# Reviews

Review of *Coherent Systems*, by Karl Schlechta, volume 2 of Studies in Logic and Practical Reasoning. Amsterdam: Elsevier, 2004, 447 pp., ISBN: 0-444-51789-8.

This book provides an in-depth and advanced treatment of so-called non-monotonic logics, which were devised in order to represent defeasible inference (i.e., that kind of inference in which reasoners draw conclusions tentatively, reserving the right to retract them in the light of further information). Throughout this book the semantic side is stressed. Non-monotonic inference operations are studied within the framework of model theory. The main focus is on so-called preference-based semantics, the study of which was initiated by Shoham [20] among others. The book, which is based on a number of papers by the author, can be described as a research monograph. It brings the reader to the front line of current research in the field of non-monotonic logic and belief revision, by showing both recent achievements and directions of future investigations.

The book is organized into nine chapters. Chapter 1 presents the basic concepts and results that are needed for what follows. Chapter 2 aims at clarifying the conceptual foundations of the different constructions of the author. Chapters 3 to 7 discuss representation results and problems for those constructions, and Chapter 8 is an attempt to integrate them into a single consistent whole. Chapter 9 summarizes the main points that have been made in the book, and ends with some suggestions as to useful ways forward.

I shall here focus on the conceptual side of the book, outlining and commenting on some of its most noticeable aspects. In a review of this nature, one cannot enter into a detailed commentary. The reader must, thus, bear in mind that the material presented below represents only a thin sampling from a rather large volume.

Chapter 2 considers various forms of common-sense reasoning, and argues that they can be based upon more basic semantical notions such as preference, size and distance, which (in some cases) are interdefinable. Much attention is paid to the core forms of common-sense reasoning that are traditionally discussed in the literature. These include: reasoning based on "what is normally the case"; default reasoning; counterfactual reasoning; revision. Here, in barest outline, is how the author analyzes them.

The first kind of reasoning is studied within the tradition of the classic work by Kraus, Lehmann and Magidor − see [10]. Rules determining what is normal are written as conditional assertions having the form $\alpha \mid\sim \beta$. This is read as "$\alpha$ normally entails $\beta$". Two evaluation rules for $\alpha \mid\sim \beta$ are introduced and compared. Both are formulated in terms of a preference relation $\prec$ (to be read as "is more normal than") on the set of possible worlds. One evaluation rule (referred to as the "minimal variant") is the standard one. It says that the relation $\alpha \mid\sim \beta$ holds whenever $\beta$ holds in every world that is minimal under the relation $\prec$ in the set of all $\alpha$-worlds. It is in general assumed that such minimal $\alpha$-worlds exist − this is known as the *Limit Assumption.*

One reason for ruling out the case where such minimal worlds do not exist, has to do with the fact that, in that specific scenario, the minimal account yields the somewhat counterintuitive result that $\alpha$ entails any proposition whatsoever. This observation leads the author to introduce and explore another account, referred to as the "limit version" (p. 17). It is intended to provide meaningful truth-conditions for $|\!\sim$ without assuming the existence of minimal $\alpha$-worlds. The details turn out to be fiddly to state concisely. For present purposes, suffice it to observe that, in the particular case where the universe of the model is totally ordered by $\prec$, the construction puts $\alpha \mid\!\sim \beta$ to be true if and only if there is some $\alpha$-world $w$ such that the material implication $\alpha \rightarrow \beta$ holds at any world $w'$ with $w' \prec w$. Interestingly enough, the limit account is also used to reveal insights into two related areas. One is belief revision theory (p. 258-261) and the other is the semantics of counterfactuals (p. 269-270).

Perhaps the best known formalism for the second type of reasoning cited above is the default logic of Reiter [17]. The account presented in *Coherent Systems* seems to be very different. Default rules are here interpreted in first-order logic by a generalized quantifier. To understand why, it is worth recalling that one of Reiter's initial motivations was "to provide[...] a representation for the 'fuzzy' quantifier 'most' or 'almost all' [...] without appealing to frequency distributions or fuzzy logics" [17, p. 83]. Therefore, a normal default rule like $\frac{\phi(x):\psi(x)}{\psi(x)}$ can be put into the form $\nabla x\phi(x) : \psi(x)$, where $\nabla$ is a restricted quantifier ranging over a "big" subset of the set of $x$ that are $\phi$s. Intuitively, the sentence $\nabla x\phi(x) : \psi(x)$ says that the elements of the latter subset are all $\psi$s. The property of being a "big" subset of another set is a qualitative notion, which can be given a precise characterization by using the notion of filter. A filter over a given set $V$ is any family $\Delta$ of subsets of $V$ such that $V \in \Delta$, $S \in \Delta$ whenever $V \supseteq S \supseteq T \in \Delta$ , and $S \cap T \in \Delta$ whenever $S, T \in \Delta$. Intuitively a filter over $V$ is thus like a family of "important" subsets of $V$ in so far as it contains $V$ and all of the supersets of any of its elements. It goes without saying that the notion of filter can also be used in a propositional setting with a possible-worlds semantics. The resulting framework is close to so-called minimal models for conditional logic as developed by Chellas [5].

The last two forms of non-monotonic reasoning are analyzed in terms of distance. The basic ingredient of the semantics is a metric space, i.e. a set $W$ of possible worlds with an associated distance function (usually called a metric) $d : W \times W \rightarrow Z$ (where $Z$ is a totally ordered set of values, e.g., the set of real numbers). The use of a metric space is quite appropriate here. According to Lewis [12], a counterfactual conditional of the form "If it were the case that $\alpha$, then it would be the case that $\beta$", symbolized as $\alpha \Rightarrow \beta$, is true in the actual world $w$ if and only if $\beta$ is true in every $\alpha$-world that is most similar to the actual world $w$. The identification of the set of most similar worlds can obviously be made *via* an abstract distance function. A similar remark applies to belief revision. In the AGM semantics for belief change (see [6]), we have $\beta \in K \star \alpha$ (sentence $\beta$ is in the result of revising theory $K$ by $\alpha$) if and only if $\beta$ is in every theory that is minimally different (among the theories containing $\alpha$) from the theory $K$. Thus, while the truth-clause for the counterfactual conditional operator considers minimal distance from the world in which the evaluation is taking place, the evaluation rule for the belief revision operator compares distances of various theories from a given theory.

The general question is whether it is possible to give a formal logical account of these various constructions. This issue is the prime focus of Chapters 3 - 7. This group of chapters is the centrepiece of the book, and discusses representation problems for structures incorporating the basic semantical notions discussed above.* It is not my purpose to summarize the various results (both positive and negative) reported in this group of chapters, and I shall

*An enlightening discussion of the difference between representation and completeness theorems can be found in Makinson [14, p. 28-29].

confine myself with a general remark on the methodology used to establish the representation results. In fact, the general strategy in proving these results is articulated around two essential steps. This is best explained using an example. Let us take the case of chapter 3, which discusses various classes of preferential structures. The first step consists in showing that any non-monotonic inference relation $|\!\sim$ satisfying a given set $X$ of rules can be represented by a choice (or selection) function $\mu$: $\mathcal{P}(W) \to \mathcal{P}(W)$ ($\mathcal{P}$= the power set operator) satisfying a corresponding set $Y$ of conditions. Here $\mu$ takes the truth-set of a statement (i.e., the set of possible worlds in which this statement is true) as input, and gives the set of its "preferred" elements as output. The second step involves showing that such a function can in turn be represented by a preference relation $\prec \subseteq W \times W$ satisfying a corresponding set $Z$ of conditions. These two kinds of representation results (which are called "algebraic" and "logical", respectively) are combined to yield a representation theorem for the inference relation $|\!\sim$ in terms of the preference relation $\prec$.

Of special interest is the fact that this two-step method can also be used to provide a similar representation result for related logics, with obvious modifications to take into account the peculiarities of the different frameworks. In particular, Chapters 4 and 6 present a number of representation results for belief revision and updating, respectively. The latter results are much like those for non-monotonic inference operations, except that the intermediary choice function takes an extra argument.† Such an adjustment is required by the fact that the syntactic construct to be represented has now two inputs. This two-step methodology is somewhat usual, and is worth mentioning. It has a number of interesting repercussions.

On the one hand, it leads the author to put a great deal of emphasis on (to put it into his own terms) "coherence properties" and their role in the study of non-monotonic reasoning. These denote in fact the properties of the intermediary selection function alluded to above. It is natural to ask what the term "coherence" means here. The author justifies his terminology by pointing out that in general such properties "describe answers to the following question: Given $\mu(X)$, can we still choose freely $\mu(Y)$, or is there some coherence between $\mu(X)$ and $\mu(Y)$ (given some relation connecting $X$ and $Y$) ?"‡(p. 7).

The reader is invited to compare the discussion with Lindström [13], Lehmann [11] and Rott [18], where a similar emphasis is put on the notion of selection function, although with a different purpose. There it is to show that the theories of non-monotonic reasoning and belief change can fruitfully be linked to the theory of rational choice developed during the last 50 years, mainly by economists. The notion of selection function is commonly used in this area too. Here $\mu$ selects, for each potential "menu" or options $\{x, y, ...\}$ open to the agent, the elements that are "best" in some unspecified sense. A variety of choice-consistency conditions are standardly placed on $\mu$, which resemble those discussed by non-monotonic logicians. For instance, the semantical counterpart of the rule for $|\!\sim$ known as "conditionalization" says that $\mu(Y) \cap X \subseteq \mu(X)$ whenever $X \subseteq Y$. This is Chernoff's basic contraction property discussed by Sen [19] among others. It says that an item that is chosen from a set $Y$ and belongs to a subset $X$ of $Y$ must be chosen from $X$ as well. Moreover, economists have also investigated connections between conditions on selection functions and the existence and properties of relations determining them.

At the same time, it should also be remembered that there are differences. One is that whereas researchers in social science consider preferences between items without assuming that they have any internal structure, logicians work with preferences between items with

---

†A similar device is employed for different purposes in Kanger [9].

‡The formulae of the quotation have been adapted to the notational convention used here.

an internal complexity; typically, these are truth-valuations. This raises new issues for consideration. In particular, it becomes natural to ask what happens if multiple copies of valuation are, or are not, allowed. Small though it may seem, this variation makes a big difference. The issue is discussed in length by Schlechta, particularly in Chapter 3, section 3.8 (but see also, e.g., Moinard and Rolland [15] and Makinson [14, p. 75-80]).

This two-step methodology also brings to light a number of problems, which would be overlooked by adopting a one-step approach. For instance, a careful examination reveals that many representation theorems will not carry over to the infinite case unless the choice function satisfies an additional property, called "definability preservation". Roughly speaking, it says that, taking the truth-set of some theory $T$ as input, $\mu$ always returns the truth-set of some other theory $T'$ as output. Now, it is natural to ask if such a condition can be dropped. This issue is the prime target of Chapter 5. Two representation results are established, one for non-monotonic inference operations, the other for distance based revision, which both hold in the absence of such an assumption. Alongside these positive results there are a number of negative results as well. These confirm, retroactively, that the definability preservation postulate is not as innocuous as it may seem.

As can be imagined, there is much more in this book. I have only been able to hint at the many logical and philosophical riches in this carefully crafted text, which naturally complements other recent books on the same subject. Considerable diversity is exhibited by current approaches to non-monotonic reasoning. *Coherent Systems* argues that there is nevertheless a core common to most of them, which can be located in the notion of preference, size and distance. When studying this book, readers may find it useful to consult the recently published David Makinson's *Bridges from Classical to Nonmonotonic Logic* [14]. There the author explores an alternative way of understanding non-monotonic logics, obtained through an in-depth analysis of the process of getting more conclusions out of a premisses set. There is, thus, a shift of emphasis from possible-worlds semantics to the notion of consequence relation. Readers with a background in Computer Science should be warned, however, that the issue of implementing non-monotonic reasoning is put to one side in these two research monographs. This is to be contrasted with, for example, the book by Antoniou [1], where default logic is put to the foreground, and issues related to its implementation are discussed as well.

I end this review with a remark on terminology, which at the same time brings to the surface a deep conceptual issue. The notions of preference, size and distance are undoubtedly useful as a tool for theoretical investigations, and help clarify how non-monotonic concepts are related to each other. It should be noted, however, that in *Coherent Systems* the author goes further. He describes himself as attempting to "reduc[e] a certain number of reasoning mechanisms to a small number of basic semantical concepts" (p. 13). The term "reduction" occurs several times, and can be misleading, for at least two reasons.

First of all, it may suggest that semantics for non-monotonic logics are vitiated by some kind of circularity. The objection is directly comparable to the one raised in the late 1970s against the framework provided by possible-worlds semantics. The truth-conditions for the modal operators do exhibit circularity, and so do the characterizations given to the non-monotonic notions. For instance, in a preference-based setting, the truth of 'if $\phi$ then *normally* $\psi$' requires the truth of $\psi$ in all of the most *normal* among the worlds that satisfy $\phi$. One issue here is whether it is appropriate to describe the circle as vicious. There is a widespread agreement among modal logicians that the answer is no. To be more specific, they agree that the circle would be vicious, only if the truth-conditions for the modal operators were indeed an attempt to reduce the modal notions to some other concepts – but they are not, so that the criticism is misplaced (*cf.* Hintikka [7, p. 367], Pörn [16, p. 13] and Jones [8, p. 69-71]).

Second, if we look more closely at the interplay between axiom schemata and conditions on frames, then we can see that the passage from the syntactic level to the semantic one is sometimes straightforward. The analysis of default reasoning in terms of size provides a good example. In general the condition that needs to be placed on a filter in order to secure the validity of the corresponding axiom schema is "obvious" (p. 367). Similar accusations have been levelled against neighbourhood semantics (or minimal models) for non-normal modal logics. These have often been criticized on the score that the properties of the neighbourhood function parallel fairly closely the corresponding axiom schemata, and thus are not particularly enlightening. Obviously, this statement needs to be tempered here. As mentioned above, in many cases, the use of a choice function (which is the non-monotonic counterpart of a neighbourhood function) is just one step in the analysis. There is, next, the task of identifying the property of the preference relation (or the metric) onto which the axiom schema must, ultimately, be mapped. This task is not always straightforward.

In view of these remarks the reader may feel tempted to ask what the non-monotonic logician's task is, if it is not to effect a reduction of the kind just described. It is tempting to say that the task is just to provide a formal framework which helps facilitate the systematic investigation of non-monotonic entailment relations between sentences. For a logician, this response (which might be termed "minimalist") is obvious. For a newcomer, it is not. Let me briefly illustrate the need for a semantic foundation of the kind found in *Coherent Systems*, by making an initial incursion into another territory – that of the so-called Unified Modeling Language (UML).[§] This is a graphical language for describing and reasoning about ontologies, which is becoming increasingly popular within the Software Engineering community. A detailed presentation of the UML goes beyond the scope of this review. For present purposes, suffice it to observe that, as such, the UML is just a graphical notation aiming at facilitating communication among researchers coming from different horizons. It is here interesting to remark that "if-then" relationships have their immediate counterparts in the UML. Indeed, a fundamental building block of so-called UML class diagrams is the "is-a" relationship between classes – graphically it is rendered as a solid directed line with a large open arrowhead, pointing to the parent class.

The use of such a graphical notation is certainly useful for communicational and pedagogical purposes. However, it carries an obvious danger if it is not supported by evidence from formalization. Often the antecedent of a given rule is itself the consequent of another rule. Once the graphical notation is used, it leads almost irresistibly to the feeling that the two rules can be chained together, by simply drawing a new arrow starting from the first class and leading directly to the third one. This is known as the principle of transitivity. Roughly it says that, if $\alpha$ implies $\beta$ and $\beta$ implies $\gamma$, then $\alpha$ implies $\gamma$. Transitivity is *prima facie* attractive, but examples from ordinary conversation convince us that it is to be avoided. The following is such an example: students are usually adult; adults are usually employed; but students are usually unemployed. It is worth considering what the abstract system based on size described above has to say about transitivity. It is easy to see that the critical implication, "If $\alpha$ then $\gamma$", is not endorsed by the formalism. This is illustrated by figure 1. "If $\alpha$ then $\beta$" holds since the $\alpha$-area (circle) is mostly included in the $\beta$-area (oval); "If $\beta$ then $\gamma$" also holds, since most of the $\beta$-area is contained in the $\gamma$-area (rectangle); but "If $\alpha$

[§] See, e.g., Booch et al. [4].

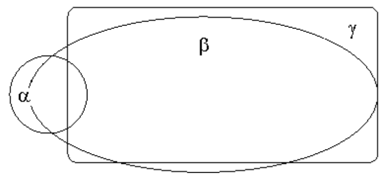


FIG. 1. Plain transitivity and filters

then $\gamma$" fails, since only a small part of the $\alpha$-area is contained within the $\gamma$-area. Of course, some restricted forms of transitivity may be acceptable. Identifying them is, precisely, one of the purposes of non-monotonic logics.¶

## References


[1] G. Antoniou. *Nonmonotonic Reasoning*. MIT Press, Cambridge Mass, 1997.

[2] F. Baader and B. Hollunder. Embedding defaults into terminological representation systems. *J. Automated Reasoning*, 14:149–180, 1995.

[3] D. Berardi, D. Calvanese, and G. De Giacomo. Reasoning on UML class diagrams using Description Logic based systems. In *Proc. of the KI'2001 Workshop on Applications of Description Logics*. CEUR Electronic Workshop Proceedings, 2001.

[4] G. Booch, J. Rumbaugh, and I. Jacobson. *The Unified Modeling Language User Guide*. Addison Wesley, 1999.

[5] B. F. Chellas. Basic conditional logic. *Journal of Philosophical Logic*, 4:133–153, 1975.

[6] P. Gärdenfors. *Knowledge in Flux*. MIT Press, 1988.

[7] J. Hintikka. Impossible possible worlds vindicated. In E. Saarinen, editor, *Game-Theoretical Semantics*, volume 5 of *Studies in Linguistics and Philosophy*, pages 367–379, Dordrecht, Holland, 1978. D. Reidel.

[8] A.J.I. Jones. *Communication and Meaning: An Essay in Applied Modal Logic*, volume 168 of *Synthese Library*. D. Reidel, Dordrecht, Holland, 1983.

[9] S. Kanger. Choice based on preference. In G. Holmström-Hintikka, S. Lindström, and R. Sliwinski, editors, *Collected Papers of Stig Kanger with Essays on his Life and Work*, volume 1, pages 214–230, Dordrecht, Holland, 2001. Kluwer Academic Publishers.

[10] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

[11] D. Lehmann. Nonmonotonic logics and semantics. *Journal of Logic and Computation*, 11:229–256, 2001.

[12] D. Lewis. *Counterfactuals*. Harvard University Press, Cambridge, 1973.

[13] S. Lindström. *A Semantic Approach to Nonmonotonic Reasoning: Inference Operations and Choice*. University of Uppsala, Uppsala Prints and Preprints in Philosophy, 1994.

[14] D. Makinson. *Bridges from Classical to Nonmonotonic Logic*, volume 5 of *Texts in Computing*. King's College Publications, London, 2005.

[15] Y. Moinard and R. Rolland. Characterizations of preferential entailments. *L. J. of the IGPL*, 10:245–72, 2002.

[16] I. Pörn. *Action Theory and Social Science: Some Formal Models*. D. Reidel, Dordrecht, Holland, 1977.

[17] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[18] H. Rott. *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*, volume 43 of *Oxford Logic Guides*. Clarendon Press, Oxford, 2001.


¶Over the past few years, a great deal of work has been done to formalize UML class diagrams. One might mention the attempt made by Berardi et al. [3] to define a mapping between UML class diagrams and so-called description logics. The semantics of "is-a" is defined in an obvious way, i.e. $C_1 \sqsubseteq C_2 := \forall x. C_1(x) \rightarrow C_2(x)$, where $\rightarrow$ denotes material implication. It is natural to ask if such a rendering is always appropriate. Some of the difficulties encountered in combining description logics with Reiter's default logic are discussed, for example, in Baader and Hollunder [2].

[19] A. Sen. Internal consistency of choice. *Econometrica*, 61:495–521, 1993.
[20] Y. Shoham. *Reasoning about Change*. MIT Press, Cambridge, USA, 1988.

Xavier Parent
*King's College London*

doi:10.1093/jigpal/jzi026