



HAL
open science

Confidence relations as a basis for uncertainty modeling, plausible reasoning and belief revision

Didier Dubois

► **To cite this version:**

Didier Dubois. Confidence relations as a basis for uncertainty modeling, plausible reasoning and belief revision. 14th Australian Joint Conference on Artificial Intelligence: Advances in Artificial Intelligence (AI 2001), Dec 2001, Adelaïde, Australia. pp.655-664. hal-03463135

HAL Id: hal-03463135

<https://hal.science/hal-03463135>

Submitted on 2 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Confidence Relations as a Basis for Uncertainty Modeling, Plausible Reasoning, and Belief Revision

Didier Dubois

IRIT-CNRS,
Université Paul Sabatier,
31062 Toulouse Cedex, France
dubois@irit.fr

Abstract. The aim of this position paper is to outline a unified view of plausible reasoning under incomplete information and belief revision, based on an ordinal representation of uncertainty. The information possessed by an agent is supposed to be made of three items: sure observations, generic knowledge and inferred contingent beliefs. The main notion supporting this approach is the confidence relation, a partial ordering of events which encodes the generic knowledge of an agent. Plausible inference is achieved by conditioning. The paper advocates the similarity between plausible reasoning with confidence relations and probabilistic reasoning. The main difference is that the ordinal approach supports the notion of accepted beliefs forming a deductively closed set, while probability theory is not tailored for it. The framework of confidence relations sheds light on the connections between some approaches to non-monotonic reasoning methods, possibilistic logic and the theory of belief revision. In particular the distinction between revising contingent beliefs in the light of observations and revising the confidence relation is laid bare.

1 Introduction

The aim of this position paper is to present a synthetic view of an approach to plausible reasoning under incomplete information with an ordinal representation of uncertainty. This approach has close connections with various works carried out more or less independently by philosophers like Ernest Adams, David Lewis, in the seventies, and Peter Gärdenfors and colleagues in the eighties, as well as several AI researchers, such as Yoav Shoham, Daniel Lehmann, Judea Pearl, Joe Halpern, Maryanne Williams, and others. Indeed, the issue of an ordinal approach to uncertainty has to do with several important topics of theoretical Artificial Intelligence, such as non-monotonic reasoning, belief revision, conditional logics, and to probabilistic reasoning as well. It seems that the ultimate aim of symbolic AI in plausible reasoning is to perform a counterpart to probabilistic inference without probabilities (Dubois and Prade 1994a).

Suppose an agent who has to reason about the current state of the world. In order to support the above thesis, we claim that the issue of plausible reasoning cannot be properly addressed without assuming that the body of information possessed by the agent contains (at least) three distinct types of items: observations pertaining to the current situation, generic knowledge about similar situations, and beliefs as to the

features of the current situation. Observations are supposed to be reliable and non conflicting. Beliefs are on the contrary taken for granted, hence brittle. Under this assumption, plausible reasoning precisely consists in inferring beliefs from (contingent) observations, and generic (background) knowledge, valid across situations. This view is classical in probability theory (De Finetti, 1974), and we claim that it also makes sense under a qualitative or an ordinal approach to plausible reasoning

2 From Confidence Relations to Accepted Beliefs

The ordinal approach presupposes that the agent's knowledge is modeled by a relation among events or propositions (built from a language), we call a *confidence relation*. Typically, it is a partial preordering on a set of propositions, that is consistent with classical deduction, expressing that some propositions are generally more plausible than (or at least as plausible as) others. The contingent observations available to the agent form a context according to which the confidence relation is conditioned. Confidence relations include comparative probability relations first introduced by De Finetti (1937) and Koopman (1940), and extensively studied by Savage (1954). All set-functions used in uncertainty modeling (probability measures, possibility measures, belief functions, etc.) generate confidence relations which are complete preorders. The set-functions studied by Friedman and Halpern (1996) under the name of "plausibility measures" generate confidence relations which are partial preorders.

A proposition is called an accepted belief for the agent if it is more plausible (in the sense of the confidence relation) than its negation, in the context of available observations. The term "accepted belief" also means that the agent considers the derived conclusions as valid, until some further observation is obtained that questions them; lastly the agent is allowed to reason with accepted beliefs as if they were true, using classical logic. Hence, by assumption, accepted beliefs form a deductively closed set of propositions. Under this assumption, a confidence relation is called an *acceptance relation* (Dubois and Prade, 1995b; Dubois et al. 1998a).

This logical closure condition for accepted beliefs has a drastic impact on the nature of acceptance relations. Basically it implies that, when a proposition A is more plausible than each of two other ones B and C, where A, B, C are mutually exclusive, considering the disjunction of B and C cannot form a proposition that is more plausible than the most plausible one A. In other words, a notion of negligibility is embedded in the acceptance relation. The closure condition, plus a few other uncontroversial ones (like monotony with respect to set inclusion) are enough to ensure the existence of a representation of the acceptance relation by means of a family of so called "comparative possibility relations" (Dubois et al. 2001). A proposition is then more plausible than another one if and only if the former is more possible than the latter in the sense of all comparative possibility relations in the family.

Comparative possibility relations have been independently introduced by David Lewis (1973) in the seventies, in the framework of modal logics of counterfactuals, and Dubois (1986) in the scope of decision theory. Their numerical counterparts have been introduced by the economist Shackle (1961), the philosopher Cohen (1977), and the systems engineer Zadeh (1978) completely independently of one another. Comparative possibility relations are very simple confidence relations because each of them is completely specified by means of a single complete preordering of elementary events

(interpretations, states of nature, possible worlds) distinguishing between normal and less normal worlds. Namely a proposition is more possible than another if the most normal situation where the former is true is more plausible than all situations where the latter is true. The case when the ordering of elementary events is partial is studied by Halpern (1997).

3 Nonmonotonic Reasoning and Default Rules

In practice, the generic knowledge possessed by an agent is often expressed by means of "if then" rules. The condition part of a rule denotes a context (in the above sense: everything the agent has observed in a given situation) and the conclusion is an accepted belief of the agent in this context. Each such rule can thus be modeled as the statement that some proposition is more plausible than another one, and a rule base can be equated to a (partially defined) acceptance relation (or a plausibility measure after Friedman and Halpern, 1996).

Each rule can also be modeled in the framework of a three-valued logic (Dubois and Prade, 1994b). In a given situation, a rule is true or false according to whether its conclusion is true or false, provided that its condition holds in this situation. Otherwise the rule takes the third truth-value which stands for "irrelevant". This is a so-called tri-event introduced by De Finetti (1937). Such generic rules form conditional knowledge bases, and the plausible inference of some proposition consists of syntactically deriving from a rule base a rule whose condition part exactly models the set of available observations, and whose conclusion part is the proposition under concern (Kraus et al, 1991). In the formal framework of acceptance relations, this syntactic inference procedure yields a plausible proposition if and only if this proposition is an accepted belief in the prescribed context, in the sense of the acceptance relation (plausibility measure after Friedman and Halpern, 1996) induced by the rule base

Indeed, under the logical closure condition, the plausible inference relation producing accepted beliefs according to a confidence relation also satisfies all postulates of preferential inference introduced by Kraus et al. (1991) for the purpose of computing what is entailed from a conditional knowledge base (except for the inference from a contradictory context), and conversely these postulates enable the confidence relation to be reconstructed. Similar properties have been laid bare in older conditional logics by Adams (1975), using infinitesimal probabilities, and, of course, Lewis (1973). A rule base also generates a family of comparative possibility relations ("rankings of models" after Lehmann and Magidor, 1992; see Dubois and Prade, 1995a).

If the family of comparative possibility relations reduces to a single one, then the plausible relation satisfies the so-called rational monotony property introduced by Makinson. This feature is characteristic of comparative possibility relations (Benferhat et al., 1997). Plausible inference with a comparative possibility relation meets Shoham (1988)'s view of nonmonotonic inference, as classical inference from the most normal situations in a given context. Plausible inference under a comparative possibility relation can be syntactically managed in possibilistic logic (Dubois et al., 1994; Lang 2000). When an acceptance relation corresponds to a family of more than one comparative possibility relations, there is a principle of information minimization that enables a unique comparative possibility relation in the family to be selected (Dubois and Prade, 1998). It is a most cautious choice ensuring a ranking of elementary events

that is as compact as possible. This selection process is at work in Pearl (1990)'s system Z and Lehmann and Magidor (1992)'s "rational closure", as well the possibilistic handling of default rule bases (Benferhat et al, 1998). Selecting a least informative comparative possibility relation in agreement with an acceptance relation is equivalent to attaching priorities to rules in a conditional knowledge base, the higher priorities being granted to the most specific rules (Pearl, 1990).

4 Plausible Inference versus Probabilistic Reasoning

The originality of the confidence relation approach to plausible inference is that, rather than starting from syntactic objects and intuitive postulates (like Lehmann and colleagues), our starting points are on the one hand the confidence relation that is thought of as a natural tool for describing an agent's uncertain knowledge, and the notion of accepted belief on the other hand. This point of view enables plausible (non-monotonic) reasoning to be cast in the general framework of uncertain reasoning, which includes probabilistic reasoning. The analogy between plausible reasoning with accepted beliefs and probabilistic reasoning is now patent (see also Paris, 1994). In probabilistic reasoning, the confidence relation stems from a probability measure or a family thereof. A set of generic rules is then encoded as a set of conditional probabilities characterizing a family of probability measures. The most popular approach in AI is currently when this family reduces to a single one, and the set of conditional probabilities defines a Bayesian network (Pearl, 1988). When the probabilistic information is incomplete, the selection of a unique probability measure often relies on the principle of maximal entropy. A Bayesian network really represents generic knowledge, like any confidence relation. This network derives either from expert domain knowledge or from statistical data. The selection of a most cautious comparative possibility relation in agreement with an acceptance relation is similar to the selection of a unique probability measure using maximal entropy (Paris, 1994).

Probabilistic inference with a Bayesian network consists in calculating the (statistical) conditional probability of a conclusion, where the conditioning event encodes the available observations (Pearl, 1988). The obtained conditional probability value is interpreted as the degree of belief of the conclusion in the current situation, assuming that this situation is a regular one in the context described by the observations. This procedure is very similar to the derivation of a plausible conclusion by conditioning an acceptance relation, or by deducing a rule from a rule base. The derived rule is valid "generally". Its conclusion is considered as an accepted belief in the current situation assuming that this situation is not an exceptional one in the context described by the observations modeled by the condition part of the derived rule. There is in fact a strong similarity between conditional probability and conditional possibility, and an ordinal form of Bayes rule exists for possibility theory (Dubois and Prade, 1998).

Of course, there are also noticeable differences between probabilistic reasoning and ordinal plausible inference:

- i) The latter does not quantify belief;
- ii) Plausible reasoning considers the most plausible situations and neglects others, while probability theory performs reasoning in the average.

iii) Lastly, probabilistic reasoning is not compatible with the notion of accepted belief.

Indeed, the conjunction of two highly probable events may fail to be highly probable and may even turn out to be very improbable. However, the arbitrary conjunction of accepted beliefs is still an accepted belief (this is because we assume that the agent considers accepted beliefs as tentatively true). This property of ordinal plausible inference has been criticized by several authors (Kyburg 1988, Poole 1991). It means that ordinal plausible inference suffers from the so-called "lottery paradox" (one can believe that any given player in a one-winner lottery game will lose with arbitrary high probability, all the more so as players are numerous, but one cannot believe that all of them will lose).

Yet, an acceptance relation can also be represented by means of a family of standard probability relations (Benferhat et al., 1999a, Snow, 1999). The corresponding probability measures are very special. They enforce a total ordering of states and are such that the probability of a state is always larger than the sum of the probabilities of less probable states. We call them "big-stepped probabilities". They are in some sense the total opposite of uniformly distributed ones (without expressing pure determinism, though). Indeed, in any context the most likely elementary event in the sense of a big-stepped probability occurs much more often than the disjunction of other elementary events.

We cannot expect to find natural sample spaces equipped with such kinds of empirically observed statistical probability functions. But one may think that for phenomena which have significant regularities, without being purely deterministic (like birds flying!), there may exist at least one partition of the sample space, the elements of which can be ordered via a big-stepped probability, and form a set of conceptually meaningful states for the agent. The existence of probabilities in strict agreement with non-monotonic inference, may also resolve the lottery paradox, that has been proposed as a counterexample to the use of classical deduction on accepted beliefs. Indeed, for big-stepped probabilities (and only for them), the set of probable beliefs $\{A, P(A/C) > 0.5\}$ remains consistent and deductively closed for any context C . In the lottery example, it is implicitly assumed that all players have equal chance of winning. The underlying probability is uniform. Hence there is no regularity at all in the lottery game: no particular occurrence is typical and randomness prevails. It is thus unlikely that an agent can come up with a set of consistent default rules about the lottery game.

On the contrary, plausible reasoning based on acceptance relations models an agent's reasoning in front of phenomena which have very regular features (but where exceptional situations may nevertheless occur). We conjecture that domains where a body of default knowledge exists can be statistically modeled by big-stepped probabilities on a meaningful partition of the sample space. If this conjecture is valid, it points out a potential link between non-monotonic reasoning and statistical data, in a knowledge discovery perspective. An open problem along this line is as follows: Given statistical data on a sample space, find the "best" partition(s) of the sample space, on which big-stepped probabilities are induced and meaningful default rules can be extracted (See Benferhat et al. 2001b for preliminary results). The difference between other rule extraction techniques and the one suggested here, is that, in our view, the presence of exceptions is acknowledged in the very definition of symbolic rules for which the proportion of such exceptions is not explicit.

5 Two Kinds of Epistemic Revision

The framework of confidence relations can also account for the AGM revision theory after Alchourron et al (1985), but it somehow questions the idea of Gärdenfors and Makinson (1988) that revision and plausible inference are two sides of the same coin. Indeed, in Gärdenfors revision theory, the agent only possesses a closed set of propositions (a belief set) and receives a sure input that can be understood as a new observation of the (static) world. So, the Gärdenfors revision theory only accounts for the evolution of the beliefs of an agent who makes a new contingent observation.

However, this is only one possible kind of epistemic revision. The distinction between contingent beliefs and generic knowledge forces to consider another meaning of revision: the revision of the generic knowledge, which is not the topic of the AGM theory. It consists of modifying the confidence relation upon arrival of new generic knowledge, for instance when the agent happens to acquire a new default rule on his domain of investigation. Then events that were thought to normally occur are now considered less normal. For instance a medical doctor does not modify his medical knowledge when he gets new test results for a patient. He just revises his beliefs about the patient state. However, a medical doctor may revise his medical knowledge when he reads a specialized book or attends a medicine conference. Several authors like Spohn (1988), Williams (1994), Boutilier and Goldszmidt (1995), Darwiche and Pearl (1997), Dubois and Prade (1997), Benferhat et al.(1999c) have considered tools and principles for generic knowledge revision, although the distinction between the two types of revision is not always so clear from reading these works. Indeed, there is no consensus on a general and systematic approach to that kind of epistemic change in the literature, and the same can be observed for problems of revision of Bayesian networks (which pertain to probability kinematics, see Domotor, 1985).

The AGM revision theory only assumes that a belief set is replaced by another belief set, and it gives minimal rationality constraints relating the prior and the posterior belief sets. Thus doing, it may wrongly suggest that the posterior belief set can indeed be derived from the prior one and the input information only. The confidence relation framework shows that this is not the case. The calculation of the posterior belief set does not use the prior belief set. The posterior belief set is built by means of plausible inference from the generic knowledge encoded in the confidence relation conditioned on the new context formed by all the available observations, including the new one. This is what we called "focusing" in previous publications (Dubois et al., 1998b).

The representation theorem of the AGM theory actually lays bare the existence of an epistemic entrenchment relation (basically the dual of a possibility relation, see Dubois and Prade, 1991) and confirms that the construction of the revised belief set can be expressed by conditioning this particular confidence relation on the input information. This strategy is the same as the one adopted when querying a Bayesian net on the basis of new observations. However, in the AGM theory, the epistemic entrenchment looks like a technical by-product of the formal construction, while we claim that this is the primitive object, and that all the belief sets are derived from it in every context. Concerning the iteration of contingent belief revision, suppose two inputs are obtained in a row. Note that since the inputs are considered as sure observations about a static world, they cannot be contradictory. When the second input arrives, a sound strategy is, in the AGM setting, to revise the original belief set (not the one

revised by the first input) by the conjunction of the two inputs. In practice, The case when two observations are inconsistent suggests that, either one of them is wrong, or they do not pertain to the same case.

Some people have claimed that in order to iterate AGM belief revision, one needs to construct not only the new belief set, but also a new epistemic entrenchment relation. In the scope of the revision of contingent beliefs induced by generic knowledge, this is questionable. On the contrary, the same epistemic entrenchment should remain across successive revisions of contingent beliefs caused by new contingent observations. Similarly, in probabilistic reasoning (Pearl, 1988), the same Bayesian network is used when new observations come in. If the epistemic entrenchment must be revised, it means that the input information is a piece of generic knowledge, and such a kind of revision is not the purpose of the AGM theory.

6 Conclusion

To sum up, the framework of confidence relations provides a unified view of non-monotonic and probabilistic reasoning. It also points out the distinction between the revision of contingent beliefs (by focusing the confidence relations on the proper context formed by the observations) and the revision of the confidence relation itself. This distinction is made clear by considering that the information possessed by an agent is made of three items: sure observations, generic knowledge and inferred contingent beliefs. From a computational point of view, plausible inference from a confidence relation can be achieved using a standard theorem-prover in propositional logic, and comes down to a sequence of consistency tests. When the confidence relation takes the form of a unique possibility relation, like in system Z and the like, the problem can be encoded in possibilistic logic, which handles prioritized propositional bases, with a complexity of $SAT * \text{Log}_2 n$ if there are n priority levels (Lang, 2001).

Future lines of research in the ordinal approach to plausible reasoning include the modeling of independence (Dubois et al. 1997, Ben Amor et al. 2000) and the study of graphical models that would be the qualitative counterpart of Bayesian networks (Benferhat et al. 1999b). Some results indicate that possibilistic logic bases, conditional knowledge bases and possibilistic nets have the same expressive power (Benferhat et al., 2001a). However it is no clear which is the most natural framework for knowledge elicitation, and for practical computation. Lastly, by bridging the gap between probability and non-monotonic reasoning, the confidence relation approach paves the way to the data-driven learning of default rules.

Acknowledgements. This paper owes much to Salem Benferhat, H el ene Fargier, J erome Lang et Henri Prade and reflects our shared concerns and joint past works.

References

- E.W. Adams (1975) *The Logic of Conditionals*. Dordrecht: D. Reidel.
- C.E.P. Alchourr on, P. G ardenfors, D. Makinson (1985), On the logic of theory change: Partial meet functions for contraction and revision. *J. of Symbolic Logic*, 50, 510-530.

- N. Ben Amor, S. Benferhat, D. Dubois, H. Geffner and H. Prade (2000) Independence in qualitative uncertainty frameworks. *Proc. 7th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2000)*, Breckenridge, Co, 235-246.
- S. Benferhat, D. Dubois, H. Prade (1997) Nonmonotonic reasoning, conditional objects and possibility theory, *Artificial Intelligence*, 92, 259-276
- S. Benferhat, D. Dubois, H. Prade (1998) Practical Handling of Exception-tainted rules and independence information in possibilistic logic. *Applied Intelligence*. 9, p. 101-127.
- S. Benferhat, D. Dubois, H. Prade (1999a) Possibilistic and standard probabilistic semantics of conditional knowledge. *J. Logic & Computation*, 9, 873-895
- S. Benferhat, D. Dubois, L. Garcia, H. Prade (1999b) Directed possibilistic graphs and possibilistic logic. In: *Information, Uncertainty and Fusion*, B. Bouchon-Meunier, R.R. Yager, L. A. Zadeh, Eds: Kluwer Academic, pp. 365-379.
- S. Benferhat, D. Dubois, O. Papini. (1999c) A sequential reversible belief revision method based on polynomials. *Proc 16th Nat Conf. on Artificial Intelligence, (AAAI'99)*, AAAI Press & The MIT Press, 733-738.
- S. Benferhat, D. Dubois, S. Kaci (2001a) Bridging logical, comparative and graphical possibilistic representation frameworks, *Sixth European Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-2001)*, LNAI series, Springer Verlag, Berlin.
- S. Benferhat, D. Dubois, S. Lagrue, H. Prade (2001b) Towards learning default rules by identifying big-stepped probabilities. *Joint 9th IFSA World Congress and 20th NAFIPS International Conference*, Vancouver, Ca.
- C. Boutilier, M. Goldszmidt M. (1995) Revision by conditionals beliefs. In: *Conditionals: From Philosophy to Computer Sciences* (G. Crocco, L. Fariñas del Cerro, A. Herzig, eds.), Oxford University Press, Oxford, UK
- L.J. Cohen (1977) *The Probable and the Provable*. Clarendon Press Oxford.
- A. Darwiche, J. Pearl J. (1997) On the logic of iterated belief revision. *Artificial Intelligence*, 89, 1997, 1-29.
- B. De Finetti (1937) La prévision, ses lois logiques et ses sources subjectives. *Ann. Inst. Poincaré*, 7, 1-68.
- B. De Finetti (1974) *Theory of Probability*, Wiley, New York.
- D. Dubois (1986) Belief structures, possibility theory and decomposable confidence measures on finite sets. *Computers and Artificial Intelligence*, 5(5), 403-416..
- Z. Domotor (1985) Probability kinematics — Conditional and entropy principles. *Synthese*, 63, 74-115.
- D. Dubois, J. Lang, and H. Prade (1994). Possibilistic logic. In D.M. Gabbay et al., editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3, p. 439-513. Oxford University Press.
- D. Dubois, L. Farinas, A. Herzig, H. Prade (1997): Qualitative relevance and independence: A roadmap. *Proc. 15th Int. Joint Conf. on A.I.*, Nagoya, 62-67. Extended version In: *Fuzzy Sets, Logics and Reasoning about Knowledge* (Dubois, D., Prade, H., Klement, E.P., eds.), Kluwer, Academic Publ., 325--359, 1999.
- D. Dubois, H. Fargier, H. Prade (1998a) Comparative uncertainty, belief functions and accepted beliefs. *Proc. of the 14th Conf. on Uncertainty in Artificial Intelligence* (G. Cooper, S. Moral, eds.), Madison, WI, Morgan & Kaufmann, San Francisco, CA, 113-120.
- D. Dubois, S. Moral and H. Prade. (1998b) Belief change rules in ordinal and numerical uncertainty theories *Handbook of Defeasible Reasoning and Uncertainty Management Systems — Vol3*. Kluwer Academic Publ., Dordrecht, The Netherlands, 311-392.

- D. Dubois, H. Fargier, H. Prade (2001) Ordinal and numerical representations of acceptance. In preparation.
- D. Dubois, H. Prade (1991) Epistemic entrenchment and possibilistic logic. *Artificial Intelligence*, 50, 223-239.
- D. Dubois, H. Prade (1994a): Non-standard theories of uncertainty in knowledge representation and reasoning. *The Knowledge Engineering Review*, 9(4), 399-416.
- D. Dubois, H. Prade (1994b) Conditional objects as nonmonotonic consequence relationships, *IEEE Trans. on Systems, Man and Cybernetics*, 24(12), 1724-1740.
- D. Dubois, H. Prade (1995a) Conditional objects, possibility theory and default rules. In *Conditionals: From Philosophy to Computer Science* (G. Crocco, L. Fariñas del Cerro, A. Herzig eds.), Oxford University Press, 311-346.
- D. Dubois, H. Prade (1995b), Numerical representations of acceptance. *Proc. of the 11th Conf. on Uncertainty in Artificial Intelligence* (P. Besnard, S. Hanks, eds.), 149-156.
- Dubois D., Prade H. (1997) A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *Int. J. of Approximate Reasoning*, 17(2/3), 295-324.
- D. Dubois, H. Prade (1998) Possibility theory: qualitative and quantitative aspects. *Handbook on Defeasible Reasoning and Uncertainty Management Systems* — Vol. 1. Kluwer Academic Publ., Dordrecht, The Netherlands, 169-226
- P. Gärdenfors (1988) *Knowledge in Flux*. MIT Press, Cambridge.
- N. Friedman, J. Halpern (1996) Plausibility measures and default reasoning. *Proc of the 13th National Conf. on Artificial Intelligence (AAAI'96)*, Portland, 1297-1304. To appear in *J. Assoc. for Comp. Mach.*
- J. Halpern (1997) Defining relative likelihood in partially-ordered preferential structures. *J. AI Research*, 7, 1-24
- B. O. Koopman (1940) The bases of probability. *Bull. Am. Math. Soc.* 46, 763-774.
- S. Kraus, D. Lehmann, M. Magidor (1990) Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44, 167-207.
- H. E. Kyburg (1988) Knowledge. In: *Uncertainty in Artificial Intelligence* vol 2, J. F. Lemmer and L. N. Kanal eds, Elsevier, 263-272.
- J. Lang (2000) Possibilistic logic: complexity and algorithms. *Handbook on Defeasible Reasoning and Uncertainty Management Systems* — Vol. 5. Kluwer Academic Publ., Dordrecht, The Netherlands, 169-226
- D. Lehmann, M. Magidor (1992) What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1), 1-60.
- D. Lewis (1973) *Counterfactuals*. Basil Blackwell, London.
- J. Paris (1994) *The Uncertain Reasoner's Companion*. Cambridge University Press, Cambridge, UK.
- D. Makinson and P. Gärdenfors (1991) Relations between the logic of theory change and nonmonotonic reasoning. In: *The Logic of Theory Change* (A. Fürmann, M. Morreau, Eds), LNAI 465, Springer Verlag, 185-205
- J. Pearl (1988). *Probabilistic Reasoning Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, CA.
- J. Pearl (1990) System Z: a natural ordering of defaults with tractable applications to default reasoning. *Proc. of the 3rd Conf. on the Theoretical Aspects of Reasoning About Knowledge (TARK'90)*, Morgan and Kaufmann, 121-135.
- D. Poole (1991) The effect of knowledge on belief: conditioning, specificity and the lottery paradox in default reasoning. *Artificial Intelligence*, 49, 281-307
- L.J. Savage (1972) *The Foundations of Statistics*. Dover, New York

- G. L. S. Shackle (1961). *Decision, Order and Time in Human Affairs*, (2nd edition), Cambridge University Press, UK.
- Y. Shoham (1988). *Reasoning About Change — Time and Causation from the Standpoint of Artificial Intelligence*, The MIT Press, Cambridge, MA.
- P. Snow (1999) Diverse confidence levels in a probabilistic semantics for conditional logics. *Artificial Intelligence* 113, 269-279.
- W. Spohn (1988) Ordinal conditional functions: A dynamic theory of epistemic states. In W. Harper, and B. Skyrms, editors, *Causation in Decision, Belief Change and Statistics*, pages 105-134.
- M.A. Williams (1994) Transmutations of knowledge systems. *Proc. of the 4th Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'94)* (J. Doyle, E. Sandewall, P. Torasso, eds.), Bonn, Germany, 1994, Morgan & Kaufmann, San Mateo, CA, 619-629.
- Zadeh L. A. (1978). Fuzzy sets as a basis for a theory of possibility, *Fuzzy Sets and Systems*, **1**, 3-28.