally, F_c should obey the principle of atomicity (cf. Theorem 5.1.1):

Definition 5.2.3. Let F_c be as described in (5.10). F_c satisfies the atomicity condition if for any $(P, \mathcal{R}), (P', \mathcal{R}^\Delta) \in \mathcal{AP}$ as in Theorem 5.1.1, $wf(F_c(P, \mathcal{R})) = wf(F_c(P', \mathcal{R}^\Delta))$.

The following proposition derives necessary conditions for a function F_c to fulfill the conditions of relevance, continuity and atomicity in special but important cases:

Proposition 5.2.3. Assume F_c as described in (5.10).

- (i) Suppose P_1, P_2 are positive distributions over two atoms a, b , with $P_1(b|a) = P_2(b|a)$, and let $\mathcal{R} = \{(b|a)[x]\}$, $x \in (0, 1)$. If F_c satisfies the relevance condition, then the weight factors α^+, α^- resp. β^+, β^- of $F_c(P_1, \mathcal{R})$ resp. $F_c(P_2, \mathcal{R})$ are equal in pairs, i.e. $\alpha^+ = \beta^+$ and $\alpha^- = \beta^-$.
- (ii) Suppose F_c satisfies the conditions of relevance, continuity and atomicity. Let $(P, \mathcal{R}) \in \mathcal{AP}$ with positive prior P such that no variable occurs both in antecedent and consequent of a conditional in \mathcal{R} and all assigned probabilities in \mathcal{R} are different from 0 and 1. Then the weight factors α^+, α^- associated in $F_c(P, \mathcal{R})$ with a conditional in \mathcal{R} only depend upon the probability x of this conditional and upon their quotient $\frac{\alpha^+}{\alpha^-}$. That means, for any $(P, \mathcal{R}), (P', \mathcal{R}') \in \mathcal{AP}$, P, P' positive,

$\mathcal{R} = \{(B|A)[x], (B_1|A_1)[x_1], \dots\}, \mathcal{R}' = \{(B'|A')[x], (B'_1|A'_1)[x'_1], \dots\}$, both sets finite, all of $x, x_i, x'_i \in (0, 1)$, no variable occurring both in antecedent and consequent of any conditional in \mathcal{R} and \mathcal{R}' , and for any weight factors α^+, α^- resp. α'^+, α'^- associated in $F_c(P, \mathcal{R})$ resp. $F_c(P', \mathcal{R}')$ with the conditional $(B|A)[x]$ resp. $(B'|A')[x]$,

$$\frac{\alpha^+}{\alpha^-} = \frac{\alpha'^+}{\alpha'^-} \quad \text{implies} \quad \alpha^+ = \alpha'^+ \quad \text{and} \quad \alpha^- = \alpha'^-.$$

Example 5.1.1 above shows that, in the cases dealt with by Proposition 5.2.3(i), all pairs α^+, α^- of weight factors have to fulfill $\frac{\alpha^+}{\alpha^-} = \frac{x}{1-x} \frac{P(a\bar{b})}{P(ab)}$. The cross ratio on the right hand side, depending only on prior and new conditional probabilities, represents exactly *relevant knowledge*. The left hand side is just the quotient of α^+ and α^- . This gives an intuitive reason for this quotient to play a key role, as it is stated in Proposition 5.2.3(ii).

Thus in the context of c-revisions, we identified clearly the parameters weight factors should be dependent on to give rise to a reasonable functional concept: $\frac{\alpha^+}{\alpha^-}$ and (the probability) x incorporate all relevant knowledge for the weight factors. Thus a reasonable functional concept for c-revisions may be realized by setting

$$\alpha^+ = F^+(x, \alpha) \quad \text{and} \quad \alpha^- = F^-(x, \alpha), \quad (5.11)$$

with two real positive functions F^+ and F^- , defined on $(0, 1) \times \mathbb{R}^+$ and related by $\frac{F^+(x, \alpha)}{F^-(x, \alpha)} = \alpha$, i.e.

$$F^+(x, \alpha) = \alpha F^-(x, \alpha). \quad (5.12)$$

As our global function $F : \mathcal{AP} \ni (P, \mathcal{R}) \mapsto P^*$ is to work for arbitrary P and \mathcal{R} , the functions F^+ and F^- are assumed to be independent of the prior and new information actually present, thus representing a fundamental inference pattern. Moreover, to yield “smooth” inferences we assume them to be continuous on $(0, 1) \times \mathbb{R}^+$. The functional concept, designed so far, should also be applied to the extreme probabilities $x \in \{0, 1\}$, incorporating classical logic as a limit case by assuming

$$F^+(0, 0) := \lim_{\substack{x \rightarrow 0 \\ \alpha \rightarrow 0}} F^+(x, \alpha) = 0, \quad F^+(1, \infty) := \lim_{\substack{x \rightarrow 1 \\ \alpha \rightarrow \infty}} F^+(x, \alpha) \in \mathbb{R}^+, \quad (5.13)$$

and

$$F^-(0, 0) := \lim_{\substack{x \rightarrow 0 \\ \alpha \rightarrow 0}} F^-(x, \alpha) \in \mathbb{R}^+, \quad F^-(1, \infty) := \lim_{\substack{x \rightarrow 1 \\ \alpha \rightarrow \infty}} F^-(x, \alpha) = 0, \quad (5.14)$$

in accordance with (5.2). The resulting posterior distribution P_F^* is a c-revision of the form

$$P_F^*(\omega) = \alpha_0 P(\omega) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} F^+(x_i, \alpha_i) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \overline{B_i}}} F^-(x_i, \alpha_i) \quad (5.15)$$

with non-negative extended real numbers $\alpha_1, \dots, \alpha_n \in \mathbb{R}^+ \cup \{0, \infty\}$ solving the n equations

$$\alpha_i = \begin{cases} \frac{x_i \sum_{\omega \models A_i \overline{B_i}} P(\omega) \prod_{\substack{j \neq i \\ \omega \models A_j B_j}} F^+(x_j, \alpha_j) \prod_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} F^-(x_j, \alpha_j)}{1 - x_i \sum_{\omega \models A_i B_i} P(\omega) \prod_{\substack{j \neq i \\ \omega \models A_j B_j}} F^+(x_j, \alpha_j) \prod_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} F^-(x_j, \alpha_j)}, & x_i \neq 0, 1 \\ 0, & x_i = 0 \\ \infty, & x_i = 1 \end{cases} \quad (5.16)$$

(see (5.1)). Note that $\alpha_i \in \mathbb{R}^+$ for $x_i \in (0, 1)$, because of the positivity of both functions F^+ and F^- and due to the P -consistency of \mathcal{R} . So the positivity condition (5.2) is satisfied, and (5.16) corresponds to the adjustment condition (5.3) here. Thus, for any n non-negative extended real numbers $\alpha_1, \dots, \alpha_n \in \mathbb{R}^+ \cup \{0, \infty\}$, $\alpha_1, \dots, \alpha_n$ is a solution of (5.16) iff $F^+(x_1, \alpha_1), F^-(x_1, \alpha_1), \dots, F^+(x_n, \alpha_n), F^-(x_n, \alpha_n)$ is a solution of (5.3) satisfying (5.2).

We summarize these remarks for the axiomatization of the second postulate (P2):

Postulate (P2): functional concept for c-revisions

There is a function $F^* : \mathcal{AP} \ni (P, \mathcal{R}) \mapsto P_c^* \in C(P, \mathcal{R})$ that assigns to each adjustment problem $(P, \mathcal{R}) \in \mathcal{AP}$ a particular c-revision P_F^* satisfying $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$. F^* is given by two real positive and continuous functions F^+ and F^- defined on $(0, 1) \times \mathbb{R}^+$, fulfilling the conditions (5.13) and (5.14) and related by (5.12), in that $P_F^* = F^*(P, \mathcal{R}) =: P *_F \mathcal{R}$ has the form (5.15) with $\alpha_1, \dots, \alpha_n \in \mathbb{R}^+ \cup \{0, \infty\}$ solving (5.16).

Define for (fixed) F^*, F^+, F^- as in (P2) and for $(P, \mathcal{R}) \in \mathcal{AP}$, $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$,

$$WQ_F(P, \mathcal{R}) := \{(\alpha_1, \dots, \alpha_n) \in \mathbb{R}^+ \cup \{0, \infty\} \mid (\alpha_1, \dots, \alpha_n) \text{ solves (5.16)}\} \quad (5.17)$$

to be set of all *weight quotients* that belong to c-revisions

$$P[\alpha_1, \dots, \alpha_n]_F(\omega) := \alpha_0 P(\omega) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i B_i}} F^+(x_i, \alpha_i) \prod_{\substack{1 \leq i \leq n \\ \omega \models A_i \bar{B}_i}} F^-(x_i, \alpha_i).$$

So $(\alpha_1, \dots, \alpha_n) \in WQ_F(P, \mathcal{R})$ iff $(F^+(x_1, \alpha_1), F^-(x_1, \alpha_1), \dots, F^+(x_n, \alpha_n), F^-(x_n, \alpha_n)) \in WF(P, \mathcal{R})$.

For any $(P, \mathcal{R}) \in \mathcal{AP}$, $F^*(P, \mathcal{R}) = P[\alpha_1, \dots, \alpha_n]_F$ is described by a particular element $(\alpha_1, \dots, \alpha_n) \in WQ_F(P, \mathcal{R})$. This disagreeable dependence on a special yet unknown solution of (5.16) may be overcome by assuming that the functions F^+ and F^- fulfill the *condition of uniqueness*:

Definition 5.2.4. Let F^+ and F^- be functions as described in (P2). F^+ and F^- satisfy the uniqueness condition iff whenever $(P, \mathcal{R}) \in \mathcal{AP}$ and $(\alpha_1, \dots, \alpha_n), (\beta_1, \dots, \beta_n) \in WQ_F(P, \mathcal{R})$ it holds that

$$P[\alpha_1, \dots, \alpha_n]_F = P[\beta_1, \dots, \beta_n]_F.$$

So, if F^+ and F^- satisfy the uniqueness condition then F^* is determined by (5.16), i.e. $F^*(P, \mathcal{R}) = P[\alpha_1, \dots, \alpha_n]_F$ for any $(\alpha_1, \dots, \alpha_n) \in WQ_F(P, \mathcal{R})$.

The functional concept (P2), describing weight factors of c-revisions, was initiated by the conditions of relevance, continuity and atomicity. The uniqueness condition ensures recovering these properties from (P2):

Proposition 5.2.4. Let $F^* : \mathcal{AP} \ni (P, \mathcal{R}) \mapsto P_F^* \in C(P, \mathcal{R})$ be as in (P2) with associated functions F^+ and F^- satisfying the uniqueness condition. Then F^* fulfills the conditions of relevance, continuity and atomicity.

We will see in Section 5.5 that the condition of uniqueness will in fact be satisfied for the special functions F^+ and F^- that will be determined by (P2) together with (P3) and (P4) (cf. Proposition 5.5.1).

Note that the conciseness of (P2) is essentially due to making use of c-revisions. So the efforts we invested in developing this conditional-logical concept begin to pay, providing now an elegant functional concept.

After having put the functional dependencies in concrete terms we are now going to study which properties the functions F^+ and F^- should have to guarantee reasonable probabilistic inferences. To simplify notation, we will usually prefer the operational $P *_F \mathcal{R}$ to the functional $F^*(P, \mathcal{R})$, where $*_F$ is described by (P2).

5.3 Logical Coherence

Surely, the adaptation scheme (5.15) will be considered sound only if the resulting posterior distribution can be used *coherently* as a prior distribution

for further adaptations. In particular, if we first adjust P only to a subset $\mathcal{R}_1 \subseteq \mathcal{R}$, and then use this posterior distribution to perform another adaptation to the full conditional information \mathcal{R} , we should obtain the same distribution as if we adjusted P to \mathcal{R} in only one step.

We state this demand for logical coherence as Postulate (P3):

Postulate (P3): logical coherence

For any distribution P and any P -consistent sets $\mathcal{R}_1, \mathcal{R}_2 \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$, the (final) posterior distribution which arises from a two-step process of adjusting P first to \mathcal{R}_1 and then adjusting this intermediate posterior to $\mathcal{R}_1 \cup \mathcal{R}_2$ is identical to the distribution resulting from directly revising P by $\mathcal{R}_1 \cup \mathcal{R}_2$.

More formally, the operator $*_F$ satisfies (P3) iff the following equation holds:

$$P *_F (\mathcal{R}_1 \cup \mathcal{R}_2) = (P *_F \mathcal{R}_1) *_F (\mathcal{R}_1 \cup \mathcal{R}_2). \quad (5.18)$$

Applying the principle of logical coherence will give us a further important result in determining the functions F^+ and F^- :

Theorem 5.3.1. *If the revision operator $*_F$ satisfies the postulate (P3) of logical coherence then*

$$F^-(0, 0) = F^+(1, \infty) = 1, \quad (5.19)$$

and F^- necessarily fulfills the functional equation

$$F^-(x, \alpha\beta) = F^-(x, \alpha) F^-(x, \beta) \quad (5.20)$$

for all $x \in (0, 1), \alpha, \beta \in \mathbb{R}^+$.

Because of (5.12), F^+ satisfies $F^+(x, \alpha\beta) = F^+(x, \alpha)F^+(x, \beta)$ iff F^- satisfies (5.20).

Theorem 5.3.1 is proved by checking conditions (5.19) and (5.20) in the very special case that P is a positive distribution over three variables a, b, c , and $\mathcal{R}_1, \mathcal{R}_2$ are given by $\mathcal{R}_1 = \{(c|a)[x]\}$ and $\mathcal{R}_2 = \{(c|b)[y]\}$. These conditions are necessary to guarantee a logical consistent behavior of the revision process for this example, and because we assumed the functions F^+, F^- to be independent of the actual case we thus proved the general validity of (5.19) and (5.20). In fact, there is little arbitrariness in choosing this special example which such a crucial meaning is assigned to. The way in which two conditionals with common conclusion should interact is one of the main issues in conditional logic and refers to the *antecedent conjunction problem* (cf.

[Nut80, Spi91]). The validity of (5.19) and (5.20) ensures a sound probabilistic treatment of this problem.

The functional equation (5.20) restricts the type of the function F^- (and that of F^+ , too,) essentially:

Proposition 5.3.1. *Let $*_F$ be a revision operator following (P2) such that F^+ and F^- satisfy (5.19) and (5.20). Then there is a (continuous) real function $c(x), x \in (0, 1)$, with*

$$\lim_{x \rightarrow 0} c(x) = 0 \quad \text{and} \quad \lim_{x \rightarrow 1} c(x) = -1 \quad (5.21)$$

such that

$$F^-(x, \alpha) = \alpha^{c(x)} \quad \text{and} \quad F^+(x, \alpha) = \alpha^{c(x)+1} \quad (5.22)$$

for any positive real α and for any $x \in (0, 1)$. Especially for $\alpha = 1$, this implies

$$F^-(x, 1) = F^+(x, 1) = 1. \quad (5.23)$$

In Theorem 5.3.1 we showed how the coherence property (5.18) determines the part the quotients α_i have to play in the revision process. We are now left with the investigation of the isolated impact of the numbers x_i which represent posterior conditional probabilities.

5.4 Representation Invariance

So far, we have largely neglected how (conditional) knowledge is represented in \mathcal{R} . Indeed, the principle of atomicity deals with *logical equivalence* of propositional formulas, but what about *probabilistic equivalences*, i.e. equivalences that are due to elementary probability calculus (see Section 1.3.2)? For instance, the sets of rules $\{(B|A)[x], A[y]\}$ and $\{AB[xy], A[y]\}$ are equivalent in this respect because each rule in one set is derivable from the rules in the other set. We surely expect the result of the revision process to be independent of the syntactic representation of probabilistic knowledge in \mathcal{R} :

Postulate (P4): representation invariance

If two P -consistent sets of probabilistic conditionals \mathcal{R} and \mathcal{R}' are probabilistically equivalent then the posterior distributions $P * \mathcal{R}$ and $P * \mathcal{R}'$ resulting from adapting the prior P to \mathcal{R} and to \mathcal{R}' , respectively, are identical.

The notion of probabilistic equivalence used here completely corresponds to that introduced in [PV90]. Using the operational notation, we are able to express (P4) more formally:

The revision operator $*_F$ satisfies (P4) iff

$$P *_F \mathcal{R} = P *_F \mathcal{R}' \quad (5.24)$$

for any two P -consistent and probabilistically equivalent sets $\mathcal{R}, \mathcal{R}' \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$.

The demand for independence of syntactic representation of probabilistic knowledge (P4) gives rise to two functional equations for $c(x)$ (cf. Proposition 5.3.1):

Proposition 5.4.1. *Let the functions F^- and F^+ describing the revision operator $*_F$ in (P2) satisfy (5.22) with a continuous function $c(x)$ fulfilling (5.21). If $*_F$ satisfies postulate (P4) resp. (5.24) then for all real $x, x_1, x_2 \in (0, 1)$ the following equations hold:*

$$c(x) + c(1 - x) = -1 \quad (5.25)$$

$$c(xx_1 + (1 - x)x_2) = -c(x)c(x_1) - c(1 - x)c(x_2) \quad (5.26)$$

The most obvious probabilistic equivalence is that of each two rules $(B|A)[x]$ and $(\bar{B}|A)[1 - x]$. This implies (5.25). Equation (5.26) is again proved by investigating a special but crucial revision problem. The relation

$$P(b|a) = P(b|ac)P(c|a) + P(b|a\bar{c})P(\bar{c}|a)$$

for arbitrarily chosen propositional variables a, b, c is fundamental to probabilistic conditionals, yielding the probabilistic equivalence of the two sets

$$\begin{aligned} \mathcal{R} &= \{(c|a)[x], (b|ac)[x_1], (b|a\bar{c})[x_2]\}, \\ \text{and } \mathcal{R}' &= \{(b|a)[y], (b|ac)[x_1], (b|a\bar{c})[x_2]\} \end{aligned}$$

with $y = xx_1 + (1 - x)x_2$ for real $x, x_1, x_2 \in (0, 1)$. The validity of (5.24) in this case necessarily implies (5.26) (see Appendix).

As a consequence of (5.25) and (5.26), we finally obtain:

Theorem 5.4.1. *If the operator $*_F$ in (P2) is to meet the demands for logical coherence (P3) and for representation invariance (P4), then F^+ and F^- necessarily have the forms*

$$F^+(x, \alpha) = \alpha^{1-x} \quad \text{and} \quad F^-(x, \alpha) = \alpha^{-x}, \quad (5.27)$$

respectively.

The continuity of the functions F^+ and F^- is essential to establish this theorem (see Appendix). This means in particular the continuous integrating of the extreme probabilities 0 and 1, that is, the seamless encompassing of classical logic.

5.5 Uniqueness and the Main Theorem

So far we have proved how the demands for logical coherence and for representation invariance constrain the functions which we assumed to underly the adjusting of P to \mathcal{R} , as described by the functional concept (P2). Applying (5.27) to (5.15) and (5.16) we recognize that the posterior distribution necessarily is of the same type (5.5) as the ME-distribution if it is to yield sound and coherent inferences.

Therefore, we nearly have reached our goal. But one step is still missing: Is this enough to *characterize* ME-inference within a conditional-logical framework? Are there possibly several different solutions of type (5.5), only one of which is the ME-distribution? And moreover, if we assume the functions F^+ and F^- to fulfill (5.27), is this *sufficient* to guarantee that the resulting operator $*_F$ satisfies logical coherence and representation invariance?

The question of uniqueness of the posterior distribution is at the center of all these problems. If it can be answered positively, we will have finished: The unique posterior distribution of type (5.5) must be the ME-distribution, $*_F$ then corresponds to ME-inference, and ME-inference is known to fulfill (5.18) and (5.24) as well as many other reasonable properties, cf. [PV90, SJ80, SJ81]. Moreover, together with (P2) uniqueness implies the conditions of relevance, atomicity and continuity (cf. Proposition 5.2.4).

The uniqueness of the solution $(\alpha_1, \dots, \alpha_n)$ of the fixpoint equation (5.6) is not clear at all. Imagine the case that the set \mathcal{R} representing new conditional knowledge contains twice the same rule in different notations. All that can be expected at best is a uniqueness statement for the product $\alpha_i \alpha_j$ of the corresponding factors. Even if we exclude such pathological cases, (5.6) is not easy to deal with at all.

But remember that we are primarily interested in the uniqueness of the posterior distribution, not in that of the solutions to (5.6). And indeed, this uniqueness is affirmed by the next theorem. In its proof, we will make use of cross-entropy as an excellently fitting measure of distance for distributions of type (5.5).

Proposition 5.5.1. *There is at most one solution of the adaptation problem ($*_{\text{prob}}$) (see page 74) of type (5.5), i.e. the functional concept defined by (5.27) satisfies the uniqueness condition.*

The following theorem summarizes our results in characterizing ME-adjustment within a conditional-logical framework:

Theorem 5.5.1 (Main Theorem). *Let $*_{\text{ME}}$ denote the ME-revision operator, that is, $*_{\text{ME}}$ assigns to a prior distribution P and some P -consistent*

set $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$ of probabilistic conditionals the one distribution $P_{ME} = P *_{ME} \mathcal{R}$ which has minimal cross-entropy with respect to P among all distributions that satisfy \mathcal{R} .

Then $*_{ME}$ yields the only adaptation of P to \mathcal{R} that obeys the principle of conditional preservation (P1), realizes a functional concept (P2) and satisfies the postulates for logical coherence (P3) and for representation invariance (P4). $*_{ME}$ is completely described by (5.5) and (5.6).

In this way, starting from a conditional-logical point of view we found a new characterization of the ME-solution to the problem

Given a distribution P and a P -consistent set of probabilistic conditionals \mathcal{R} , which way is the best for revising P by \mathcal{R} ?

The thorough embedding of the problem within the conditional-logical framework presented here conveys a clear understanding what actually makes the ME-distribution to be the best choice – ME-inference and probabilistic conditionals fit perfectly well.

In concluding this chapter, let us summarize the steps we have taken to obtain our conditional-logic characterization of ME-inference: The principle of conditional preservation was used to determine the structure of the posterior distribution. Then we assumed that a (continuous) functional concept extending classical logic should underly the adjustment process, and we isolated the crucial parameters which this concept should depend on. It was represented by means of two functions F^+ and F^- accomplishing the discrimination between these elementary events satisfying the antecedent of a rule that also satisfy its conclusion and those events that do not. So they constitute the decisive components for the extent of distortion the prior distribution is to be exposed to under adjustment.

Only two further preconditions were necessary to arrive at the desired characterization: logical coherence and independence of syntactical representation. While the latter postulate is usually considered to be fundamental to any reasonable inference procedure, the former one introduces a new aspect to reasoning, comparing inferences based on different theories, or epistemic states, respectively. This will be pursued in a more general framework later on, see Section 7.1.

Note that no exceptional demands had to be made, and the ME-solution arose in a rather natural way. Moreover, the proof of Proposition 5.5.1, which states the uniqueness of the solution, illustrates how perfectly well the approach presented here realizes ME-inference in an understandable manner, without imposing any external and abstract minimality demand. Actually, the proper idea of minimality is being made explicit by the four postulates.

6. Reasoning at Optimum Entropy

Now that the ME-principles have been motivated and characterized in detail it is time to ask how reasoning at optimum entropy actually works. We will throw light on this question from two different sides.

The first part of this chapter (cf. also [KI98b]) investigates how inferring at optimum entropy fits the formal framework for nonmonotonic reasoning of [Mak94]. In particular, we show that any inference operation based on ME-reasoning is *cumulative* and satisfies *loop*. Moreover, it turns out to be *supraclassical* with respect to *standard probabilistic consequence* which obviously generalizes classical consequence within a probabilistic setting (cf. Section 6.1). We also focus on the relationships between nonmonotonic and conditional ME-reasoning. Once more, it becomes obvious that material implication and conditionals differ substantially. To make the differences clear, we extend the conditional probabilistic language we are working in so as to contain probabilistic formulas corresponding to material implication, too. We show that conditionalization in its usual sense relates to material implication, whereas the connections between nonmonotonic reasoning and conditionals are more complex.

Though the ME-methods genuinely manipulate knowledge of a numerical nature, ME-reasoning is not easily understood by observing the probabilities in change. ME-logic is not truth-functional, as fuzzy logic is, nor is its aim to raise or to lower probabilities, as in the framework of upper and lower probabilities, and there is no straightforward calculation algorithm, as for Bayesian networks. ME-inferring rather makes use of the intensional structures of probabilistic knowledge (cf. [Par94, SJ80, KI98a]), so it seems to be better classified and appreciated by describing its formal properties as a nonmonotonic inference operation.

Nevertheless, some examples and practical inference schemes in simple but typical cases are important to illustrate ME-inference beyond formal results; they will be presented in the second part of this chapter (see also [KI97b]). The representation of the ME-distribution central to the argumentation in Chapter 5 (see equations (5.5), (5.6) and (5.7), page 76) then turns out to be not only of theoretical but also of practical use, allowing us to calculate

ME-probability values explicitly. For instance, we will show how knowledge is propagated transitively, and we will deal with *cautious cut* and *cautious monotonicity*. These inference schemes, however, are global, not local, i.e. all knowledge available has to be taken into account in their premises to give correct results. But they provide useful insights into the practice of ME-reasoning.

6.1 Probabilistic Consequence and Probabilistic Inference

The inference operation Cn of classical logic satisfies the following characteristic conditions:

- (i) *inclusion or reflexivity*: $X \subseteq Cn(X)$;
- (ii) *idempotence*: $Cn(X) = Cn(Cn(X))$;
- (iii) *cut*: $X \subseteq Y \subseteq Cn(X)$ implies $Cn(Y) \subseteq Cn(X)$;
- (iv) *monotonicity*: $X \subseteq Y$ implies $Cn(X) \subseteq Cn(Y)$,

where $X, Y \subseteq \mathcal{L}$.

So Cn is a *consequence operation*, in the sense of Tarski. Its rigidity, however, appreciated in mathematics and establishing its fundamental meaning for logical deduction, makes it a poor candidate to represent commonsense reasoning – in particular, monotonicity prohibits defeasible conclusions (cf. [DP96]). Nevertheless, many nonstandard logics have elements of classical logic built in in one way or another, and Makinson [Mak94] and Kraus, Lehmann, Magidor [KLM90] make use of it as a reference point to formulate their principles for nonmonotonic reasoning.

Within a probabilistic framework, the classical consequence operation has to be modified appropriately to provide a suitable tool for the description of formal inferences.

Semantically, Cn may be described by $Cn(X) = \{A \in \mathcal{L} \mid X \models A\}$, where \models demands taking account of *every* model of X . So we define the *standard probabilistic consequence operation* $Cn_{prob} : 2^{(\mathcal{L}|\mathcal{L})^{prob}} \rightarrow 2^{(\mathcal{L}|\mathcal{L})^{prob}}$ by virtue of

$$Cn_{prob}(\mathcal{R}) = \{\phi \in (\mathcal{L}|\mathcal{L})^{prob} \mid P \models \phi \text{ for all } P \in Mod(\mathcal{R})\},$$

with associated *standard probabilistic consequence relation* $\models^{prob} \subseteq 2^{(\mathcal{L}|\mathcal{L})^{prob}} \times (\mathcal{L}|\mathcal{L})^{prob}$

$$\mathcal{R} \models^{prob} \phi \text{ iff } \phi \in Cn_{prob}(\mathcal{R}).$$

If $\mathcal{R} \subseteq (\mathcal{L}|\mathcal{L})^{prob}$ is inconsistent, then $Mod(\mathcal{R}) = \emptyset$, hence $Cn_{prob}(\mathcal{R}) = (\mathcal{L}|\mathcal{L})^{prob}$ for inconsistent \mathcal{R} . Paris and Vencovska [PV98, Par94] present

a sequent calculus for \models^{prob} . Because standard probabilistic consequence is based on considering all models it is a consequence relation in the sense of Tarski.

Proposition 6.1.1. *Cn_{prob} satisfies inclusion, idempotence, cut and monotonicity.*

Thus Cn_{prob} takes the part of a classical logic within probabilistic logics. In general, however, probabilistic models are ill-famed for their extreme flexibility, so an inference operation based on *all* possible models is much too conservative and will yield little more but quite trivial results. In analogy to [Mak94], we define *probabilistic inference operations* and *probabilistic inference relations*, respectively, to study nonclassical inferences:

Definition 6.1.1. *A probabilistic inference operation is an operation*

$$C : 2(\mathcal{L}|\mathcal{L})^{prob} \rightarrow 2(\mathcal{L}|\mathcal{L})^{prob}$$

on sets of probabilistic conditionals. Its associated probabilistic inference relation is given by $\vdash \subseteq 2(\mathcal{L}|\mathcal{L})^{prob} \times (\mathcal{L}|\mathcal{L})^{prob}$, $\mathcal{R} \vdash \phi$ iff $\phi \in C(\mathcal{R})$. C is called complete iff it specifies for each set $\mathcal{R} \subseteq (\mathcal{L}|\mathcal{L})^{prob}$ with $C(\mathcal{R}) \neq \emptyset$ and $C(\mathcal{R}) \neq (\mathcal{L}|\mathcal{L})^{prob}$, a unique distribution, i.e. iff there is a (unique) distribution $Q_{\mathcal{R}}$ such that $C(\mathcal{R}) = Th(Q_{\mathcal{R}})$.

For fixed prior P define the *ME-inference operation* $C_P^{ME} : 2(\mathcal{L}|\mathcal{L})^{prob} \rightarrow 2(\mathcal{L}|\mathcal{L})^{prob}$ by

$$C_P^{ME}(\mathcal{R}) = \begin{cases} Th(P *_{ME} \mathcal{R}) & \text{iff } \mathcal{R} \text{ is } P - \text{consistent} \\ (\mathcal{L}|\mathcal{L})^{prob} & \text{else} \end{cases} \quad (6.1)$$

where $*_{ME}$ is the ME-revision operator (cf. Section 2.5 and Chapter 5, in particular Theorem 5.5.1). Obviously, C_P^{ME} is a complete probabilistic inference operation. The corresponding *ME-inference relation* is denoted by \vdash_P^{ME} .

6.2 Basic Properties of the ME-Inference Operation

In this section, we list some properties of ME-inference which are particularly important in the framework of nonmonotonic reasoning and belief revision (see Section 2.1, page 13). Most of these properties are already well-known (cf. [PV90, SJ81]) or easily proved, respectively (see Appendix).

One property that has proved to be crucial for characterizing ME-inference is that of *logical coherence* (cf. Section 5.3):

$$(P *_{ME} \mathcal{R}) *_{ME} (\mathcal{R} \cup \mathcal{S}) = P *_{ME} (\mathcal{R} \cup \mathcal{S}) \quad (6.2)$$

for any sets $\mathcal{R}, \mathcal{S} \subseteq (\mathcal{L}|\mathcal{L})^{prob}$ of probabilistic conditionals.

Furthermore, the uniqueness of the solution of type (5.5) yields an easy corollary which meets fundamental demands for “good revisions”:

Proposition 6.2.1. $P *_ME \mathcal{R} = P$ if and only if $P \models \mathcal{R}$.

The next proposition shows that ME-inference satisfies important postulates for nonmonotonic inference operations:

Proposition 6.2.2. *The ME-inference operation C_P^{ME} satisfies*

1. *inclusion:* $\mathcal{R} \subseteq C_P^{ME}(\mathcal{R})$.
2. *idempotence:* $C_P^{ME}(C_P^{ME}(\mathcal{R})) = C_P^{ME}(\mathcal{R})$.
3. *cumulativity:* $\mathcal{R} \subseteq \mathcal{S} \subseteq C_P^{ME}(\mathcal{R})$ implies $C_P^{ME}(\mathcal{R}) = C_P^{ME}(\mathcal{S})$.

As can be seen at once, cumulativity is equivalent to

$$\text{If } P *_ME \mathcal{R} \models \mathcal{S} \text{ then } P *_ME (\mathcal{R} \cup \mathcal{S}) = P *_ME \mathcal{R}$$

which is stated as Principle 5 in [PV90].

In the works of Makinson [Mak94] and Kraus, Lehmann and Magidor [KLM90], cumulativity takes a central position as a very fundamental property of inference processes: cumulativity ensures the inference process to be stable, in that adding of derivable knowledge does not alter the set of nonmonotonic conclusions. As stated above (following [Mak94]), cumulativity is a *pure* condition, without any reference to logical connectives in the underlying language and thus being applicable to probabilistic logics in a straightforward way. Besides the numerous postulates for nonmonotonic inference relations in [Mak94] and [KLM90] that are based on classical structures, there is another interesting pure condition called *loop* (see (2.2) in Section 2.1, page 13):

Proposition 6.2.3. C_P^{ME} satisfies the loop-property:

If $\mathcal{R}_1, \dots, \mathcal{R}_m \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$ with

$$\mathcal{R}_{i+1} \subseteq C_P^{ME}(\mathcal{R}_i) \quad \text{for } 1 \leq i \leq m-1 \quad \text{and} \quad \mathcal{R}_m \subseteq C_P^{ME}(\mathcal{R}_1),$$

then $C_P^{ME}(\mathcal{R}_i) = C_P^{ME}(\mathcal{R}_j)$ for all $i, j = 1, \dots, m$.

ME-inference \vdash_P^{ME} is not transitive, but loop guarantees unambiguity for a sequence $\mathcal{R}_1 \vdash_P^{ME} \mathcal{R}_2 \vdash_P^{ME} \dots \vdash_P^{ME} \mathcal{R}_m \vdash_P^{ME} \mathcal{R}_1$ of nonmonotonic derivations.

In the rest of this section, we will investigate interrelationships between ME-inference and standard probabilistic consequence.

Proposition 6.2.4. C_P^{ME} is supraclassical, that is, $Cn_{prob}(\mathcal{R}) \subseteq C_P^{ME}(\mathcal{R})$ for all $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$.

The proof of this proposition is straightforward because $P *_ME \mathcal{R}$ either is a model of \mathcal{R} , or $C_P^{ME}(\mathcal{R}) = (\mathcal{L} \mid \mathcal{L})^{prob}$.

Supraclassicality implies further important properties: Because C_P^{ME} is cumulative and supraclassical, it also satisfies *full absorption*:

$$Cn_{prob}C_P^{ME} = C_P^{ME} = C_P^{ME}Cn_{prob}$$

So in particular (cf. [Mak94]), C_P^{ME} fulfills *right weakening*:

$$\mathcal{R} \subseteq C_P^{ME}(\mathcal{S}) \text{ implies } Cn_{prob}(\mathcal{R}) \subseteq C_P^{ME}(\mathcal{S})$$

and *left logical equivalence*:

$$\text{if } Cn_{prob}(\mathcal{R}) = Cn_{prob}(\mathcal{S}) \text{ then } C_P^{ME}(\mathcal{R}) = C_P^{ME}(\mathcal{S})$$

Therefore ME-inference respects standard deductional structures. But ME-inference fails *distribution*:

$$C_P^{ME}(\mathcal{R}) \cap C_P^{ME}(\mathcal{S}) \not\subseteq C_P^{ME}(Cn_{prob}(\mathcal{R}) \cap Cn_{prob}(\mathcal{S})).$$

This can easily be seen as follows: Consider the case $P = P_0(a, b, c)$, where $P_0(a, b, c)$ is the uniform distribution over three propositional variables a, b and c . Let $\mathcal{R} = \{(b|a)[x_1], (c|a)[x_2]\}$ and $\mathcal{S} = \{(b|a)[x_2], (c|a)[x_1]\}$ with $x_1 \neq x_2$. Then $(bc|a)[x_1x_2] \in C_P^{ME}(\mathcal{R}) \cap C_P^{ME}(\mathcal{S})$ (cf. Proposition 6.4.5 in Section 6.4), but $(bc|a)[x_1x_2] \notin C_P^{ME}(Cn_{prob}(\mathcal{R}) \cap Cn_{prob}(\mathcal{S}))$.

6.3 ME-Logic and Conditionals

In Chapter 5, it is shown that ME-reasoning may be characterized as a sound and consistent method for handling probabilistic conditionals. Thus the notion of conditionals may be considered fundamental to ME-logic, and the property of *conditionalization*

$$\text{If } \{A_i\} \cup \{B\} \vdash C \text{ then } \{A_i\} \vdash B \rightarrow C, \quad (6.3)$$

as stated in standard, non-probabilistic terms in [Mak94], is of special interest. Here \rightarrow is meant to represent material implication which has no straightforward counterpart in probabilistic logics. Namely, if ϕ, ψ are probabilistic formulas (facts or conditionals), then $\phi \rightarrow \psi$ cannot be taken as, for instance, $\neg\phi \vee \psi$ because neither negation nor disjunction are defined in $(\mathcal{L} \mid \mathcal{L})^{prob}$.

\rightarrow is also clearly different from the operator $(\cdot|\cdot)$ used for probabilistic conditionals, not only for syntactical differences. The satisfaction relation \models between a distribution P and a conditional $(B|A)[x]$ actually involves an adjustment process, namely shifting P to $P_A = P(\cdot|A)$, whereas material implication should be satisfied *within* P .

In the sequel, we will investigate conditionalization with respect to a suitably generalized material implication $A \rightarrow B[x]$ in probabilistics as well as with respect to the proper probabilistic conditionals of the form $(B|A)[x]$.

So for the moment, let us extend our language $(\mathcal{L} | \mathcal{L})^{prob}$ so as to also contain all formulas of the type $A \rightarrow B[x]$:

$$(\mathcal{L} | \mathcal{L})^{prob, \rightarrow} = (\mathcal{L} | \mathcal{L})^{prob} \cup \{A \rightarrow B[x] \mid A, B \in \mathcal{L}, x \in [0, 1]\}.$$

We modify C_P^{ME} slightly to take its values in $2^{(\mathcal{L}|\mathcal{L})^{prob, \rightarrow}}$:

$$C_P^{ME} : 2^{(\mathcal{L}|\mathcal{L})^{prob}} \rightarrow 2^{(\mathcal{L}|\mathcal{L})^{prob, \rightarrow}}.$$

In correspondence to classical material implication, and in compatibility with the use of the conditional operator $(\cdot|\cdot)$, we define the semantics of $A \rightarrow B[x]$ in a probabilistic setting as

$$P \models A \rightarrow B[x] \quad \text{iff} \quad P \models A[1] \text{ implies } P \models B[x]$$

for a probabilistic distribution P . Note that $P \models A \rightarrow B[x]$ is not equivalent to $P \models (\neg A \vee B)[x]$, i.e. to $P(\neg A \vee B) = x$. Rather we have $\models^{prob} A \rightarrow B[x]$ iff $Mod(A[1]) \subseteq Mod(B[x])$. So our probabilistic interpretation of \rightarrow generalizes what Adams called *strict consequence* in [Ada66, p. 274]. In particular, $\models^{prob} A \rightarrow B[x]$ implies $P \models (B|A)[x]$ for all distributions P with $P(A) \neq 0$.

Using this notion of (probabilistic) material implication, ME-inference satisfies conditionalization in the sense of (6.3):

Proposition 6.3.1. *Let P be a distribution. Whenever $B[x] \in C_P^{ME}(\mathcal{R} \cup \{A[1]\})$ then $A \rightarrow B[x] \in C_P^{ME}(\mathcal{R})$.*

But note that Proposition 6.3.1 is false for probabilistic conditionals: in general, $B[x] \in C_P^{ME}(\mathcal{R} \cup \{A[1]\})$ does not imply $(B|A)[x] \in C_P^{ME}(\mathcal{R})$. Instead, we have the following connections between ME-inference and conditional implication:

Proposition 6.3.2. *Let P be a distribution.*

For any $\mathcal{R} \subseteq (\mathcal{L} | \mathcal{L})^{prob}$, $(B|A)[x] \in (\mathcal{L} | \mathcal{L})^{prob}$, we have

1. $(P *_{ME} \{A[1]\}) = P_A = P(\cdot|A)$.
2. $P *_{ME} (\mathcal{R} \cup \{A[1]\}) = P_A *_{ME} \mathcal{R}$, i.e. $C_P^{ME}(\mathcal{R} \cup \{A[1]\}) = C_{P_A}^{ME}(\mathcal{R})$.

3. $(P *_{ME} \mathcal{R}) *_{ME} \{A[1]\} \models B[x] \text{ implies } P *_{ME} \mathcal{R} \models (B|A)[x].$

Proposition 6.3.2 shows that $P *_{ME} (\mathcal{R} \cup \{A[1]\})$ and $(P *_{ME} \mathcal{R}) *_{ME} \{A[1]\}$ differ in general, and that only the latter of these distributions reveals a behavior that is compatible to conditional implication.

Proposition 6.3.2(2) links up two inference operations based on different probability distributions in a special case. Such relationships are of crucial meaning for studying *iterated belief revision* and for investigating belief revision and nonmonotonic reasoning in a more general framework (cf. Chapters 2 and 7).

6.4 ME-Deduction Rules

We will now leave the abstract level of argumentation and turn to concrete inference patterns. Once more, it must be emphasized, however, that ME-infering is a *global*, not a *local* method: Only if *all* knowledge available is taken into account, the results of ME-inference are reliable to yield best expectation values. Thus it is not possible to use only partial information for reasoning, and then continue the process of adjusting from the obtained intermediate distribution with the information still left. It is important that in the two-step adjustment process $(P * \mathcal{R}_1) * (\mathcal{R}_1 \cup \mathcal{R}_2)$ dealt with in the coherence postulate (P3) (see Section 5.3) the second adaptation step uses full information $\mathcal{R}_1 \cup \mathcal{R}_2$. In fact, the distributions $(P *_{ME} \mathcal{R}_1) *_{ME} (\mathcal{R}_1 \cup \mathcal{R}_2)$ and $(P *_{ME} \mathcal{R}_1) *_{ME} \mathcal{R}_2$ differ in general.

For this reason, the deduction rules to be presented in the sequel do not provide a convenient (and complete) calculus for ME-reasoning. But they effectively illustrate the reasonableness of that technique by calculating explicitly inferred probabilities of rules in terms of given (or learned, respectively) probabilities. In contrast to this, the inference patterns for deriving lower and upper bounds for probabilities presented in [DPT90] and [TGK92] are local, but they are afflicted with all problems typical to methods for inferring intervals, not single values (cf. Chapter 5).

It must be pointed out that in principle, ME-reasoning is feasible for many consistent probabilistic representation and adaptation problems by iterative propagation. This is realized, for instance, by the probabilistic expert system shells SPIRIT [RM96] and PIT [SF97] far beyond the scope of the few inference patterns given below (also see Example 6.4.3; cf. Chapter 9).

We will use the following notation:

$$\frac{\mathcal{R} : (B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]}{(B_1^*|A_1^*)[x_1^*], \dots, (B_m^*|A_m^*)[x_m^*]}$$

iff $\mathcal{R} = \{(B_1|A_1)[x_1], \dots, (B_n|A_n)[x_n]\}$ and $P_0 *_{ME} \mathcal{R} \models \{(B_1^*|A_1^*)[x_1^*], \dots, (B_m^*|A_m^*)[x_m^*]\}$, where P_0 is a uniform distribution of suitable size.

6.4.1 Chaining Rules

Proposition 6.4.1 (Transitive Chaining). *Suppose a, b, c to be propositional variables, $x_1, x_2 \in [0, 1]$. Then*

$$\frac{\mathcal{R} : (b|a)[x_1], (c|b)[x_2]}{(c|a)[\frac{1}{2}(2x_1x_2 + 1 - x_1)]} \quad (6.4)$$

Example 6.4.1. Suppose the propositional variables a, b, c are given the meanings $a=$ Being young, $b=$ Being single, and $c=$ Having children, respectively. We know (or assume) that young people are usually singles (with probability 0.9) and that mostly, singles do not have children (with probability 0.85), so that $\mathcal{R} = (b|a)[0.9], (\bar{c}|b)[0.85]$. Using (6.4) with $x_1 = 0.9$ and $x_2 = 0.85$, ME-reasoning yields $(\bar{c}|a)[0.815]$ (the negation of c makes no difference). Therefore from the knowledge stated by \mathcal{R} we may conclude that the probability of an individual not having children if (s)he is young is best estimated by 0.815.

In many cases, however, rules must not be simply connected transitively as in Proposition 6.4.1 because definite exceptions are present. Let us consider the famous "Tweety the penguin"-example.

Example 6.4.2. Most birds fly, i.e. $(fly|bird)[x_1]$ with a probability x_1 between 0.5 and 1, penguins are definitely birds, $(bird|penguin)[1]$, but no one has ever seen a flying penguin, so $(fly|penguin)[x_2]$ with a probability x_2 very close to 0. What may be inferred about Tweety who is known to be a bird and a penguin?

The crucial point in this example is that two pieces of evidence apply to Tweety, one being more specific than the other. The next proposition shows that ME-reasoning is able to cope with *categorical specificity*.

Proposition 6.4.2 (Categorical Specificity). *Suppose a, b, c to be propositional variables, $x_1, x_2 \in [0, 1]$. Then*

$$\frac{\mathcal{R} : (b|a)[x_1], (b|c)[x_2], (a|c)[1]}{(b|ac)[x_2]} \quad (6.5)$$

Actually, (6.5) is a general probabilistic deduction scheme, not only due to ME-reasoning: Specific information dominates more general information,

if the specificity relation $(a|c)[1]$ is categorical. This is an easy probabilistic calculation. Proposition 6.4.2, however, is proved here by using the α -factors of the corresponding ME-distribution to illustrate the interdependencies between the three conditionals. If the probability of the specificity conditional $(a|c)$ lies somewhere in between 0 and 1, the equational system determining the α_i 's becomes more complicated. But it can be solved by iteration, e.g. by the aid of SPIRIT (cf. [RM96]), if the conditional probabilities involved are numerically specified (cf. Example 6.4.3 below). Within a qualitative probabilistic context, Adam's ϵ -semantics [Ada75] presents a method to handle exceptions and to take account of subclass specificity. Goldszmidt, Morris and Pearl [GMP90] showed how reasoning based on infinitesimal probabilities may be improved by using ME-principles.

Example 6.4.3. A knowledge base is to be built up representing “Typically, students are adults”, “Usually, adults are employed” and “Mostly, students are not employed” with probabilistic degrees of uncertainty $0.99(< 1)$, 0.8 and 0.9 , respectively. Let a, s, e denote the propositional variables $a =$ Being an Adult, $s =$ Being a Student, and $e =$ Being Employed. The quantified conditional information may be written as $\mathcal{R} = \{(a|s)[0.99], (e|a)[0.8], (\bar{e}|s)[0.9]\}$. From this, SPIRIT (cf. [RM96]) calculates $P^*(\bar{e}|as) = 0.8991 \approx 0.9$. So the more specific information s dominates a clearly, but not completely.

Thus ME-inference solves in an elegant way the problem of conflicting evidence. Specific information dominates more general knowledge by virtue of the inherent mechanisms, without any external preferential or hierarchical structures as in [KLM90, Bre89], and without rankings as in [Gef92, GP92]. The weight of a rule is encoded by its conditional-logical structure and its probability, its interactions with other rules being given implicitly. It is only the application of the ME-principle which combines the probabilistic rules to yield inferences, thus allowing a convenient modularity of knowledge representation.

6.4.2 Cautious Monotonicity and Cautious Cut

Obviously, ME-inference acts nonmonotonically: conjoining the antecedent of a conditional with a further literal may alter the probability of the conditional dramatically (cf. Example 6.4.3). But a weak form of monotonicity is reasonable and can indeed be proved:

Proposition 6.4.3 (Cautious Monotonicity). *Suppose a, b, c to be propositional variables, $x_1, x_2 \in [0, 1]$. Then*

$$\frac{\mathcal{R} : (b|a)[x_1], (c|a)[x_2]}{(c|ab)[x_2]} \quad (6.6)$$

(6.6) illustrates how ME-propagation respects conditional independence (cf. [SJ81]): $P^*(c|ab) = P^*(c|a) = x_2$ if $(b|a)[x_1]$, $(c|a)[x_2]$ is the only available knowledge.

The cautious monotonicity inference rule deals with *adding* information to the antecedent. Another important case arises if literals in the antecedent have to be *deleted*. Of course we cannot expect the classical cut rule to hold. But, as in the case of monotonicity, a *cautious cut rule* may be proved:

Proposition 6.4.4 (Cautious Cut). *Suppose a, b, c to be propositional variables, $x_1, x_2 \in [0, 1]$. Then*

$$\frac{\mathcal{R} : (c|ab)[x_1], (b|a)[x_2]}{(c|a)[\frac{1}{2}(2x_1x_2 + 1 - x_2)]} \quad (6.7)$$

(6.7) is *cautious* in that the probability of $(c|a)$ is a (simple) function of the probabilities assigned to $(c|ab)$ and $(b|a)$. By observing the equivalence $(b|a) \equiv (ab|a)$, (6.7) may be taken as an immediate consequence of the transitive chaining rule (6.4).

6.4.3 Conjoining Literals in Antecedent and Consequent

The following deduction schemes deal with various cases of inferring probabilistic conditionals with literals in antecedents and consequents being conjoined. Three of them – *Conjunction Left*, *Conjunction Right*, (ii) and (iii) – are treated in [TGK92] under similar names, thus allowing a direct comparison of ME-inference to probabilistic local bounds propagation. Cautious monotonicity (6.6) may be found in that paper, too, where it is denoted as *Weak Conjunction Left*. We will omit the straightforward proofs.

Proposition 6.4.5 (Conjunction Right). *Suppose a, b, c to be propositional variables, $x_1, x_2 \in [0, 1]$. Then the following ME-inference rules hold:*

$$\begin{aligned} (i) \quad & \frac{\mathcal{R} : (b|a)[x_1], (c|a)[x_2]}{(bc|a)[x_1x_2]} \\ (ii) \quad & \frac{\mathcal{R} : (b|a)[x_1], (c|ab)[x_2]}{(bc|a)[x_1x_2]} \\ (iii) \quad & \frac{\mathcal{R} : (b|a)[x_1], (c|b)[x_2]}{(bc|a)[x_1x_2]} \end{aligned}$$

Proposition 6.4.6 (Conjunction Left). *Suppose a, b, c to be propositional variables, $x_1, x_2 \in [0, 1]$. Then*

$$\frac{\mathcal{R} : (b|a)[x_1], (bc|a)[x_2]}{(c|ab)\left[\frac{x_2}{x_1}\right]}$$

6.4.4 Reasoning by Cases

The last inference scheme presented in this section will show how probabilistic information obtained by considering exclusive cases is being processed at maximum entropy:

Proposition 6.4.7 (Reasoning by cases). *Suppose a, b, c to be propositional variables, $x_1, x_2 \in [0, 1]$. Then*

$$\frac{\mathcal{R} : (c|ab)[x_1], (c|a\bar{b})[x_2]}{(c|a)\left[x_1\left(1 + \frac{x_1^{x_1}(1-x_1)^{1-x_1}}{x_2^{x_2}(1-x_2)^{1-x_2}}\right)^{-1} + x_2\left(1 + \frac{x_2^{x_2}(1-x_2)^{1-x_2}}{x_1^{x_1}(1-x_1)^{1-x_1}}\right)^{-1}\right]}$$

To date, no deduction scheme is known for the interesting *antecedent conjunction problem*

$$\frac{\mathcal{R} : (c|a)[x_1], (c|b)[x_2]}{(c|ab)[??]}$$

(cf. [Nut80, Spi91]). Such a scheme would reveal clearly how ME-inference combines evidences. Note that the consistent handling of the antecedent conjunction problem plays a crucial role for characterizing ME-inference (see Theorem 5.3.1, page 86, and the remarks following it).

7. Belief Revision and Nonmonotonic Reasoning – Revisited

Nonmonotonic reasoning and belief revision are closely related in that they both deal with reasoning under uncertainty and try to reveal sensible lines of reasoning in response to incoming information (cf. Chapter 2, in particular, Section 2.3). As we already pointed out, the crucial difference between both areas is the role of the knowledge base which is only implicit in nonmonotonic reasoning, but explicit and in fact in the focus of interest in belief revision. So the correspondences between axioms of belief change and properties of nonmonotonic inference operations are usually elaborated only in the case that revisions are based on a fixed theory (cf. [MG91]), and very little work has been done to incorporate iterated belief revision in that framework.

Within the context of ME-inference, iterated adjustments arise quite naturally and are dealt with in a satisfactory manner (cf. Equation (6.2) and Proposition 6.3.2). The crucial point here is that the ME-operator $*_{ME}$ actually is a full revision operator taking *two* entries, namely a distribution P on its left and a (compatible and consistent) set of conditionals on its right. Nonmonotonic reasoning and belief revision usually focus on handling its right entry, while considering its left entry – i.e. the theory inferences are based on – to be given.

In this chapter, we will exploit the relationships between nonmonotonic reasoning and belief revision further by considering epistemic states and sets of conditionals instead of theories and propositional beliefs. We will provide a more general framework that not only allows a more accurate representation of belief revision via nonmonotonic formalisms, but also gives, vice versa, an important impetus to handle iterated revisions. So we will generalize the notion of an inference operation and introduce *universal inference operations* as a suitable counterpart to (full) revision operators in nonmonotonic logics.

In particular, we will show that the property of *logical coherence* which was identified as one of the axioms for characterizing ME-inference (cf. Section 5.3) may be considered as a strong version of *cumulativity* for universal inference operations.

Leaving the classical framework also allows a more accurate view on iterated revision by differentiating between *simultaneous* and *successive revision*. The former will be seen to handle genuine revisions appropriately, while the latter may also model *updating*. This distinction is based on clearly separating background or generic knowledge from evidential or contextual knowledge, a feature that is listed in [DP96] as one of three basic requirements a *plausible exception-tolerant inference system* has to meet. Moreover, we will show that in a probabilistic framework, it is also possible to treat revision as different from *focusing* without giving up the assumption of having a single, distinguished probability distribution as base for inferences.

Parts of the results presented in this chapter were already published in [KI01] and [KI98b].

7.1 Universal Inference Operations

In Section 6.1, we called a probabilistic inference operation C *complete* if it specifies for each set $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$ with $C(\mathcal{R}) \neq \emptyset$, $C(\mathcal{R}) \neq (\mathcal{L} \mid \mathcal{L})^{prob}$, a unique distribution $Q_{\mathcal{R}}$ such that $C(\mathcal{R}) = Th(Q_{\mathcal{R}})$ (cf. Definition 6.1.1, page 93). That is to say that to each set of probabilistic conditionals yielding non-trivial inferences a suitable model is associated by which the corresponding inferences can be described. Hence we assumed the probabilistic inference operation C to be *model-based* (cf. [Her91, Thi89]).

This definition also makes sense in a more general framework dealing with epistemic states if we choose a suitable language $(\mathcal{L} \mid \mathcal{L})^*$. For instance, for probability distributions and ordinal conditional functions, we take $(\mathcal{L} \mid \mathcal{L})^* = (\mathcal{L} \mid \mathcal{L})^{prob}$ and $(\mathcal{L} \mid \mathcal{L})^{OCF}$, respectively, and in a purely qualitative setting, we assume $(\mathcal{L} \mid \mathcal{L})^* = (\mathcal{L} \mid \mathcal{L})$. In any case, \mathcal{L} is a propositional language over an alphabet \mathcal{V} . Let $\mathcal{E}^* = \mathcal{E}_{\mathcal{V}}^*$ denote the set of epistemic states using $(\mathcal{L} \mid \mathcal{L})^*$ for representation of (conditional) beliefs (see Section 3.1). For an epistemic state $\Psi \in \mathcal{E}^*$, we have

$$Th^*(\Psi) = \{\phi \in (\mathcal{L} \mid \mathcal{L})^* \mid \Psi \models \phi\}$$

which is assumed to describe Ψ uniquely, up to representation equivalence (cf. the *uniqueness assumption* (3.4), p. 31). So epistemic states are considered as models of sets of conditionals $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$:

$$Mod^*(\mathcal{R}) = \{\Psi \in \mathcal{E}^* \mid \Psi \models \mathcal{R}\}$$

This allows us to extend semantical entailment to sets of conditionals by setting

$$\mathcal{R}_1 \models^* \mathcal{R}_2 \quad \text{iff} \quad Mod^*(\mathcal{R}_1) \subseteq Mod^*(\mathcal{R}_2),$$

and to define a (monotonic) consequence operation $Cn^* : 2^{(\mathcal{L}|\mathcal{L})^*} \rightarrow 2^{(\mathcal{L}|\mathcal{L})^*}$ by

$$Cn^*(\mathcal{R}) = \{\phi \in (\mathcal{L}|\mathcal{L})^* \mid \mathcal{R} \models^* \phi\},$$

in analogy to classical consequence. Two sets of conditionals $\mathcal{R}_1, \mathcal{R}_2 \subseteq (\mathcal{L}|\mathcal{L})^*$ are called (*epistemically*) *equivalent* iff $Mod^*(\mathcal{R}_1) = Mod^*(\mathcal{R}_2)$.

In the sequel, we will consider (*conditional*) *inference operations*

$$C : 2^{(\mathcal{L}|\mathcal{L})^*} \rightarrow 2^{(\mathcal{L}|\mathcal{L})^*} \quad (7.1)$$

associating with each set of conditionals a set of inferred conditionals. This generalizes the notion of inference operations given in Section 2.1 since propositional facts may be considered as degenerated conditionals.

Definition 7.1.1. A conditional inference operation C is called *complete* iff it specifies for each set $\mathcal{R} \subseteq (\mathcal{L}|\mathcal{L})^*$ with $C(\mathcal{R}) \neq \emptyset, C(\mathcal{R}) \neq (\mathcal{L}|\mathcal{L})^*$, a complete epistemic state $\Psi_{\mathcal{R}}$, i.e. iff there is an epistemic state $\Psi_{\mathcal{R}}$ such that $C(\mathcal{R}) = Th^*(\Psi_{\mathcal{R}})$.

Definition 7.1.2. A universal inference operation \mathbf{C} assigns a complete conditional inference operation

$$C_{\Psi} : 2^{(\mathcal{L}|\mathcal{L})^*} \rightarrow 2^{(\mathcal{L}|\mathcal{L})^*}$$

to each epistemic state $\Psi \in \mathcal{E}^*$:

$$\mathbf{C} : \Psi \mapsto C_{\Psi}.$$

\mathbf{C} is said to be *reflexive* (idempotent, cumulative) iff all its involved inference operations have the corresponding property.

If $\mathbf{C} : \Psi \mapsto C_{\Psi}$ is a universal inference operation, C_{Ψ} is complete for each $\Psi \in \mathcal{E}^*$. That means, for each set $\mathcal{R} \subseteq (\mathcal{L}|\mathcal{L})^*$, $C_{\Psi}(\mathcal{R})$ is either \emptyset or $(\mathcal{L}|\mathcal{L})^*$, or it specifies completely an epistemic state $\Phi_{\Psi, \mathcal{R}}$:

$$C_{\Psi}(\mathcal{R}) = Th^*(\Phi_{\Psi, \mathcal{R}}) \quad (7.2)$$

Define the set of all such epistemic states by

$$\mathcal{E}^*(C_{\Psi}) = \{\Phi \in \mathcal{E}^* \mid \exists \mathcal{R} \subseteq (\mathcal{L}|\mathcal{L})^* : C_{\Psi}(\mathcal{R}) = Th^*(\Phi)\}.$$

Definition 7.1.3. A universal inference operation \mathbf{C} preserves consistency iff for each epistemic state $\Psi \in \mathcal{E}^*$ and for each consistent set $\mathcal{R} \subseteq (\mathcal{L}|\mathcal{L})^*$, $C_{\Psi}(\mathcal{R}) \neq \emptyset$ and $C_{\Psi}(\mathcal{R}) \neq (\mathcal{L}|\mathcal{L})^*$. In a quantitative setting, when Ψ is represented by a conditional valuation function V , we further presuppose \mathcal{R} to be V -consistent.

Here consistency of a set $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ means that there is an epistemic state representing \mathcal{R} . When the (prior) epistemic state is represented by a conditional valuation function V , \mathcal{R} is said to be V -consistent iff there is a conditional valuation function V' which is V -consistent and represents \mathcal{R} (see Definitions 4.5.1 and 5.1.1).

Definition 7.1.4. *A universal inference operation \mathbf{C} is founded iff for each epistemic state Ψ and for any $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$, $\Psi \models \mathcal{R}$ implies $C_\Psi(\mathcal{R}) = Th^*(\Psi)$.*

The property of foundedness establishes a close and intuitive relationship between an epistemic state Ψ and its associated inference operation C_Ψ , distinguishing Ψ as its stable starting point. In particular, if \mathbf{C} is founded then $C_\Psi(\emptyset) = Th^*(\Psi)$ (this property is called *faithfulness* in [KI01]).

As to the universal inference operation \mathbf{C} , foundedness ensures injectivity, as can easily be proved:

Proposition 7.1.1. *If \mathbf{C} is founded, then it is injective.*

In standard (i.e. one-dimensional) nonmonotonic reasoning, as it was developed in [Mak94] and [KLM90], cumulativity occupies a central and fundamental position, claiming the inferences of a set \mathcal{S} that “lies in between” another set \mathcal{R} and its nonmonotonic consequences $C(\mathcal{R})$ to coincide with $C(\mathcal{R})$ (cf. equation (2.1), p. 12).

To establish a similar well-behavedness of \mathbf{C} with respect to *epistemic states*, we introduce suitable relations to compare epistemic states with one another.

Definition 7.1.5. *Let $\mathbf{C} : \Psi \mapsto C_\Psi$ be a universal inference operation. For each epistemic state Ψ , define a relation \sqsubseteq_Ψ on $\mathcal{E}^*(C_\Psi)$ by setting*

$$\Phi_1 \sqsubseteq_\Psi \Phi_2$$

iff there are sets $\mathcal{R}_1 \subseteq \mathcal{R}_2 \subseteq (\mathcal{L} \mid \mathcal{L})^$ such that*

$$Th^*(\Phi_1) = C_\Psi(\mathcal{R}_1) \quad \text{and} \quad Th^*(\Phi_2) = C_\Psi(\mathcal{R}_2)$$

For founded universal inference operations, we have in particular $C_\Psi(\emptyset) = Th^*(\Psi)$ for all $\Psi \in \mathcal{E}^*$, so Ψ is a minimal element of $\mathcal{E}^*(C_\Psi)$ with respect to \sqsubseteq_Ψ :

Proposition 7.1.2. *If \mathbf{C} is founded, then for all $\Psi \in \mathcal{E}^*$ and for all $\Phi \in \mathcal{E}^*(C_\Psi)$, it holds that $\Psi \sqsubseteq_\Psi \Phi$.*

We will now generalize the notion of cumulativity for universal inference relations:

Definition 7.1.6. *A universal inference operation \mathbf{C} is called strongly cumulative iff for each $\Psi \in \mathcal{E}^*$ and for any epistemic states $\Phi_1, \Phi_2 \in \mathcal{E}^*(C_\Psi)$, $\Psi \sqsubseteq_\Psi \Phi_1 \sqsubseteq_\Psi \Phi_2$ implies: whenever $\mathcal{R}_1 \subseteq \mathcal{R}_2 \subseteq (\mathcal{L} \mid \mathcal{L})^*$ such that $Th^*(\Phi_1) = C_\Psi(\mathcal{R}_1)$ and $Th^*(\Phi_2) = C_\Psi(\mathcal{R}_2)$, then $Th^*(\Phi_2) = C_\Psi(\mathcal{R}_2) = C_{\Phi_1}(\mathcal{R}_2)$.*

Strong cumulativity describes a relationship between inference operations based on different epistemic states, thus linking up the inference operations of \mathbf{C} . In the definition above, Φ_1 is an epistemic state intermediate between Ψ and Φ_2 , with respect to the relation \sqsubseteq_Ψ , and strong cumulativity claims that the inferences based on Φ_1 coincide with the inferences based on Ψ within the scope of Φ_2 .

The next proposition is immediate:

Proposition 7.1.3. *Let \mathbf{C} be a universal inference operation which is strongly cumulative. Suppose $\Psi \in \mathcal{E}^*$, $\Phi \in \mathcal{E}^*(C_\Psi)$ such that $Th^*(\Phi) = C_\Psi(\mathcal{R})$, $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$. Then*

$$C_\Psi(\mathcal{S}) = C_\Phi(\mathcal{S})$$

for any $\mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})^$ with $\mathcal{R} \subseteq \mathcal{S}$.*

The following theorem justifies the name “strong cumulativity”: It states that strong cumulativity actually generalizes cumulativity for an important class of universal inference operations:

Theorem 7.1.1. *If \mathbf{C} is founded, then strong cumulativity implies cumulativity.*

Universal inference operations $\mathbf{C} : \Psi \mapsto C_\Psi$ are a proper counterpart of revision operators $*$ in nonmonotonic reasoning by virtue of setting

$$\Psi * \mathcal{R} = \Phi_{\Psi, \mathcal{R}}$$

(cf. (7.2) above) for $\Psi \in \mathcal{E}^*$, $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$, that is

$$C_\Psi(\mathcal{R}) = Th(\Psi * \mathcal{R}) \tag{7.3}$$

Using this notation and the uniqueness assumption (3.4), foundedness means

$$\Psi * \mathcal{R} = \Psi \quad \text{if} \quad \Psi \models \mathcal{R}$$

(cf. the stability postulate (CR2) in Section 4.1). Strong cumulativity is equivalent to

$$\Psi * (\mathcal{R} \cup \mathcal{S}) = (\Psi * \mathcal{R}) * (\mathcal{R} \cup \mathcal{S}), \quad (7.4)$$

that is, to what was called *logical coherence* in the framework of probabilistic inference (see equation (5.18) in Section 5.3). Equation (6.2) and Proposition 6.2.1 in Section 6.2 show that ME-inference defines a founded and strongly cumulative universal operation $\mathbf{C}^{ME} : P \mapsto C_P^{ME}$. Both properties, foundedness and strong cumulativity, may serve to control iterated revisions.

It is worth noticing that strong cumulativity, introduced here as a generalization of cumulativity for binary probabilistic revision operators and satisfied by ME-inference, may be considered as a set-theoretical version of postulate (C1) of the DP-postulates (cf. p. 22) for iterated belief change (cf. [DP97a]).

7.2 Simultaneous and Successive Revision

Investigating revision in the generalized framework of epistemic states and (quantified) conditionals allows a deeper insight into the mechanisms underlying the revision process. As a crucial difference to propositional belief revision, it is possible to distinguish between revising *simultaneously* and *successively*: In general, we have

$$\Psi * (\mathcal{R} \cup \mathcal{S}) \neq (\Psi * \mathcal{R}) * \mathcal{S} \quad (7.5)$$

(this is well-known for ME-inference); instead, we may only postulate strong cumulativity or logical coherence, respectively,

$$\Psi * (\mathcal{R} \cup \mathcal{S}) = (\Psi * \mathcal{R}) * (\mathcal{R} \cup \mathcal{S}),$$

which is essentially weaker. The failure revealed in (7.5) is responsible e.g. for the unpleasant complexity of ME-reasoning (cf. [Par94]). No cutting down to local propagation rules is possible here in general, but, on the other hand, we observe a greater variety of revision types. In fact, (7.5) allows us to incorporate knowledge on different levels:

Suppose the (already revised) epistemic state $\Psi * \mathcal{R}$ reflects our knowledge, and we learn the conditionals in \mathcal{S} to hold. Three (generally) different ways to revise $\Psi * \mathcal{R}$ by \mathcal{S} are imaginable:

- If we decide that \mathcal{R} and \mathcal{S} represent knowledge on the same level, then we should accept $\Psi * (\mathcal{R} \cup \mathcal{S})$ as revised epistemic state.
- Maybe we regard \mathcal{S} as successive to \mathcal{R} ; then $(\Psi * \mathcal{R}) * \mathcal{S}$ is supposed to reflect the new state of belief.

- A third type of revision arises if one considers \mathcal{S} as belonging to Ψ , perhaps representing additional background or generic knowledge. Then a suitable revision can be performed by calculating $(\Psi * \mathcal{S}) * \mathcal{R}$.

The first of these revision types realizes genuine revision: $\Psi * \mathcal{R}$ is revised by learning additional evidential knowledge. The second type is more in the sense of *updating*: We do not expect explicitly the knowledge in \mathcal{R} to hold any longer, rather we concede that some change in the world may have occurred so that \mathcal{S} now overrides \mathcal{R} . The third of the revision types above deals with a possible change in generic knowledge and raises a new perspective in the framework of belief change (see Example 7.5.1).

Thus in a generalized framework, different types of belief change may be realized by making use of one and the same (binary) revision operator, or universal inference operation, respectively, in different ways, allowing a convenient homogeneity of methods. For doing so, it is necessary, however, to represent background (generic, prior) knowledge as separated from evidential knowledge pertaining to the given situation (what is often regarded as an essential prerequisite for efficient plausible reasoning, cf. [DP96, DP97b]). In the following section, we will throw some formal light on belief revision and updating when it is possible to distinguish between knowledge on different levels.

7.3 Separating Background from Evidential Knowledge

Epistemic states provide a convenient, stable and rich representation of knowledge and may serve as an excellent starting point to perform a belief change operation. Yet they have no history, all (uncertain or conditional) knowledge is considered on the same level, no distinction is made between explicit and implicit knowledge, or between generic and evidential knowledge. Generic knowledge may be regarded as constraints imposed on epistemic states, so a proper handling is possible by considering *sets* of epistemic states (cf. e.g. [Voo96a]). If, however, one is not willing to give up the convenience of having a single epistemic state as “best” knowledge base, a way to overcome the restrictions described above may be offered by taking (properly defined) *belief bases* as primitive representations of epistemic knowledge, from which epistemic states may be calculated.

Definition 7.3.1. A belief base is a pair (Ψ, \mathcal{R}) , where Ψ is an epistemic state (background knowledge), and $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ is a set of (quantified) conditionals representing contextual (or evidential) knowledge.

Usually, evidential knowledge is restricted to certain facts, representing knowledge about a present *case*. This is generalized here by considering the evidence \mathcal{R} as reflecting knowledge about the *context* under consideration, thus being again of an epistemic nature. Certain knowledge is dealt with as a borderline case, but is not of major interest. So fluctuation of knowledge is modelled more naturally: Prior knowledge serves as a base for obtaining an adequate probabilistic description of the present context which may be used again as background knowledge for further change operations.

The transition from belief bases to epistemic states can be achieved by an adequate universal inference operation \mathbf{C} , or a binary belief change operator $*$, respectively:

$$*(\Psi, \mathcal{R}) := \Psi * \mathcal{R} \quad \text{with} \quad C_{\Psi}(\mathcal{R}) = Th^*(\Psi * \mathcal{R}) \quad (7.6)$$

For this to be well-defined, we have to ensure that both $C_{\Psi}(\mathcal{R}) \neq \emptyset$ and $C_{\Psi}(\mathcal{R}) \neq (\mathcal{L} \mid \mathcal{L})$. In this book, we will only deal with consistent belief change, assuming \mathbf{C} to preserve consistency and \mathcal{R} to be consistent with the prior knowledge Ψ , if necessary. Though struggling with inconsistencies is certainly a challenging subject, the concentration on handling consistent beliefs in the present framework will help to get a clear first view on the topic.

In the following, we will develop postulates for revising belief bases (Ψ, \mathcal{R}) by new conditional information $\mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})^*$, yielding a new belief base $(\Psi, \mathcal{R}) \circ \mathcal{S}$, in the sense of the AGM-postulates.

Due to distinguishing background knowledge from context information, we are able to compare the knowledge stored in different belief bases:

Definition 7.3.2. *A pre-ordering \sqsubseteq on belief bases is defined by*

$$(\Psi_1, \mathcal{R}_1) \sqsubseteq (\Psi_2, \mathcal{R}_2) \quad \text{iff} \quad \Psi_1 = \Psi_2 \quad \text{and} \quad \mathcal{R}_2 \models^* \mathcal{R}_1$$

(Ψ_1, \mathcal{R}_1) and (Ψ_2, \mathcal{R}_2) are \sqsubseteq -equivalent,

$$(\Psi_1, \mathcal{R}_1) \equiv_{\sqsubseteq} (\Psi_2, \mathcal{R}_2),$$

iff $(\Psi_1, \mathcal{R}_1) \sqsubseteq (\Psi_2, \mathcal{R}_2)$ and $(\Psi_2, \mathcal{R}_2) \sqsubseteq (\Psi_1, \mathcal{R}_1)$.

Therefore $(\Psi_1, \mathcal{R}_1) \equiv_{\sqsubseteq} (\Psi_2, \mathcal{R}_2)$ iff $\Psi_1 = \Psi_2$ and \mathcal{R}_1 and \mathcal{R}_2 are semantically equivalent, i.e. iff both belief bases reflect the same epistemic (background and contextual) knowledge. If the universal inference operation \mathbf{C} , resp. $*$, satisfies left logical equivalence (cf. p. 13) with respect to Cn^* , then $(\Psi_1, \mathcal{R}_1) \equiv_{\sqsubseteq} (\Psi_2, \mathcal{R}_2)$ implies $\Psi_1 * \mathcal{R}_1 = \Psi_2 * \mathcal{R}_2$.

The following postulates do not make use of the universal inference operation but are to characterize pure belief base revision by the revision operator \circ :

Postulates for conditional base revision:

Let Ψ be an epistemic state, and let $\mathcal{R}, \mathcal{R}_1, \mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ be sets of conditionals.

(CBR1) $(\Psi, \mathcal{R}) \circ \mathcal{S}$ is a belief base.

(CBR2) If $(\Psi, \mathcal{R}) \circ \mathcal{S} = (\Psi, \mathcal{R}_1)$ then $\mathcal{R}_1 \models^* \mathcal{S}$.

(CBR3) $(\Psi, \mathcal{R}) \sqsubseteq (\Psi, \mathcal{R}) \circ \mathcal{S}$.

(CBR4) $(\Psi, \mathcal{R}) \circ \mathcal{S}$ is a minimal belief base (with respect to \sqsubseteq) among all belief bases satisfying (PR1)-(PR3).

(CBR1) is the most fundamental axiom and coincides with the demand for *categorical matching* (cf. [GR94]). (CBR2) is generally called *success*: the new context information is now represented (up to epistemic equivalence). (CBR3) states that revision should preserve prior knowledge. Thus it is crucial for revision in contrast to update. Finally, (CBR4) is in the sense of *informational economy* (cf. [Gär88]): No unnecessary changes should occur. Admittedly, our postulates are much simpler than those proposed by Hansson (see, for instance, [Han89, Han91], and [GR94, p. 61]). They are, however, not based upon classical logic. So they are more adequate in the framework of general epistemic states.

The following characterization may be proved easily:

Theorem 7.3.1. *The revision operator \circ satisfies the axioms (CBR1) – (CBR4) iff*

$$(\Psi, \mathcal{R}) \circ \mathcal{S} \equiv_{\sqsubseteq} (\Psi, \mathcal{R} \cup \mathcal{S}). \quad (7.7)$$

So from (CBR1)-(CBR4), other properties of the revision operator also follow in a straightforward manner which are usually found among characterizing postulates:

Proposition 7.3.1. *Suppose the revision operator \circ satisfies (7.7). Then it fulfills the following properties:*

- (i) If $\mathcal{R} \models^* \mathcal{S}$, then $(\Psi, \mathcal{R}) \circ \mathcal{S} \equiv_{\sqsubseteq} (\Psi, \mathcal{R})$;
- (ii) If $(\Psi_1, \mathcal{R}_1) \sqsubseteq (\Psi_2, \mathcal{R}_2)$ then $(\Psi_1, \mathcal{R}_1) \circ \mathcal{S} \sqsubseteq (\Psi_2, \mathcal{R}_2) \circ \mathcal{S}$;
- (iii) $((\Psi, \mathcal{R}) \circ \mathcal{S}_1) \circ \mathcal{S}_2 \equiv_{\sqsubseteq} (\Psi, \mathcal{R}) \circ (\mathcal{S}_1 \cup \mathcal{S}_2)$,

where $(\Psi, \mathcal{R}), (\Psi_1, \mathcal{R}_1), (\Psi_2, \mathcal{R}_2)$ are belief bases and $\mathcal{S}, \mathcal{S}_1, \mathcal{S}_2 \subseteq (\mathcal{L} \mid \mathcal{L})^*$.

(i) shows a minimality of change, while (ii) is stated in [Gär88] as a monotonicity postulate. (iii) deals with the handling of non-conflicting iterated revisions.

Here we investigate revision merely under the assumption that the new information is compatible with what is already known. Belief revision based on classical logics is nothing but expansion in this case, and Theorem 7.3.1 indeed shows that revision of belief bases should reasonably mean expanding contextual knowledge. Note that revising a belief base (Ψ, \mathcal{R}) by $\mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ also induces a change of the corresponding belief state $\Psi^* = \Psi * \mathcal{R}$ to $(\Psi^*)' = *((\Psi, \mathcal{R}) \circ \mathcal{S})$. According to Theorem 7.3.1, if the (underlying) universal inference operation \mathbf{C} satisfies left logical equivalence, then the only reasonable revision operation (as specified by (CBR1)-(CBR4)) is given on the belief state level by

$$*((\Psi, \mathcal{R}) \circ \mathcal{S}) = \Psi * (\mathcal{R} \cup \mathcal{S}) \quad (7.8)$$

and therefore,

$$Th^*((\Psi, \mathcal{R}) \circ \mathcal{S}) = C_\Psi(\mathcal{R} \cup \mathcal{S})$$

This parallels the result for the classical belief revision theory, with the inference operation C_Ψ replacing the classical consequence operation (cf. [Gär88]; see also Theorem 2.2.1, p. 14). Nevertheless, we prefer using the more general term “revision” to “expansion” here. For if we consider the epistemic states generated by the two belief bases $\Psi * \mathcal{R}$ and $((\Psi, \mathcal{R}) \circ \mathcal{S}) = \Psi * (\mathcal{R} \cup \mathcal{S})$, we see that the epistemic status that $\Psi * \mathcal{R}$ assigns to conditionals occurring in \mathcal{S} will normally differ from those in Ψ as well as from those in $\Psi * (\mathcal{R} \cup \mathcal{S})$ with expanded contextual knowledge. So the belief in the conditionals in \mathcal{S} is actually *revised*.

7.4 Revising Epistemic States by Sets of Conditionals

Introducing belief bases in Section 7.3 opened up the possibility to perform genuine revisions in a clear way, namely by extending evidential or contextual knowledge learned about a *static* world. In general, however, a revision of an epistemic state by (sets of) conditionals may also be triggered by information referring to changes in the world, thus demanding actually for *updating* the epistemic state. Typical situations for updating occur when knowledge about a prior world is to be adapted to more recent information (e.g. a demographic model gained from statistical data of past periods should be brushed up by new data, see, for instance, the florida murderers-examples 3.5.1 and 4.5.1, or the Example 7.5.1 below).

In the sequel, we will list reasonable postulates for a (general) revision of epistemic states by sets of conditionals, matching both the frameworks of (genuine) revision and of updating. The postulates partly generalize those for revising an epistemic state by a single conditional listed in Section 4.1,

page 55. Furthermore, we will compare these postulates to properties of the underlying universal inference operation, thus continuing to study the close connection between belief change on one hand and nonmonotonic inference operations, on the other hand, which has been known already for a couple of years (see, for instance, [MG91, GR94]) within the framework of classical logics.

Postulates for revising epistemic states by sets of conditionals:

(CSR1) $\Psi * \mathcal{R} \models \mathcal{R}$.

(CSR2) If $\Psi \models \mathcal{R}$ then $\Psi * \mathcal{R} = \Psi$.

(CSR3) If \mathcal{R}_1 and \mathcal{R}_2 are equivalent, then $\Psi * \mathcal{R}_1 = \Psi * \mathcal{R}_2$.

(CSR4) If $\Psi * \mathcal{R}_1 \models \mathcal{R}_2$ and $\Psi * \mathcal{R}_2 \models \mathcal{R}_1$ then $\Psi * \mathcal{R}_1 = \Psi * \mathcal{R}_2$.

(CSR5) $\Psi * (\mathcal{R}_1 \cup \mathcal{R}_2) = (\Psi * \mathcal{R}_1) * (\mathcal{R}_1 \cup \mathcal{R}_2)$.

Postulates (CSR1), (CSR2) and (CSR3) constitute basic properties of epistemic belief change. (CSR4) states that two revising procedures with respect to sets \mathcal{R}_1 and \mathcal{R}_2 should result in the same epistemic knowledge base if each revision represents the new information of the other. This property is called *reciprocity* in the framework of nonmonotonic logics (cf. [Mak94]) and appears as axiom (U6) in the work of Katsuno and Mendelzon [KM91b]. (CSR5) is the postulate for logical coherence and deals with *iterative revision*. It demands that at least, updating any intermediate epistemic state $\Psi * \mathcal{R}_1$ by the full information $\mathcal{R}_1 \cup \mathcal{R}_2$ should result in the same epistemic state as revising Ψ by $\mathcal{R}_1 \cup \mathcal{R}_2$ in one step. The rationale behind this axiom is that if the information about the new world drops in in parts, updating any provisional state of belief by the full information should result unambiguously in a final belief state.

Note that in general, the revisions $(\Psi * \mathcal{R}_1) * \mathcal{R}_2$ and $(\Psi * \mathcal{R}_1) * (\mathcal{R}_1 \cup \mathcal{R}_2)$ will differ because the first is not supposed to maintain prior contextual information, \mathcal{R}_1 . As was already mentioned, (CSR5) is a set-theoretical version of axiom (C1) in [DP97a] (see also page 22). (CSR5) has proved to be a crucial property for the characterization of ME-inference (cf. Chapter 5) but actually goes back to [SJ81].

Postulates for reasonable revisions or updatings, respectively, based on inference processes are also proposed in [PV92].

For representing revision operations satisfying the postulates stated above, we will make use of the relationship between binary revision operators and universal inference operations:

Proposition 7.4.1. *Suppose revision is being realized via a universal inference operation as in (7.3).*

- (i) $*$ satisfies (CSR1) iff \mathbf{C} is reflexive.
- (ii) $*$ satisfies (CSR2) iff \mathbf{C} is founded.
- (iii) $*$ satisfies (CSR3) iff \mathbf{C} satisfies left logical equivalence.
- (iv) Assuming reflexivity resp. the validity of (CSR1), $*$ satisfies (CSR4) iff \mathbf{C} is cumulative.
- (v) Assuming foundedness resp. the validity of (CSR2), $*$ satisfies (CSR5) iff \mathbf{C} is strongly cumulative.

The proofs are immediate. From this proposition, a representation result follows in a straightforward manner:

Theorem 7.4.1. *If $*$ is defined by (7.3), it satisfies all of the postulates (CSR1)-(CSR5) iff the universal inference operation \mathbf{C} is reflexive, founded, strongly cumulative and satisfies left logical equivalence.*

7.5 Revision versus Updating

In this section, we will try to get a clearer view on formal parallels and differences, respectively, between (genuine) revision and updating. For an adequate comparison, we have to observe the changes of belief *states* that are induced by revision of belief *bases*. Observing (7.6), (CBR2) and (CBR3) translate into

$$(\mathbf{CBR2}') \quad *((\Psi, \mathcal{R}) \circ \mathcal{S}) \models \mathcal{S}.$$

$$(\mathbf{CBR3}') \quad *((\Psi, \mathcal{R}) \circ \mathcal{S}) \models \mathcal{R}.$$

While (CBR2') parallels (CSR1), (CBR3') establishes a crucial difference between revision and updating: revision preserves prior knowledge while updating does not, neither in a classical nor in a generalized framework.

The intended effects of revision and updating on a belief state $\Psi * \mathcal{R}$ that is generated by a belief base (Ψ, \mathcal{R}) are made obvious by – informally! – writing

$$(\Psi * \mathcal{R}) \circ \mathcal{S} = \Psi * (\mathcal{R} \cup \mathcal{S}) \neq (\Psi * \mathcal{R}) * \mathcal{S} \quad (7.9)$$

(cf. (7.8)). This reveals clearly the difference, but also the relationship between revision and updating: Revising $\Psi * \mathcal{R}$ by \mathcal{S} results in the same state of belief as updating Ψ by (the full contextual information) $\mathcal{R} \cup \mathcal{S}$. Note also the difference in case that $\Psi * \mathcal{R} \models \mathcal{S}$: If $*$ satisfies (CSR2) then $(\Psi * \mathcal{R}) * \mathcal{S} = \Psi * \mathcal{R}$, but, in general, $\Psi * (\mathcal{R} \cup \mathcal{S})$ will differ from $\Psi * \mathcal{R}$. This does not violate the cumulativity of $*$, or of C_Ψ , respectively, because \mathcal{S} is not supposed to include \mathcal{R} (if it does, then (strong) cumulativity yields equality). Rather it reveals

clearly the distinction between having \mathcal{S} to be incorporated as explicit constraints and having \mathcal{S} only satisfied in a revised epistemic state (cf. [GP96, page 88]).

The representation of an epistemic state by a belief base, however, is not unique, different belief bases may generate the same belief state (the same holds for classical belief bases, cf. [Han89], [GR94, p. 48]). So we could not define genuine epistemic revision on belief states, but had to consider belief bases in order to separate background and context knowledge unambiguously. It is interesting to observe, however, that strong cumulativity, together with foundedness, ensures at least a convenient independence of revision from background knowledge: If two belief bases (Ψ_1, \mathcal{R}) , (Ψ_2, \mathcal{R}) with different prior knowledge but the same contextual knowledge give rise to the same belief state

$$\Psi_1 * \mathcal{R} = \Psi_2 * \mathcal{R},$$

then – assuming strong cumulativity and foundedness –

$$\begin{aligned} \Psi_1 * (\mathcal{R} \cup \mathcal{S}) &= (\Psi_1 * \mathcal{R}) * (\mathcal{R} \cup \mathcal{S}) \\ &= (\Psi_2 * \mathcal{R}) * (\mathcal{R} \cup \mathcal{S}) \\ &= \Psi_2 * (\mathcal{R} \cup \mathcal{S}). \end{aligned}$$

So strong cumulativity and foundedness guarantee a particular well-behavedness with respect to inference, updating and revision.

In the following example, we will illustrate revision and updating in a probabilistic environment, using ME-inference as the proper universal inference operation.

Example 7.5.1. A psychologist has been working with addicted people for a couple of years. His experiences concerning the propositions

$\mathcal{V} : a : \text{addicted to } \underline{a}lcohol$
 $d : \text{addicted to } \underline{d}rugs$
 $y : \text{being } \underline{y}oung$

may be summarized by the following distribution P that expresses his belief state probabilistically (where negation is indicated by barring the corresponding letter):

ω	$P(\omega)$	ω	$P(\omega)$	ω	$P(\omega)$	ω	$P(\omega)$
ady	0.050	$\bar{a}dy$	0.333	$ad\bar{y}$	0.053	$\bar{a}d\bar{y}$	0.053
$a\bar{d}y$	0.093	$\bar{a}\bar{d}y$	0.102	$a\bar{d}\bar{y}$	0.225	$\bar{a}\bar{d}\bar{y}$	0.091

The following probabilistic conditionals may be entailed from P :

$$\begin{array}{ll}
 (d|a)[0.242] & (\text{i.e. } P(d|a) = 0.242) \\
 (d|\bar{a})[0.666] & (\text{i.e. } P(d|\bar{a}) = 0.666) \\
 (a|y)[0.246] & (a|\bar{y})[0.660] \\
 (d|y)[0.662] & (d|\bar{y})[0.251]
 \end{array}$$

Now the psychologist is going to change his job: He will be working in a clinic for people addicted only to alcohol and/or drugs. He is told that the percentage of persons addicted to alcohol, but also addicted to drugs, is higher than usual and may be estimated by 40 %.

So the information the psychologist has about the “new world” is represented by

$$\mathcal{R} = \{a \vee d[1], (d|a)[0.4]\}.$$

The distribution P from above is now supposed to represent background or prior knowledge, respectively. So the psychologist revises or updates, respectively, P by \mathcal{R} using ME-inference and obtains $P^* = P *_{ME} \mathcal{R}$ as new belief state:

ω	$P^*(\omega)$	ω	$P^*(\omega)$	ω	$P^*(\omega)$	ω	$P^*(\omega)$
ady	0.099	$\bar{a}dy$	0.425	$ad\bar{y}$	0.105	$\bar{a}d\bar{y}$	0.066
$a\bar{d}y$	0.089	$\bar{a}\bar{d}y$	0.0	$a\bar{d}\bar{y}$	0.216	$\bar{a}\bar{d}\bar{y}$	0.0

After having spent a couple of days in the new clinic, the psychologist realized that this clinic was for young people only. So he had to revise genuinely his knowledge about his new sphere of activity and arrived at the revised belief state $*_{ME}((P, \mathcal{R}) \circ y[1]) = P *_{ME} (\mathcal{R} \cup y[1]) =: P_1^*$ shown in the following table:

ω	$P_1^*(\omega)$	ω	$P_1^*(\omega)$	ω	$P_1^*(\omega)$	ω	$P_1^*(\omega)$
ady	0.120	$\bar{a}dy$	0.700	$ad\bar{y}$	0.0	$\bar{a}d\bar{y}$	0.0
$a\bar{d}y$	0.180	$\bar{a}\bar{d}y$	0.0	$a\bar{d}\bar{y}$	0.0	$\bar{a}\bar{d}\bar{y}$	0.0

(7.10)

7.6 Focusing in a Probabilistic Framework

Focusing means applying generic knowledge to a reference class appropriate to describe the context of interest (cf. [DP97b, DP96]). In a probabilistic setting, focusing is best done by conditioning which, however, is used for revision, too. So revision and focusing are supposed to coincide in the framework of Bayesian probabilities though they differ conceptually: Revision

is not only *applying knowledge*, but means incorporating a new constraint so as to *refine knowledge*.

Dubois and Prade argued that the assumption of having a uniquely determined probability distribution to represent the available knowledge at best is responsible for that flaw, and they recommend to use upper and lower probabilities to permit a proper distinction (cf. [DP96, DP97b]).

Making use of ME-inference $*_{ME}$, however, it is indeed possible to realize this conceptual difference appropriately without giving up the convenience of having a single distribution for inferences. To make this clear, we have to consider belief changes induced by some certain information $A[1]$, that is, we learn proposition A with certainty. The following proposition reveals the difference between revision by $A[1]$, as realized according to (7.8), and focusing to A by conditioning (cf. also Proposition 6.3.2 in Section 6.3).

Proposition 7.6.1. *Let P be a distribution, $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$ a P -consistent set of probabilistic conditionals, and suppose $A[1]$ to be a certain probabilistic fact.*

- (i) $P *_{ME} \{A[1]\} = P(\cdot|A)$;
in particular, $(P *_{ME} \mathcal{R}) *_{ME} A[1] = (P *_{ME} \mathcal{R})(\cdot|A)$.
- (ii) $*_{ME}((P, \mathcal{R}) \circ \{A[1]\}) = P *_{ME} (\mathcal{R} \cup \{A[1]\}) = P(\cdot|A) *_{ME} \mathcal{R}$.

Both parts of this proposition may be proved by using the representation (5.5) of the ME-distribution.

Example 7.6.1 (continued). The distribution obtained in (7.10) by *revision* is different from that one the psychologist would have obtained by *focusing* his knowledge represented by $P^* = P *_{ME} \mathcal{R}$ on a young patient, which is given by $P^* *_{ME} \{y[1]\} = P^*(\cdot|y) =: P_y^*$:

ω	$P_y^*(\omega)$	ω	$P_y^*(\omega)$	ω	$P_y^*(\omega)$	ω	$P_y^*(\omega)$
ady	0.162	$\bar{a}dy$	0.693	$ad\bar{y}$	0.0	$\bar{a}d\bar{y}$	0.0
$a\bar{d}y$	0.145	$\bar{a}\bar{d}y$	0.0	$a\bar{d}\bar{y}$	0.0	$\bar{a}\bar{d}\bar{y}$	0.0

Proposition 7.6.1 and Examples 7.5.1 and 7.6.1 show that, in a (generalized) probabilistic framework, a proper distinction between focusing and revision is possible. This difference is akin to the one between “conditioning” and “constraining” elaborated by Voorbraak [Voo96a] for classes of probability functions (for a criticism of conditioning *sets* of probability measures, cf. [GH98]). Paris and Vencovská [PV92] also consider focusing by using uncertain information in the context of various inference processes.

8. Knowledge Discovery by Following Conditional Structures

In many cases, knowledge bases for expert systems consist of rules, i.e., of conditional statements. In the previous chapters, we investigated in detail the formal properties of conditionals, how to represent them appropriately and how to handle them under change of beliefs. Solving these problems is a necessary prerequisite to arrive at a satisfactory representation and processing of knowledge. When designing an expert system, however, at first one has to face another crucial problem: Where do all the rules come from? How to find a set of rules representing relevant knowledge in an exhaustive way? Besides human expertise, also experimental data may be available. Incorporating the detailed experiences of an expert into the knowledge base usually is an indispensable task in knowledge acquisition. Extracting and providing information from databases, however, may essentially help to support, automate and improve this process.

Data mining and *knowledge discovery*, respectively, mean finding new and relevant information in databases. Usually, knowledge discovering is understood as the more comprehensive task, including preparing and cleaning the available data and interpreting the results revealed by the actual data mining process, aiming at discovering interesting patterns in data (cf. [FPSS96, FPSSU96, FU⁺96]).

In this chapter, we will focus on this central part of knowledge discovery within a probabilistic framework, where we assume experimental data to be represented by a probability distribution. This means that we will deal with relatively “small” data mining problems with respect to the number of variables or propositions involved. By using *clustering techniques* (see, for instance, [AGGR98]) and considering *LEG-networks* as an appropriate tool to split up large probability distributions in a system of local distributions (see Chapter 9), however, the problem of discovering relevant relationships amongst variables can be reduced to mining manageable distributions.

Relationships amongst variables and sets of variables may be expressed by *association rules* (cf. [AIS93, MS98, AMS⁺96, SA95, Bol96], and see below) which are a special kind of probabilistic conditionals. Relevance of such rules is usually measured by considering their *confidence*, which is nothing but a

conditional probability, and their *support*, which is the number of cases in the database that a rule is based upon (cf. [AIS93]). These are certainly plausible indicators for a rule to interest the user. If, for instance, the manager of a supermarket wants to improve the layout of his store, he will be interested in knowing how many customers buying product A also buy product B, and which percentage of total sales those transactions constitute.

When designing the knowledge base of an expert system, however, relevance of rules depends on the representation and inference methods used. Within a probabilistic framework, knowledge of conditional independences is of particular importance. Moreover, when using ME-methods, we would best find a set of conditionals that represents the distribution under consideration by means of ME-propagation. The corresponding rules will be considered not only technically relevant, but also relevant in a fundamental, information-theoretical sense. By observing that ME-inference obeys the principle of conditional preservation (cf. Section 5.1 and Section 4.5), we have to search for a set of conditionals with respect to which a given distribution is indifferent (cf. Definition 4.5.1). An approach to solve this problem is developed in Section 8.2 and constitutes the main contribution of this chapter. We will illustrate this method in several examples. In particular, we will show how it may help to find a suitable ME-optimal set of rules.

We will start with recalling some results from general probabilistic knowledge discovery.

8.1 Probabilistic Knowledge Discovery

Mining statistical databases roughly serves three purposes: Firstly, one is interested in finding relevant *association rules*, i.e. expressions of the form $A \rightarrow B$ where A, B are disjoint subsets of an *item set* \mathcal{I} (cf. [AIS93]). The database yields a relative frequency distribution Pr over \mathcal{I} , and the *support* and the *confidence*, respectively, of such a rule $A \rightarrow B$ is simply defined as $Pr(A \cup B)$ and $Pr(B|A)$, respectively. Effective algorithms are available to find significant association rules even in large databases (see [AIS93, MS98, AMS⁺96, SA95, Bol96]).

As a second task, statistical databases and probability distributions, respectively, are investigated in search of *causal structures* which can be represented as *Markov graphs*, or as *directed acyclic graphs* (see [SGS93]). The notion of conditional independence is fundamental to those techniques, and an important application is the discovery of Bayesian networks from data (see, for instance, [SGS93, Jen96, Hec96]).

Thirdly, if the number of variables is not too large, one may consider the resulting distribution as representing important inherent relationships between propositional variables in a more logical sense and search for interesting conditional rules revealing these relationships. This approach may be considered as dual to the second one: Here *dependencies*, not *independencies*, are to be discovered. This is often done for building up the knowledge base of an expert system. Here a special type of conditional, the so-called *single-elementary conditional*, is of particular interest, since it is supposed to pinpoint relevant information:

Definition 8.1.1. *A conditional $(B|A)$ is called a single-elementary conditional if the antecedent A is an elementary conjunction, i.e. a conjunction of literals, and the conclusion B is a single literal, and if neither $AB \equiv \perp$ nor $A\bar{B} \equiv \perp$. The basic conditional $\psi_{\omega, \omega'} = (\text{form}(\omega) | \text{form}(\omega, \omega'))$ (cf. Definition 3.4.2) is a basic single-elementary conditional if ω and ω' are neighboring worlds, i.e. interpretations differing with respect to exactly one atom.*

Traditionally, single-elementary conditionals are appreciated for representing knowledge very clearly and intelligibly. Thus they occupy an outstanding position for knowledge representation and reasoning. Association rules are similar to single-elementary conditionals. There are, however, structural differences: Antecedent and conclusion of an association rule may be considered as elementary conjunctions with only positive literals. The conclusion of a single-elementary conditional consists of only one literal, i.e., it contains exactly one item. Basic single-elementary conditionals are single-elementary conditionals with antecedents of maximal length.

Usually, single-elementary conditionals are considered important if the corresponding probability is significantly high, that is, near to 1 within a certain distance $\epsilon \geq 0$. But high probability alone does not suffice to make a conditional really *relevant*: If $P(B|A) > 1 - \epsilon$ then there will normally be a lot of other variables V such that $P(B|A\dot{V}) > 1 - \epsilon$, too. So the problem of discovering relevant single-elementary conditionals may be conceived as finding single-elementary conditionals where the antecedent is as short as possible (with respect to the number of occurring literals). This problem is dealt with by the author and others in [KIR96, Ger97], and [Sch98] ; moreover, [Sch98] also searches for exceptions to such shortest single-elementary conditionals. In both cases, parts of the computed conditionals were used for probabilistic knowledge representation via ME-inference (cf. [RKI97b]). These sets of conditionals, however, were not optimally suitable because no specific feature of ME-methods was taken into consideration in the discovery algorithms. In particular, ME-propagation yields a distribution which is indifferent with respect to the set of learned conditionals (see Section 5.1). So, as a necessary condition for an appropriate set \mathcal{R} of conditionals ME-representing a

given distribution P , we may claim that P be indifferent with respect to \mathcal{R} (see Section 3.6.1). An approach to solve this problem is developed in the next section within a more general setting, considering conditional valuation functions instead of probability distributions.

8.2 Discovering Conditional Structures

Given some conditional valuation function $V : \mathcal{L} \rightarrow (\mathcal{A}, \oplus, \odot, 0^{\mathcal{A}}, 1^{\mathcal{A}})$, two questions arise at once:

- What knowledge does V represent? What are the propositional and the conditional beliefs held in V ?
- Which subset of these conditionals (including facts) is distinguished in the sense that V is in accordance with the conditional structures it imposes on possible worlds?

The first question means answering queries $(B|A)[x], x = ?$, by calculating $V(B|A)$. The second question amounts to finding a set $\mathcal{R} \subseteq (\mathcal{L} | \mathcal{L})^{(*)}$ such that V is a c-representation with respect to \mathcal{R} . Ideally, we would have V to be a faithful c-representation, i.e. we are searching for a set $\mathcal{R} \subseteq (\mathcal{L} | \mathcal{L})^{(*)}$ such that $V \models \mathcal{R}$ and $\ker V = \ker \sigma_{\mathcal{R}}$, or $\ker_0 V = \ker_0 \sigma_{\mathcal{R}}$, respectively. Assuming faithfulness means presupposing that no equation $V(\hat{\omega}) = 1^{\mathcal{A}}$ is fulfilled accidentally, but that any of these equations is induced by \mathcal{R} (cf. also the *Faithfulness condition* in [SGS93, pp. 35f.]).

At the end of Section 3.5, we made some first considerations concerning this generally complicated and expensive task.

In this section, as the main result of this chapter, we will present an approach to computing sets of conditionals that underly the knowledge represented by some conditional valuation function $V : \mathcal{L} \rightarrow \mathcal{A}$. As a crucial prerequisite, we will assume that this knowledge is representable by a set of *single-elementary conditionals*. This restriction to single-elementary conditionals is important, but should not be considered as a heavy drawback bearing in mind the expressibility of these conditionals.

Suppose $\mathcal{R} \subseteq (\mathcal{L} | \mathcal{L})$ is an existing, but unknown set of *single-elementary conditionals*, such that $\ker \sigma_{\mathcal{R}} = \ker V$, and $\ker V$ is known. In the following, we will present a method for determining or approximating, respectively, sets $\mathcal{S} \subseteq (\mathcal{L} | \mathcal{L})$ such that $\ker V = \ker \sigma_{\mathcal{S}}$.

Each conditional in \mathcal{R} is presupposed to be single-elementary, so we set

$$\mathcal{R} = \{(b_1|A_1), \dots, (b_n|A_n)\} \quad (8.1)$$

where A_i are elementary conjunctions and b_i are literals, $1 \leq i \leq n$. Without loss of generality, only to simplify notation, we assume all literals b_i to be positive (cf. Lemma 3.5.2).

Let

$$\sigma_{\mathcal{R}} : \widehat{\Omega} \rightarrow \mathcal{F}_{\mathcal{R}} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \dots, \mathbf{a}_n^+, \mathbf{a}_n^- \rangle$$

denote a conditional structure homomorphism with respect to \mathcal{R} (cf. Equation (3.19) in Section 3.5, page 42).

For each atom $v \in \mathcal{L}$, choose an arbitrary, but fixed numbering of the remaining atoms $\{w \mid w \neq v\} = (w_0, w_1, \dots, w_{\#(atoms)-1})$. Then *basic single-elementary conditionals* are conditionals of the form

$$\psi_{v,l} = (v \mid \bigwedge_j w_j^{\epsilon_j}) \quad (8.2)$$

with $\epsilon_j \in \{0, 1\}$, $w_j^1 := w_j$, $w_j^0 := \overline{w_j}$, $0 \leq j \leq \#(atoms) - 1$ and $l = \sum_j \epsilon_j 2^j$. We will abbreviate the antecedent of $\psi_{v,l}$ by $C_{v,l}$:

$$C_{v,l} := \bigwedge_j w_j^{\epsilon_j}, \quad l = \sum_j \epsilon_j 2^j \quad (8.3)$$

(the numbering w_j depends on v). Let

$$\mathcal{B} = \{\psi_{v,l} \mid v \text{ atom in } \mathcal{L}, 0 \leq l \leq 2^{\#(atoms)-1} - 1\}$$

denote the set of all basic single-elementary conditionals in $(\mathcal{L} \mid \mathcal{L})$, and let

$$\mathcal{F}_{\mathcal{B}} = \langle \mathbf{b}_{v,l}^+, \mathbf{b}_{v,l}^- \mid v \text{ atom in } \mathcal{L}, 0 \leq l \leq 2^{\#(atoms)-1} - 1 \rangle$$

be the free abelian group corresponding to \mathcal{B} (cf. Section 3.5) with conditional structure homomorphism

$$\sigma_{\mathcal{B}} = \prod_{v,l} \sigma_{v,l} : \widehat{\Omega} \rightarrow \mathcal{F}_{\mathcal{B}}, \quad (8.4)$$

$$\sigma_{v,l}(\omega) = \begin{cases} \mathbf{b}_{v,l}^+ & \omega = C_{v,l}v \\ \mathbf{b}_{v,l}^- & \omega = C_{v,l}\bar{v} \\ 1 & \text{else} \end{cases} \quad (8.5)$$

Lemma 8.2.1. $\sigma_{\mathcal{B}}$ is injective, i.e. $\ker \sigma_{\mathcal{B}} = \{1\}$.

So $\sigma_{\mathcal{B}}$ provides the most finely grained conditional structure on $\widehat{\Omega}$: No different elements $\widehat{\omega}_1 \neq \widehat{\omega}_2$ are equivalent with respect to \mathcal{B} .

We define a homomorphism

$$g : \mathcal{F}_{\mathcal{B}} \rightarrow \mathcal{F}_{\mathcal{R}}$$

via

$$g(\mathbf{b}_{v,l}^+) = \prod_{\substack{1 \leq i \leq n \\ \psi_{v,l} \sqsubseteq (b_i | A_i)}} \mathbf{a}_i^+ \quad \text{and} \quad g(\mathbf{b}_{v,l}^-) = \prod_{\substack{1 \leq i \leq n \\ \psi_{v,l} \sqsubseteq (b_i | A_i)}} \mathbf{a}_i^- \quad (8.6)$$

where \sqsubseteq is defined according to Definition 3.4.1 in Section 3.4, p. 36. Note that we presupposed all conditionals in \mathcal{R} to have positive literals in the consequent. Otherwise, the definition of g has to be altered appropriately, but the results to be presented in the sequel will still hold. The prerequisite of dealing with single-elementary conditionals, however, is essential for the following. The next lemma provides an easy, but far-reaching characterization for the relation \sqsubseteq to hold between single-elementary conditionals:

Lemma 8.2.2. *Let $(b|A)$ and $(d|C)$ be two single-elementary conditionals. Then*

$$(d|C) \sqsubseteq (b|A) \quad \text{iff} \quad C \leq A \text{ and } b = d \quad (8.7)$$

Remark 8.2.1. The preceding lemma may be slightly generalized to hold for conditionals $(b|A)$ and $(d|C)$ where A and C are disjunctions of elementary conjunctions not containing b resp. d .

Using Lemma 8.2.2, we have

$$g(\mathbf{b}_{v,l}^+) = \prod_{\substack{1 \leq i \leq n \\ b_i = v, C_{v,l} \leq A_i}} \mathbf{a}_i^+ \quad \text{and} \quad g(\mathbf{b}_{v,l}^-) = \prod_{\substack{1 \leq i \leq n \\ b_i = v, C_{v,l} \leq A_i}} \mathbf{a}_i^- \quad (8.8)$$

It is important to note that for different atoms v and v' , only different \mathbf{a}_i^+ occur in $g(\mathbf{b}_{v,l}^+)$ and $g(\mathbf{b}_{v',l'}^+)$, respectively, by Lemma 8.2.2 (analogically for \mathbf{a}_i^- and $g(\mathbf{b}_{v,l}^-)$ and $g(\mathbf{b}_{v',l'}^-)$, respectively). Moreover, each \mathbf{a}_i^+ and \mathbf{a}_i^- , respectively, occurs at most once in each $g(\mathbf{b}_{v,l}^+)$ and $g(\mathbf{b}_{v,l}^-)$, respectively. This will be used several times in the sequel.

g establishes a connection between the conditional structures with respect to \mathcal{B} and to – the still unknown, but existing – \mathcal{R} :

Theorem 8.2.1. *Let $g : \mathcal{F}_{\mathcal{B}} \rightarrow \mathcal{F}_{\mathcal{R}}$ be defined as in 8.6. Then $\sigma_{\mathcal{R}} = g \circ \sigma_{\mathcal{B}}$.*

The property of all conditionals in \mathcal{R} to be single-elementary is crucial for the proof of this theorem. Only in this case it is guaranteed that each \mathbf{a}_i^+ or \mathbf{a}_i^- , respectively, occurring in $\sigma_{\mathcal{R}}(\omega)$, also occurs exactly once in $g \circ \sigma_{\mathcal{B}}(\omega)$.

Theorem 8.2.1 provides immediately a method for determining $\ker g$ by considering $\sigma_{\mathcal{B}}$ and $\ker \sigma_{\mathcal{R}}$ which is assumed to be known:

Corollary 8.2.1. $\hat{\omega} \in \ker \sigma_{\mathcal{R}}$ iff $\sigma_{\mathcal{B}}(\hat{\omega}) \in \ker g$.

Proposition 8.2.1. Let $\hat{\omega} = \omega_1^{r_1} \cdot \dots \cdot \omega_m^{r_m} \in \hat{\Omega}$.

Then $\sigma_{\mathcal{B}}(\omega_1^{r_1} \cdot \dots \cdot \omega_m^{r_m}) \in \ker g$ iff for all atoms v in \mathcal{L} ,

$$\prod_{\substack{1 \leq k \leq m \\ \omega_k = C_{v,l} v}} (\mathbf{b}_{v,l}^+)^{r_k}, \quad \prod_{\substack{1 \leq k \leq m \\ \omega_k = C_{v,l} \bar{v}}} (\mathbf{b}_{v,l}^-)^{r_k} \in \ker g. \quad (8.9)$$

So each (generating) element of $\ker \sigma_{\mathcal{R}}$ gives rise to an equation modulo $\ker g$ for the generators $\mathbf{b}_{v,l}^+, \mathbf{b}_{v,l}^-$ of $\mathcal{F}_{\mathcal{B}}$.

Lemma 8.2.3. Let v be an atom of the language \mathcal{L} .

$\prod_{1 \leq k \leq m} (\mathbf{b}_{v,l_k}^+)^{r_k} \in \ker g$ or $\prod_{1 \leq k \leq m} (\mathbf{b}_{v,l_k}^-)^{r_k} \in \ker g$ iff for all $(b_i | A_i) \in \mathcal{R}$, $1 \leq i \leq n$, such that $v = b_i$ it holds that $\sum_{k: C_{v,l_k} \leq A_i} r_k = 0$.

This lemma shows a complete symmetry between the generators $\mathbf{b}_{v,l}^+$ and $\mathbf{b}_{v,l}^-$ occurring in elements of $\ker g$ (which is also obvious by the definition of g). So the superscripts may be omitted if not explicitly needed. Formally, let $\mathbf{b}_{v,l}$ denote the quotient of $\mathbf{b}_{v,l}^+$ and $\mathbf{b}_{v,l}^-$:

$$\mathbf{b}_{v,l} = \frac{\mathbf{b}_{v,l}^+}{\mathbf{b}_{v,l}^-} \quad (8.10)$$

Then the following corollary holds:

Corollary 8.2.2. Let v be an atom of the language \mathcal{L} .

$$\begin{aligned} \prod_{1 \leq k \leq m} (\mathbf{b}_{v,l_k}^+)^{r_k} \in \ker g & \quad \text{iff} \quad \prod_{1 \leq k \leq m} (\mathbf{b}_{v,l_k}^-)^{r_k} \in \ker g \\ & \quad \text{iff} \quad \prod_{1 \leq k \leq m} (\mathbf{b}_{v,l_k})^{r_k} \in \ker g \end{aligned}$$

The idea of the procedure to be described in the sequel is to explore the relations mod $\ker g$ holding between the group elements $\mathbf{b}_{v,l} \in \mathcal{F}_{\mathcal{B}}$ with the aim to define a finite sequence of sets $\mathcal{S}^{(0)}, \mathcal{S}^{(1)}, \dots$ of conditionals such that

$$\ker \sigma_{\mathcal{S}^{(0)}} \subseteq \ker \sigma_{\mathcal{S}^{(1)}} \subseteq \dots \subseteq \ker \sigma_{\mathcal{R}} \quad (8.11)$$

Thus the sequence $\mathcal{S}^{(0)}, \mathcal{S}^{(1)}, \dots$ tries to approximate \mathcal{R} . We will first present the fundamental idea of the method and develop the necessary theoretical results in the rest of this section. In the next section, the procedure will be explained and illustrated in detail by examples.

Step 0:

We start with setting

$$\mathcal{S}^{(0)} = \mathcal{B} \quad (8.12)$$

Lemma 8.2.1 states $\ker \sigma_{\mathcal{S}^{(0)}} = 1$, so (8.11) trivially holds. Let \equiv_g denote the equivalence relation mod $\ker g$ on $\mathcal{F}_{\mathcal{B}}$, i.e. $\mathbf{b}_1 \equiv_g \mathbf{b}_2$ iff $g(\mathbf{b}_1) = g(\mathbf{b}_2)$ for any two elements $\mathbf{b}_1, \mathbf{b}_2 \in \mathcal{F}_{\mathcal{B}}$, where g is defined as in (8.6). For each (generating) element $\hat{\omega} = \omega_1^{r_1} \cdots \omega_m^{r_m}$ of $\ker \sigma_{\mathcal{R}}$, set up an equation modulo $\ker g$:

$$\sigma_{\mathcal{B}}(\hat{\omega}) \equiv_g 1,$$

this means, according to (3.21),

$$1 = \prod_{v,l} (\mathbf{b}_{v,l}^+)^{\sum_{k:\sigma_{v,l}(\omega_k)=\mathbf{b}_{v,l}^+} r_k} \prod_{v,l} (\mathbf{b}_{v,l}^-)^{\sum_{k:\sigma_{v,l}(\omega_k)=\mathbf{b}_{v,l}^-} r_k}, \quad (8.13)$$

and split up these equations according to Proposition 8.2.1 and Corollary 8.2.2.

Step 1:

First, eliminate from $\mathcal{B} = \mathcal{S}^{(0)}$ all $\psi_{v,l}$ for which there is an equation $\mathbf{b}_{v,l} \equiv_g 1$:

$$\mathcal{S}^{(1)'} = \mathcal{S}^{(0)} - \{\psi_{v,l} \in \mathcal{S}^{(0)} \mid \mathbf{b}_{v,l} \equiv_g 1 \text{ is known}\}$$

The equations modulo $\ker g$ further partition $\mathcal{S}^{(1)'}$ into equivalence classes $[\mathbf{b}_{v,l}]_g = \{\mathbf{b}_{v,l'} \in \mathcal{S}^{(1)'} \mid \mathbf{b}_{v,l} \equiv_g \mathbf{b}_{v,l'} \text{ is known}\}$, $\mathbf{b}_{v,l} \in \mathcal{S}^{(1)'}$ (only $\mathbf{b}_{v,l'}$ with the same v occur in $[\mathbf{b}_{v,l}]_g$, according to Lemma 8.2.2). For each such equivalence class $[\mathbf{b}_{v,l}]_g$, set

$$D_{v,j_v}^{(1)} = \bigvee_{\mathbf{b}_{v,l'} \in [\mathbf{b}_{v,l}]_g} C_{v,l'}$$

and

$$\varphi_{v,j_v}^{(1)} = (v \mid D_{v,j_v}^{(1)}) = \bigsqcup_{\mathbf{b}_{v,l'} \in [\mathbf{b}_{v,l}]_g} \psi_{v,l'},$$

where $j_v = 1, 2, \dots$ is a proper finite numbering. Now we set

$$\mathcal{S}^{(1)} = \{\varphi_{v,j_v}^{(1)} \mid v \text{ atom}, j_v = 1, 2, \dots\} \quad (8.14)$$

Define homomorphisms

$$h^{(1)} : \mathcal{F}_{\mathcal{S}^{(0)}} \rightarrow \mathcal{F}_{\mathcal{S}^{(1)}} \quad \text{and} \quad g^{(1)} : \mathcal{F}_{\mathcal{S}^{(1)}} \rightarrow \mathcal{F}_{\mathcal{R}}$$

via

$$h^{(1)}(\mathbf{b}_{v,l}) = \begin{cases} 1 & \text{if } \mathbf{b}_{v,l} \equiv_g 1 \\ \mathbf{s}_{v,j_v}^{(1)} & \text{if } \psi_{v,l} \sqsubseteq \varphi_{v,j_v}^{(1)} \end{cases} \quad (8.15)$$

$$g^{(1)}(\mathbf{s}_{v,j_v}^{(1)}) = \prod_{\substack{1 \leq i \leq n \\ v=b_i, D_{v,j_v}^{(1)} \leq A_i}} \mathbf{a}_i \quad (8.16)$$

Note that we omit the superscripts $+$ and $-$ for simplicity of notation. This is in accordance with (8.10), because $\mathbf{b}_{v,l}^+$ and $\mathbf{b}_{v,l}^-$ are dealt with in a completely symmetrical way, and (8.15) and (8.15) also hold for the corresponding quotients. $h^{(1)}$ is well-defined, because each $\mathbf{b}_{v,l} \not\equiv_g 1$ is contained in exactly one equivalence class $[\mathbf{b}_{v,l'}]_g$ and thus $\mathbf{b}_{v,l} \sqsubseteq \varphi_{v,j_v}^{(1)}$ for exactly one $\varphi_{v,j_v}^{(1)}$.

So $h^{(1)}$ models the transition from $\mathcal{S}^{(0)}$ to $\mathcal{S}^{(1)}$, and $g^{(1)}$ relates $\mathcal{F}_{\mathcal{S}^{(1)}}$ to $\mathcal{F}_{\mathcal{R}}$ as g does for $\mathcal{F}_{\mathcal{B}}$. The following lemma shows that a similar equation as given in Theorem 8.2.1 still holds:

Lemma 8.2.4. *Let $\mathcal{S}^{(1)}, h^{(1)}, g^{(1)}$ be defined as above. Then the following relationships hold:*

- (i) $g = g^{(1)} \circ h^{(1)}$.
- (ii) $\sigma_{\mathcal{S}^{(1)}} = h^{(1)} \circ \sigma_{\mathcal{S}^{(0)}}$.
- (iii) $g^{(1)} \circ \sigma_{\mathcal{S}^{(1)}} = \sigma_{\mathcal{R}}$.

Corollary 8.2.3. $\ker \sigma_{\mathcal{S}^{(0)}} \subseteq \ker \sigma_{\mathcal{S}^{(1)}} \subseteq \ker \sigma_{\mathcal{R}}$.

This first step usually reduces \mathcal{B} considerably and shows the general pattern of modifying the set of conditionals under consideration by defining appropriate homomorphisms $h^{(1)}$ and $g^{(1)}$, respectively. This will be pursued in the next step, too. As an important difference to Step 1, however, we will no longer deal with basic single-elementary conditionals. More general, $\mathcal{S}^{(1)}$ is a set of conditionals $\varphi_{v,j_v}^{(1)}$ with a single atom v in the conclusion, and the antecedent $D_{v,j_v}^{(1)}$ of $\varphi_{v,j_v}^{(1)}$ is a disjunction of elementary conjunctions not containing v .

Due to Lemma 8.2.4(i), we have

$$g\left(\prod_{1 \leq k \leq m} (\mathbf{b}_{v,l_k})^{r_k}\right) = 1 \quad \text{iff} \quad g^{(1)}\left(\prod_{1 \leq k \leq m} h^{(1)}(\mathbf{b}_{v,l_k})^{r_k}\right) = 1$$

Thus by replacing each \mathbf{b}_{v,l_k} by its image $h^{(1)}(\mathbf{b}_{v,l_k})$, we will now explore equivalence mod $\ker g^{(1)}$ between the generators $\mathbf{s}_{v,j_v}^{(1)}$ of $\mathcal{F}_{\mathcal{S}^{(1)}}$. Note that while neither \mathcal{R} nor g are known, the homomorphisms $h^{(t)}$ will approximate \mathcal{R} in a constructive way.

Step 2:

Step 1 already revealed the basic idea of the method to be presented: (Single-elementary) Conditionals are joined by \sqcup in accordance with the equations modulo $\ker g$. The second step also follows this idea for more complex equations. First we specify its starting point:

Prerequisites: Suppose $\mathcal{S}^{(t)}$ is a set of conditionals $\varphi_{v,j}^{(t)}$ with a single atom v in the conclusion, and the antecedent $D_{v,j}^{(t)}$ of $\varphi_{v,j}^{(t)}$ is a disjunction of elementary conjunctions not containing v . Let $\mathcal{F}_{\mathcal{S}^{(t)}} = \langle \mathbf{s}_{v,j}^{(t)+}, \mathbf{s}_{v,j}^{(t)-} \rangle_{v,j}$ be the free abelian group associated with $\mathcal{S}^{(t)}$, and let $g^{(t)} : \mathcal{F}_{\mathcal{S}^{(t)}} \rightarrow \mathcal{F}_{\mathcal{R}}$ be the homomorphism defined by

$$g^{(t)}(\mathbf{s}_{v,j}^{(t)}) = \prod_{\substack{1 \leq i \leq n \\ v = b_i, D_{v,j}^{(t)} \leq A_i}} \mathbf{a}_i$$

such that

$$g^{(t)} \circ \sigma_{\mathcal{S}^{(t)}} = \sigma_{\mathcal{R}}.$$

Let $\equiv_{g^{(t)}}$ mean $\equiv \text{mod } \ker g^{(t)}$.

In this step, we exploit equations of the form

$$\mathbf{s}_{v,j_0}^{(t)} \equiv_{g^{(t)}} \mathbf{s}_{v,j_1}^{(t)} \cdots \mathbf{s}_{v,j_m}^{(t)} \quad (8.17)$$

to modify $\mathcal{S}^{(t)}$ appropriately. To obtain this modified set $\mathcal{S}^{(t+1)}$,

1. eliminate $\varphi_{v,j_0}^{(t)}$ from $\mathcal{S}^{(t)}$;
2. replace each $\varphi_{v,j_k}^{(t)}$ by

$$\varphi_{v,j_k}^{(t+1)} = \varphi_{v,j_0}^{(t)} \sqcup \varphi_{v,j_k}^{(t)} = (v \mid D_{v,j_0}^{(t)} \vee D_{v,j_k}^{(t)}),$$

for $1 \leq k \leq m$. Set $D_{v,j_k}^{(t+1)} = D_{v,j_0}^{(t)} \vee D_{v,j_k}^{(t)}$, $1 \leq k \leq m$;

3. retain all other $\varphi_{w,l}^{(t)}$, i.e.

$$\varphi_{w,l}^{(t+1)} = \varphi_{w,l}^{(t)} \quad \text{for } w \neq v \text{ or } l \notin \{j_0, j_1, \dots, j_m\}.$$

This also includes the case $m = 0$, i.e. $\varphi_{v,j_0}^{(t)} \equiv_{g^{(t)}} 1$; in this case, (2) is vacuous and therefore is left out.

Define homomorphisms $h^{(t+1)} : \mathcal{F}_{\mathcal{S}^{(t)}} \rightarrow \mathcal{F}_{\mathcal{S}^{(t+1)}}$ and $g^{(t+1)} : \mathcal{F}_{\mathcal{S}^{(t+1)}} \rightarrow \mathcal{F}_{\mathcal{R}}$ by

$$h^{(t+1)}(\mathbf{s}_{w,l}^{(t)}) = \begin{cases} \prod_{1 \leq k \leq m} \mathbf{s}_{v,j_k}^{(t+1)} & \text{if } w = v, l = j_0 \\ \mathbf{s}_{v,j_k}^{(t+1)} & \text{if } w = v, l = j_k, 1 \leq k \leq m \\ \mathbf{s}_{w,l}^{(t+1)} & \text{else} \end{cases}$$

and

$$g^{(t+1)}(\mathbf{s}_{w,l}^{(t+1)}) = \prod_{\substack{1 \leq i \leq n \\ w=b_i, D_{w,l}^{(t+1)} \leq A_i}} \mathbf{a}_i.$$

We now prove a statement equivalent to Lemma 8.2.4 for this general case:

Lemma 8.2.5. *Let $\mathcal{S}^{(t+1)}, h^{(t+1)}, g^{(t+1)}$ be defined as above. Then the following relationships hold:*

- (i) $g^{(t+1)} \circ h^{(t+1)} = g^{(t)}.$
- (ii) $h^{(t+1)} \circ \sigma_{\mathcal{S}^{(t)}} = \sigma_{\mathcal{S}^{(t+1)}}.$
- (iii) $g^{(t+1)} \circ \sigma_{\mathcal{S}^{(t+1)}} = \sigma_{\mathcal{R}}.$

So the new set $\mathcal{S}^{(t+1)}$ is apt to continue the set chain in the sense of (8.11):

Corollary 8.2.4. *With the same notation as in Lemma 8.2.5, it holds that*

$$\ker \sigma_{\mathcal{S}^{(t)}} \subseteq \ker \sigma_{\mathcal{S}^{(t+1)}} \subseteq \ker \sigma_{\mathcal{R}} \quad (8.18)$$

By replacing each group element $\mathbf{s}_{v,l}^{(t)}$ by $h^{(t+1)}(\mathbf{s}_{v,l}^{(t)})$, equations holding modulo $\ker g^{(t)}$ are transformed into equations modulo $\ker g^{(t+1)}$:

$$g^{(t)}\left(\prod_k (\mathbf{s}_{v,l_k}^{(t)})^{r_k}\right) = 1 \quad \text{iff} \quad g^{(t+1)}\left(\prod_k h^{(t+1)}(\mathbf{s}_{v,l_k}^{(t)})^{r_k}\right) = 1,$$

due to Lemma 8.2.5(i).

By repeating step 2, the original equations modulo $\ker g$ are modified and solved, if possible, defining a sequence of sets $\mathcal{S}^{(t)}$ of conditionals such that

$$\ker \sigma_{\mathcal{S}^{(0)}} \subseteq \dots \subseteq \ker \sigma_{\mathcal{S}^{(t)}} \subseteq \ker \sigma_{\mathcal{S}^{(t+1)}} \subseteq \dots \subseteq \ker \sigma_{\mathcal{R}},$$

as desired, together with homomorphisms $g^{(t)}$ describing their relationship to \mathcal{R} .

Suppose that no further reduction of equations modulo $\ker g^{(t)}$ according to step 2 is possible, and the procedure halts. So we arrive at a set $\mathcal{S}^{(*)}$ of conditionals $\varphi_{v,j}^{(*)}$ with a single atom v in the conclusion, and the antecedent $D_{v,j}^{(*)}$ of $\varphi_{v,j}^{(*)}$ is a disjunction of elementary conjunctions not containing v . Let $\mathcal{F}_{\mathcal{S}^{(*)}} = \langle \mathbf{s}_{v,j}^{(*)+}, \mathbf{s}_{v,j}^{(*)-} \rangle_{v,j}$ be the free abelian group associated with $\mathcal{S}^{(*)}$, and let $g^{(*)} : \mathcal{F}_{\mathcal{S}^{(*)}} \rightarrow \mathcal{F}_{\mathcal{R}}$ be the homomorphism defined by

$$g^{(*)}(\mathbf{s}_{v,j}^{(*)}) = \prod_{\substack{1 \leq i \leq n \\ v=b_i, D_{v,j}^{(*)} \leq A_i}} \mathbf{a}_i$$

such that

$$g^{(*)} \circ \sigma_{\mathcal{S}^{(*)}} = \sigma_{\mathcal{R}}.$$

Now, one of the following two situations may occur:

- Either, there are still non-trivial equations modulo $\ker g^{(*)}$. In this case, $\mathcal{S}^{(*)}$ is only an approximation of $\ker \sigma_{\mathcal{R}}$, and $\ker \sigma_{\mathcal{S}^{(*)}}$ is included in it.
- Or, all equations modulo $\ker g$ could be solved successfully, so no non-trivial equations modulo $\ker g^{(*)}$ are left. That is, for any $\hat{\omega} \in \ker \sigma_{\mathcal{R}}$, $1 = \sigma_{\mathcal{R}}(\hat{\omega}) = g^{(*)} \circ \sigma_{\mathcal{S}^{(*)}}(\hat{\omega})$ holds trivially, i.e. due to $\sigma_{\mathcal{S}^{(*)}}(\hat{\omega}) = 1$. But this means $\ker \sigma_{\mathcal{R}} \subseteq \ker \sigma_{\mathcal{S}^{(*)}} \subseteq \ker \sigma_{\mathcal{R}}$, so

$$\ker \sigma_{\mathcal{S}^{(*)}} = \ker \sigma_{\mathcal{R}}$$

and the procedure ends up with full success.

In general, the techniques described in steps 1 and 2 will not suffice to eliminate all equations modulo $\ker g$, and we will be left with more complex equations modulo $\ker g^{(t)}$ of the form

$$\prod_k (\mathbf{s}_{v,j_k}^{(t)})^{r_k} \equiv_{g^{(t)}} \prod_l (\mathbf{s}_{v,j_l}^{(t)})^{s_l}, \quad (8.19)$$

all $r_k, s_l > 0$. The great variety of relationships possibly holding between the conditionals involved makes it difficult, if not impossible in general, to construct a new appropriate set $\mathcal{S}^{(t+1)}$ of conditionals in a straightforward way.

Nevertheless, the method developed so far already illustrates the central idea of how to find the conditionals whose structures a conditional valuation function V follows: By investigating relations between the numerical values of V , the effects of conditionals are analyzed and isolated, and conditionals are joined suitably so as to fit the conditional structures inherent to V . The operations on conditionals are based on equations between group elements representing these conditionals. So the formal framework for conditionals developed in Chapter 3 once again proved useful, providing the possibility of *calculating relevant conditionals* from e.g. probability distributions. We will illustrate how this works by considering examples in the next section.

Though at the present state, the method is not guaranteed to terminate successfully, we will find that in many cases, it will yield a useful approximation of the unknown set \mathcal{R} of conditionals. Treating equations of form (8.19) is a topic of our ongoing research, and results will be published in a further paper.

8.3 Examples – ME-Knowledge Discovery

We will now illustrate the method described in the previous section by two probabilistic examples. Given a probability distribution P , we will show how to calculate a set \mathcal{S} (or \mathcal{S}^{prob} , respectively) of (probabilistic) conditionals such that $P = P_0 *_{ME} \mathcal{S}^{prob}$, where $*_{ME}$ is the ME-operator and P_0 is a suitable uniform distribution. That is, we are going to solve what is known as the *inverse maxent problem*.

Due to the fact that P is necessarily indifferent with respect to such a set \mathcal{S} , analyzing the relationships between the numerical values in P will help to find such an \mathcal{S} , as is explained in the previous section and as will be carried out in the following. Note that by assuming P to be a faithful representation of some set \mathcal{R}^{prob} , i.e. $P(\hat{\omega}) = 1$ iff $\sigma_{\mathcal{R}}(\hat{\omega}) = 1$, we have $P(\hat{\omega}) = 1$ iff $\sigma_{\mathcal{B}}(\hat{\omega}) \equiv_g 1$, according to Corollary 8.2.1.

We consider formulas involving the three atomic propositions a, b, c , interpreted by

- a being a student
- b being young
- c being single (i.e. unmarried)

in two different settings, represented by two distributions. We list the twelve basic single-elementary conditionals $\psi_{v,l}$ of \mathcal{B} :

$$\begin{array}{lll} \psi_{a,0}=(a \mid \bar{b}\bar{c}) & \psi_{b,0}=(b \mid \bar{a}\bar{c}) & \psi_{c,0}=(c \mid \bar{a}\bar{b}) \\ \psi_{a,1}=(a \mid \bar{b}c) & \psi_{b,1}=(b \mid \bar{a}c) & \psi_{c,1}=(c \mid \bar{a}b) \\ \psi_{a,2}=(a \mid b\bar{c}) & \psi_{b,2}=(b \mid a\bar{c}) & \psi_{c,2}=(c \mid ab\bar{c}) \\ \psi_{a,3}=(a \mid bc) & \psi_{b,3}=(b \mid ac) & \psi_{c,3}=(c \mid abc) \end{array}$$

with corresponding generators $\mathbf{b}_{v,l}^+, \mathbf{b}_{v,l}^-$ of $\mathcal{F}_{\mathcal{B}}$.

Example 8.3.1. The first distribution P_1 over a, b, c is given as follows:

ω	$P_1(\omega)$	ω	$P_1(\omega)$
abc	0.3028	$\bar{a}bc$	0.2133
$ab\bar{c}$	0.0336	$\bar{a}b\bar{c}$	0.0237
$a\bar{b}c$	0.0421	$\bar{a}\bar{b}c$	0.1712
$a\bar{b}\bar{c}$	0.0421	$\bar{a}\bar{b}\bar{c}$	0.1712

By calculating ratios of probabilities of neighboring worlds, we observe immediately

$$P_1(\bar{a}\bar{b}c) = P_1(\bar{a}\bar{b}\bar{c}), \quad P_1(a\bar{b}c) = P_1(a\bar{b}\bar{c}), \quad P_1\left(\frac{abc}{ab\bar{c}}\right) = P_1\left(\frac{\bar{a}bc}{\bar{a}b\bar{c}}\right),$$

and we will now show how these numerical relationships can be used to calculate a set \mathcal{S} of conditionals that may impose such a structure on P .

So let

$$K_1 = \left\langle \frac{\bar{a}\bar{b}c}{\bar{a}\bar{b}\bar{c}}, \frac{\bar{a}\bar{b}c}{\bar{a}\bar{b}\bar{c}}, \frac{abc \cdot \bar{a}\bar{b}c}{\bar{a}\bar{b}\bar{c} \cdot \bar{a}\bar{b}c} \right\rangle$$

Step 0: We start with $\mathcal{S}^{(0)} = \mathcal{B}$; the generators of K_1 give rise to the following equations modulo $\ker g$:

$$\begin{aligned} 1 &\equiv_g \sigma_{\mathcal{B}} \left(\frac{\bar{a}\bar{b}c}{\bar{a}\bar{b}\bar{c}} \right) &= \frac{\mathbf{b}_{a,1}^- \mathbf{b}_{b,1}^- \mathbf{b}_{c,0}^+}{\mathbf{b}_{a,0}^- \mathbf{b}_{b,0}^- \mathbf{b}_{c,0}^-} \\ 1 &\equiv_g \sigma_{\mathcal{B}} \left(\frac{\bar{a}\bar{b}c}{\bar{a}\bar{b}\bar{c}} \right) &= \frac{\mathbf{b}_{a,1}^+ \mathbf{b}_{b,3}^- \mathbf{b}_{c,2}^+}{\mathbf{b}_{a,0}^+ \mathbf{b}_{b,2}^- \mathbf{b}_{c,2}^-} \\ 1 &\equiv_g \sigma_{\mathcal{B}} \left(\frac{abc \cdot \bar{a}\bar{b}c}{\bar{a}\bar{b}\bar{c} \cdot \bar{a}\bar{b}c} \right) &= \frac{\mathbf{b}_{a,3}^+ \mathbf{b}_{b,3}^+ \mathbf{b}_{c,3}^+ \cdot \mathbf{b}_{a,2}^- \mathbf{b}_{b,0}^+ \mathbf{b}_{c,1}^-}{\mathbf{b}_{a,2}^+ \mathbf{b}_{b,2}^+ \mathbf{b}_{c,3}^+ \cdot \mathbf{b}_{a,3}^- \mathbf{b}_{b,1}^+ \mathbf{b}_{c,1}^+} \end{aligned}$$

Considering these equations for each atom a, b, c separately and omitting the $\{+, -\}$ -signs, we obtain

$$\begin{aligned} \mathbf{b}_{a,1} &\equiv_g \mathbf{b}_{a,0}, & \mathbf{b}_{a,3} &\equiv_g \mathbf{b}_{a,2} \\ \mathbf{b}_{b,1} &\equiv_g \mathbf{b}_{b,0}, & \mathbf{b}_{b,3} &\equiv_g \mathbf{b}_{b,2} \\ \mathbf{b}_{c,0} &\equiv_g \mathbf{b}_{c,2} \equiv_g 1, & \mathbf{b}_{c,3} &\equiv_g \mathbf{b}_{c,1} \end{aligned}$$

(cf. Proposition 8.2.1 and Corollary 8.2.2).

Step 1: We eliminate the basic single-elementary conditionals $\psi_{c,0}$ and $\psi_{c,2}$ from $\mathcal{S}^{(0)} = \mathcal{B}$, and join conditionals according to the equations above; we obtain as conditionals in $\mathcal{S}^{(1)}$

$$\begin{aligned} \psi_{a,0} \sqcup \psi_{a,1} &=: \varphi_{a,0}^{(1)} = (a|\bar{b}) \\ \psi_{a,2} \sqcup \psi_{a,3} &=: \varphi_{a,1}^{(1)} = (a|b) \\ \psi_{b,0} \sqcup \psi_{b,1} &=: \varphi_{b,0}^{(1)} = (b|\bar{a}) \\ \psi_{b,2} \sqcup \psi_{b,3} &=: \varphi_{b,1}^{(1)} = (b|a) \\ \psi_{c,1} \sqcup \psi_{c,3} &=: \varphi_{c,1}^{(1)} = (c|b) \end{aligned}$$

These are all conditionals in $\mathcal{S}^{(1)}$, with corresponding elements $\mathbf{s}_{v,l}^{(1)} \in \mathcal{F}_{\mathcal{S}^{(1)}}$. All of the equations modulo $\ker g$ set up in step 0 are transformed into trivial equations modulo $\ker g^{(1)}$. Calculating the probabilities of these conditionals in P_1 , we obtain

$$\begin{aligned} P_1(a|\bar{b}) &\approx 0.2; & P_1(a|b) &\approx 0.6; \\ P_1(b|\bar{a}) &\approx 0.4; & P_1(b|a) &\approx 0.8; \\ P_1(c|b) &\approx 0.9. \end{aligned}$$

By using an ME-tool (like SPIRIT, cf. [RKI97b]), we see that actually

$$P_1 = P_0 *_{ME} \{(a|\bar{b})[0.2], (a|b)[0.6], (b|\bar{a})[0.4], (b|a)[0.8], (c|b)[0.9]\},$$

and these conditionals represent knowledge incorporated by P which is relevant, in particular, with respect to ME-inference.

Example 8.3.2. The second distribution P_2 over a, b, c is given as follows:

ω	$P_2(\omega)$	ω	$P_2(\omega)$
abc	0.1950	$\bar{a}bc$	0.1528
$ab\bar{c}$	0.1758	$\bar{a}b\bar{c}$	0.1378
$a\bar{b}c$	0.0408	$\bar{a}\bar{b}c$	0.1081
$a\bar{b}\bar{c}$	0.0519	$\bar{a}\bar{b}\bar{c}$	0.1378

Here important relationships between probabilities are revealed by

$$P_2(\bar{a}b\bar{c}) = P_2(\bar{a}\bar{b}\bar{c}), \quad P_2\left(\frac{abc}{ab\bar{c}}\right) = P_2\left(\frac{\bar{a}bc}{\bar{a}b\bar{c}}\right), \quad P_2\left(\frac{a\bar{b}c}{a\bar{b}\bar{c}}\right) = P_2\left(\frac{\bar{a}\bar{b}c}{\bar{a}\bar{b}\bar{c}}\right),$$

so that

$$K_2 = \left\langle \frac{\bar{a}b\bar{c}}{\bar{a}\bar{b}\bar{c}}, \frac{abc \cdot \bar{a}b\bar{c}}{ab\bar{c} \cdot \bar{a}bc}, \frac{\bar{a}bc \cdot \bar{a}\bar{b}\bar{c}}{\bar{a}\bar{b}\bar{c} \cdot \bar{a}bc} \right\rangle$$

Again we start with $\mathcal{S}^{(0)} = \mathcal{B}$. The generators of K_2 yield the following equations modulo $\ker g$:

$$\begin{aligned} 1 &\equiv_g \sigma_{\mathcal{B}} \left(\frac{\bar{a}b\bar{c}}{\bar{a}\bar{b}\bar{c}} \right) = \frac{\mathbf{b}_{a,2}^- \mathbf{b}_{b,0}^+ \mathbf{b}_{c,1}^-}{\mathbf{b}_{a,0}^- \mathbf{b}_{b,0}^- \mathbf{b}_{c,0}^-} \\ 1 &\equiv_g \sigma_{\mathcal{B}} \left(\frac{abc \cdot \bar{a}b\bar{c}}{ab\bar{c} \cdot \bar{a}bc} \right) = \frac{\mathbf{b}_{a,3}^+ \mathbf{b}_{b,3}^+ \mathbf{b}_{c,3}^+ \cdot \mathbf{b}_{a,2}^- \mathbf{b}_{b,0}^+ \mathbf{b}_{c,1}^-}{\mathbf{b}_{a,2}^+ \mathbf{b}_{b,2}^+ \mathbf{b}_{c,3}^- \cdot \mathbf{b}_{a,3}^- \mathbf{b}_{b,1}^+ \mathbf{b}_{c,1}^+} \\ 1 &\equiv_g \sigma_{\mathcal{B}} \left(\frac{\bar{a}bc \cdot \bar{a}\bar{b}\bar{c}}{\bar{a}\bar{b}\bar{c} \cdot \bar{a}bc} \right) = \frac{\mathbf{b}_{a,1}^+ \mathbf{b}_{b,3}^- \mathbf{b}_{c,2}^+ \cdot \mathbf{b}_{a,0}^- \mathbf{b}_{b,0}^- \mathbf{b}_{c,0}^-}{\mathbf{b}_{a,0}^+ \mathbf{b}_{b,2}^- \mathbf{b}_{c,2}^- \cdot \mathbf{b}_{a,1}^- \mathbf{b}_{b,1}^- \mathbf{b}_{c,0}^+} \end{aligned}$$

Considering these equations for each atom a, b, c separately and omitting the $\{+, -\}$ -signs, we obtain

$$\begin{aligned} \mathbf{b}_{a,0} &\equiv_g \mathbf{b}_{a,1} \equiv_g \mathbf{b}_{a,2} \equiv_g \mathbf{b}_{a,3} \\ \mathbf{b}_{c,0} &\equiv_g \mathbf{b}_{c,1} \equiv_g \mathbf{b}_{c,2} \equiv_g \mathbf{b}_{c,3} \\ \mathbf{b}_{b,0} &\equiv_g 1 \\ \mathbf{b}_{b,3} &\equiv_g \mathbf{b}_{b,1} \mathbf{b}_{b,2} \end{aligned}$$

(cf. Proposition 8.2.1 and Corollary 8.2.2).

Eliminating and joining conditionals according to steps 1 and 2, we obtain

$$\mathcal{S}_2 = \{(a|\top), (c|\top), (b|a), (b|c)\}$$

and

$$\mathcal{S}_2^{prob} = \{(a|\top)[0.4635], (c|\top)[0.4967], (b|a)[0.8], (b|c)[0.7]\}$$

represents P_2 via ME-inference, i.e. $P_2 = P_0 *_{ME} \mathcal{S}_2^{prob}$.

8.4 Open Problems and Concluding Remarks

The approach to ME-optimal knowledge discovery developed in this chapter makes an important new contribution to the field; nevertheless, a lot of work remains to be done:

1. How are the crucial relationships inherent to a given P found? It seems advisable to investigate the orders of magnitudes of the probabilities instead of precise rational numbers. In many cases, considering ratios $\frac{P(\omega)}{P(\omega')}$ with neighboring worlds ω, ω' will help to find important relationships.
2. The sets of conditionals discovered in the examples 8.3.1 and 8.3.2 are not really ME-optimal because they contain redundant conditionals: It is straightforward to check that $P_1 = P_0 *_{ME} \{(b|a)[0.8], (c|b)[0.9]\}$ and $P_2 = P_0 *_{ME} \{(b|a)[0.8], (b|c)[0.7]\}$. Eliminating redundant conditionals from the resulting set of conditionals is still an open problem.
3. Last not least, more complex equations of the type (8.19) still have to be dealt with.

We are, however, optimistic in that the method presented here may be extended to also treat the more difficult equations in problem 3, and that it can be modified appropriately to yield an even more effective algorithm, tailor-made to ME-propagation and avoiding redundant conditionals, thus solving problem 2.

The applicability of the method presented in this chapter neither depends on the presupposition of V being a faithful c-representation nor on having a complete description of $\ker V$ available: Each numerical relationship found amongst the values of V corresponds to an element of $\ker V$ and may be used to set up equations for the group elements in $\mathcal{F}_{\mathcal{B}}$ modulo $\ker g$. The generators of $\ker V$ are particularly appropriate for this task, in that they yield basic equations, but any other element will do, too. If V fails to be a faithful c-representation of some suitable set of conditionals, then too many equations modulo $\ker g$ will have to be solved trivially. In this case, backtracking will be necessary, undoing the last joining of conditionals.

Moreover, the techniques of Section 8.2 can also be applied when a prior epistemic state has to be taken into account. For instance in a probabilistic framework, when P is the actual distribution to be investigated and P' is some prior distribution, we can compute a set $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$ with $P' *_{ME} \mathcal{R} = P$ in the same way as above, namely by applying the algorithm to the normalized function P/P' .

9. Algorithms and Implementations

In this chapter, we present a selection of various computational approaches to quantified uncertain reasoning and probabilistic knowledge discovery.

9.1 ME-Reasoning with Probabilistic Conditionals

Probability theory provides a powerful and mathematically founded, non-heuristic framework for uncertain reasoning, but, due to their exponential complexity, probability distributions are not easily dealt with. Efficient algorithms are necessary to represent probabilistic dependencies between a large number of variables or atoms, respectively, and to incorporate new information so as to achieve a revised or instantiated probabilistic state of belief.

To reduce the complexity of probability distributions, ME-reasoning can make use of so-called *LEG-networks*, where LEG stands for *local event group* (cf. [Lem82, Lem83]). LEG-networks are hypergraphs with its hyperedges (*LEG's*) consisting of sets of atomic propositions (or events, or propositional variables, respectively). To each LEG, a *component marginal distribution* is associated. Like the clique trees of Bayesian networks (cf. [LS88, Nea90, Jen96]), they allow local computations and propagations of probabilities.

The expert system shell *SPIRIT*¹ uses LEG-networks for representing sets of probabilistic rules and reasoning via the principle of maximum entropy (cf. [RM96, RKI97a, RKI97b, Mey98]). Given a set of probabilistic conditionals \mathcal{R} , a hypergraph is constructed whose hyperedges consist exactly of variables occurring in one conditional in \mathcal{R} , respectively. Usually, this hypergraph fails to be acyclic, so a covering hypertree is generated which allows a decomposition of the associated ME-distribution $P^* = P_0 *_{ME} \mathcal{R}$ (cf. [Lem83, Mal89]). Learning of the conditionals in \mathcal{R} is done locally on the LEG's by approximating iteratively the Lagrange factors α_i (cf. (5.5), p. 76) which yield a potential representation of P^* . The global propagation is

¹ available at <http://www.fernuni-hagen.de/BWLOR/forsch.htm>

then carried out by successively adjusting neighboring LEG's by a procedure similar to *iterative proportional scaling* (cf. [Lau82, MR92]), passing through the structure of the hypertree. Answering queries is done by modifying the hypertree appropriately and propagating the values of the instantiated variables (or atoms, respectively). It is worth noticing that the Lagrange factors α which are so meaningful for the theoretical results presented in this book are also of crucial importance for an efficient computation. For a detailed description of the algorithm, cf. [Mey98].

So LEG-networks provide an efficient method to perform local computations for ME-reasoning, reducing its complexity (cf. [Par94]). Their meaning is similar to that of Bayesian networks for probabilistic reasoning in general. There are, however, crucial differences between Bayesian reasoning and ME-reasoning:

- Instead of imposing external assumptions of (conditional) independencies on the variables, ME-reasoning follows the internal structure of conditionals to install independencies intensionally: Independence is only assumed when no other information is available.
- ME-reasoning does not require the specification of large amounts of conditional probabilities to build up a Bayesian network. It rather fills up the necessary values in an information-theoretically optimal manner. Instead, it offers the possibility of specifying knowledge by conditionals in an intuitive way, requiring only to list and quantify relationships which are considered relevant by the user.

Thus, ME-reasoning seems to combine ideally sound probabilistic reasoning with intuitive knowledge representation, providing a powerful machinery to realize commonsense and expert reasoning in demanding domains like medical diagnosis. Another ME-system, LEXMED², is already used to support physicians in diagnosing appendicitis in a German hospital (cf. [SE99]). LEXMED is based on the system shell PIT ([FS96, SF97]). PIT not only accepts precise probabilities, but also allows one to specify intervals of probability values for the conditionals. An approach to combine ME-reasoning with probabilistic logic programming techniques is presented in [LKI99].

9.2 Probabilistic Knowledge Discovery

Within the field of probabilistic knowledge discovery, conditionals of a simple syntax have proved to be of particular importance. These conditionals

² Homepage of LEXMED: <http://lexmed.fh-weingarten.de>

usually have conjunctions of atoms or literals in their antecedents and consequents, avoiding disjunctions in order to pinpoint relevant information. *Association rules* (cf. [AIS93]) only make use of positive literals, whereas *single-elementary conditionals* have only one (positive) literal as its consequence (cf. Definition 8.1.1).

In [KIR96], we dealt with discovering *relevant* and *significant* single-elementary rules in a given probability distribution P . There, the *significance* of a probabilistic rule is measured simply by its probability with respect to a threshold ϵ : $(B|A)$ is called (ϵ) -*significant* if $P(B|A) > 1 - \epsilon$. And *relevance* aims at presenting relationships in a concise way, that is, by shortening the (conjunctive) antecedent of a rule without giving up significance. To find significant single-elementary conditionals, one has to check for each elementary conjunction the corresponding ratios of probabilities to its neighboring conjunctions. To introduce the notion of relevance, we gave a criterion which elementary conjunctions should be investigated to bring forth rules with a particular short antecedent. The level of significance ϵ can be chosen by the user, so as to allow investigations on different levels of abstraction.

The algorithm presented in [KIR96] starts with “long” rules, successively shortening the elementary conjunctions under consideration. In contrast to this, the algorithm in [Sch98] begins with short conjunctions, extending the antecedents of rules in search of exceptions. The implemented program in [Sch98] also checks the quality of the set of discovered probabilistic rules by measuring the information-theoretical distance between the corresponding ME-distribution and the original distribution P . In [Mül98], the idea of a *structural interestingness* of probabilistic rules is discussed, and a program to read data from databases and to find association rules was implemented.

9.3 Possibilistic Belief Revision

Another approach to realize quantified uncertain reasoning is made in [Hoc99] by means of *possibilistic logic*: Instead of assigning *one* degree of uncertainty to each proposition, as in probabilistic logic, possibilistic logic allows one to specify the epistemic attitude towards a proposition by *two* values, a degree of *necessity* and a degree of *possibility* which are usually assumed to range within the unit interval (cf. [DLP94]). Possibility and necessity measures are both determined by *possibility distributions* $\pi : \Omega \rightarrow [0, 1]$. Possibilistic logic aims at capturing qualitative epistemic relationships between propositions. Indeed, possibility distributions are very similar to ordinal conditional functions (cf. [DP94]) – in particular, the degree of possibility of a disjunction is the maximum of the degrees of possibility of the disjuncts.

In [Hoc99], a program was implemented that computes a representation of a set of propositional formulas, each equipped with a degree of necessity, by a possibility distribution. The resulting knowledge base can be modified via the belief change operations of *expansion*, *revision* and *contraction* (cf. Section 2.2). In possibility theory, a deduction theorem similar to that in classical logics holds, allowing the program to derive (new) possibilistic knowledge by making use of a *possibilistic resolution calculus*. As a special feature of possibilistic logic, however, the program is capable of tolerating inconsistencies, and to take the *degree of inconsistency* of a knowledge base into regard when deriving possibilistic information. All the theoretical background for possibilistic deduction and possibilistic change operations is explained in [Hoc99] in detail.