Two Paradigms of Nonmonotonic Reasoning

Alexander Bochman
Computer Science Department,
Holon Academic Institute of Technology, Israel
e-mail: bochmana@hait.ac.il

Abstract

We provide a conceptual description of the field of nonmonotonic reasoning as comprising two essentially different theories, preferential and explanatory nonmonotonic reasoning. The relationship between the two constitutes the main theoretical problem of nonmonotonic reasoning, and its solution should hopefully provide an impetus for the future development of the field.

Introduction

Studies in nonmonotonic reasoning have given rise to two different approaches that we will call *preferential* and *explanatory* nonmonotonic reasoning. The explanatory approach includes default and modal nonmonotonic logics, as well as logic programming with negation as failure (see, e.g., (Marek & Truszczyński 1993; Bochman 2005)). The preferential approach was initiated in (Gabbay 1985) on the logical side, and in (Shoham 1988) on the AI side. It encompasses nonmonotonic inference relations and a general theory of belief change (see (Bochman 2001; Schlechta 2004)).

Differences between the two approaches can be found on a number of levels. To begin with, there are two senses in which a logical formalism can be nonmonotonic. First, it can be *locally* nonmonotonic in that its rules do not admit addition of new premises, that is, they do not satisfy Strengthening the Antecedent. Second, it may be *globally* nonmonotonic in that adding further rules to a system may invalidate previous conclusions. These two kinds of nonmonotonicity are largely independent. Thus, preferential inference relations (Kraus, Lehmann, & Magidor 1990) are locally nonmonotonic (*Birds fly* does not imply *Penguins fly*). However, they are globally monotonic, since addition of new conditionals does not invalidate previous conclusions. On the other hand, default logic (Reiter 1980) exemplifies the combination of local monotonicity with global nonmonotonicity. Any default theory can be extended with default rules obtained from existing ones by strengthening their premises; they will not change the set of extensions. On the other hand, adding arbitrary new rules to the default theory may result in creating new extensions, so nonmonotonic conclusions made earlier will not, in general, be preserved.

Despite this difference, one of the main incentives behind the preferential approach has been the hope that default logic and other formalisms can be subsumed by the preferential approach under some ingenious notion of preference. Unfortunately, subsequent studies have raised grave doubts about this hope. Thus, the nonmonotonic semantics of default logic has turned out to violate even the most basic postulates of cumulative inference (see (Makinson 1989)).

In a hindsight, this outcome should have been expected, since the selection of intended models in the explanatory approach is not preferential in a usual sense; rather, such models are determined as models satisfying certain justification conditions with respect to the rules. On a most abstract level, they are expressible as fixed points of an operator which is not even monotonic. Accordingly, the relevant preference appears to be a trivial, zero-one preference that differentiates only right models from bad ones.

Both preferential and explanatory nonmonotonic reasoning can be seen as theories of a reasoned use of assumptions. Now, preferential reasoning treats such assumptions as *defaults*, namely as normality assumptions we can use whenever there is no evidence to the contrary. This presumptive reading has a semantic counterpart in the notion of *normality*: defaults should hold for normal circumstances, and the theory tells us to assume that the world is as normal as is compatible with the known facts. This naturally creates a preferential setting, in which the normality of models is measured by the set of defaults they support (see below).

Explanatory reasoning assigns, however, a different role to assumptions. Following (Poole 1989), we can call such assumptions *conjectures*. Conjectures are assumptions that we make in order to explain observations. The supposition of normality is not essential here; we make conjectures only if there is evidence that requires them for explanation, in contrast to defaults that can be freely assumed whenever possible. As was argued by Poole, the distinction between defaults and conjectures is closely related to the distinction between prediction and explanation: while we use defaults in order to predict facts that are yet unknown, conjectures are invoked when we have to explain known facts. Untreated syphilis can explain paresis, though the syphilis assumption does not have predictive force of deriving paresis.

Unfortunately, the above distinction has been obscured, because all the main formalisms of nonmonotonic reasoning, including default logic, have claimed their rights and responsibility on representation of normality defaults. Thus, Reiter has suggested in (Reiter 1980) that we can identify such defaults with default rules of the form $A{:}B/B$, appropriately called normal default rules.

On our view, the preferential approach provides a more adequate analysis of normality defaults. The examples in the literature that reveal a discrepancy between Reiter's normal defaults $A{:}B/B$ and corresponding preferential conditionals $A \mid\!\sim B$ point out in favor of the latter. So the criticism raised against default logic in the preferential camp seems justified so far as we are talking about which notion better reflects our understanding of normality. Still, there is nothing wrong or nonintuitive about having both $A : B/B$ and $A : \neg B/\neg B$ in a default theory, though it is certainly counterintuitive to treat such rules simultaneously as normality defaults. In the setting of default logic, such rules say simply that, when A holds, both B and $\neg B$ are equally admissible conjectures. This indicates, however, that default logic has a subject of its own that should not be extrapolated to the entire field of nonmonotonic reasoning.

Nonmonotonic Reasoning and Logic

Nonmonotonic reasoning changes the ways logic is used. To begin with, there are many reasons to believe that nonmonotonic reasoning cannot be expressed in the form of a logical inference in some ingenious 'nonmonotonic logic'. Logical systems provide only a more or less tight *framework* in which nonmonotonic reasoning can be represented and performed. This pertains not only to monotonic logical systems, but also to so-called nonmonotonic logics.

Many nonmonotonic formalisms have two components. The first is a logical framework, e.g., classical logic in circumscription, or a modal logic in modal nonmonotonic logics. The second, nonmonotonic, component determines which of the possible models should be considered as intended ones. In ordinary logical systems, the semantics determines the logical consequences of a theory, but also, and most importantly, it provides an interpretation for the syntax itself. Namely, it provides propositions and rules of a syntactic formalism with *meaning*, and its theories with informational content. By its very design, however, the nonmonotonic semantics is defined as a certain *subset* of the set of possible models, and consequently it does not determine, in turn, an interpretation of the syntax. As a result, two radically different theories may have the same nonmonotonic semantics. Furthermore, such a difference cannot be viewed as apparent, since it may well be that by adding further rules or facts to both these theories, we obtain new theories that already have different nonmonotonic models.

The above situation is quite similar to the distinction between meaning (intension) and extension of logical names, a distinction that is fundamental for modern logic. Nonmonotonic semantics provides, in a sense, the extensional content of a theory in a particular context of its use. In order to determine the meaning, or informational content, of a theory, we have to consider the underlying logic. This requires a *clear separation of logical and nonmonotonic aspects of nonmonotonic reasoning*. The separation suggests the following understanding of nonmonotonic reasoning:

NMR = Logic + Nonmonotonic Semantics

Logic and its associated logical semantics are responsible for providing the meaning of the rules of the formalism, while the nonmonotonic semantics provides us with nonmonotonic consequences of a theory. The benefits of this separation will be described below. As we will see, however, both approaches to nonmonotonic reasoning will have to be elaborated in order to comply with this schema.

Preferential Nonmonotonic Reasoning

The main problem nonmonotonic reasoning deals with is that default assumptions are often incompatible with one another, or with known facts. In such cases of conflict we must have a reasoned choice. The preferential approach follows here the slogan *"Choice presupposes preference"*. Namely, it asserts that the choice of assumptions should be made by establishing preferences among admissible options. This makes preferential approach a special case of a general methodology that is at least as old as the decision theory.

Epistemic States

It has been suggested in (Bochman 2001) that a general representation of preferential nonmonotonic reasoning can be given in terms of epistemic states, defined below.

Definition 1. *An* epistemic state *is a triple* $(\mathcal{S}, l, \prec)$*, where* $\mathcal{S}$ *is a set of* (admissible) belief states*,* $\prec$ *a preference relation on* $\mathcal{S}$*, while* l *is a labeling function assigning a deductively closed* belief set *to every state from* $\mathcal{S}$*.*

On the intended interpretation, belief states are generated by allowable combinations of default assumptions. The preference relation is usually based on the supposition that a belief state generated by a larger set of defaults is preferred to a state generated by a smaller set. Additional preferences arise from priorities among defaults themselves (see below).

Epistemic states are epistemic because they say nothing directly about what is actually true, but only what is believed (or presumed) to hold. This makes them relatively stable entities; change in facts does not necessary lead to change in epistemic states. The actual assumptions made in particular situations are obtained by choosing preferred belief states that are consistent with the facts.

Prioritization A formal description of epistemic states generated by default bases provides us with characteristic properties of epistemic states arising in particular situations. An epistemic state is *base-generated* by a set of propositions Δ with respect to a classical consequence relation Th if

- $\mathcal{S}$ is the set $\mathcal{P}(\Delta)$ of subsets of Δ;
- $l(\Gamma) = \mathrm{Th}(\Gamma)$, for any $\Gamma \subseteq \Delta$;
- $\prec$ is monotonic on $\mathcal{P}(\Delta)$: if $\Gamma \subset \Phi$, then $\Gamma \prec \Phi$.

The preference order on belief states is usually derived from priorities among individual defaults. This task is reducible to a problem of combining a set of preference relations into a single 'consensus' preference order. Let Δ be ordered by a *priority relation* $\lhd$ which will be assumed to be a strict partial order: $\alpha \lhd \beta$ will mean that α *is prior to* β.

Recall that defaults are presumed to hold insofar as it is consistent to do so. Hence any default δ determines a primary preference relation $\preccurlyeq_\delta$ on $\mathcal{P}(\Delta)$ by which admissible belief sets containing the default are preferred to belief sets that do not contain it: $\Gamma \preccurlyeq_\delta \Phi \;\equiv\;$ if $\delta \in \Gamma$ then $\delta \in \Phi$.

Now the problem of finding a global preference order amounts to constructing an operator that maps the set of preference relations $\{\preccurlyeq_\delta |\ \delta \in \Delta\}$ to a single preference relation on $\mathcal{P}(\Delta)$. By the results from (Andreka, Ryan, & Schobbens 2002), if this operator is required to satisfy the Arrow's conditions, it should be defined via the following *lexicographic rule*:

$$\Gamma \prec \Phi \;\equiv\; \Gamma \neq \Phi \wedge (\forall \alpha \in \Gamma \backslash \Phi)(\exists \beta \in \Phi \backslash \Gamma)(\beta \lhd \alpha)$$

(Lifschitz 1985) was the first to use this construction in prioritized circumscription, while (Geffner 1992) employed it for defining preference relations on sets of defaults.

Despite its virtues, the above definition of preference still does not cover some applications, because it is based on absolute, unconditional priorities. Thus, inheritance hierarchies (see below) turn out to be representable in terms of priorities that are conditional upon presence of other defaults.

A conditional priority relation on Δ is a ternary relation $\alpha \lhd_\Gamma \beta$ (where $\Gamma \subseteq \Delta$) saying that α *is prior to* β *with respect to* Γ. Conditional priorities can also be used for defining preference relations as follows:

$$\Gamma \prec \Phi \;\equiv\; \Gamma \neq \Phi \;\&\; (\forall \alpha \in \Gamma \backslash \Phi)(\exists \beta \in \Phi \backslash \Gamma)(\beta \lhd_\Gamma \alpha)$$

It turns out that, under some natural additional conditions, this generalized preference relation has the same nice features as in the absolute case.

Nonmonotonic inference and its kinds

All nonmonotonic inference relations presuppose a two-step selection procedure: for an evidence A, we take the set $\langle A \rangle$ of belief states that are consistent with A and choose preferred elements in this set. A *skeptical inference* (or prediction) is obtained when we infer only what is supported by each of the preferred states. Namely, B is a skeptical conclusion from A if each preferred belief set that is consistent with A, taken together with A itself, implies B.

Definition 2. *B is a* skeptical consequence *of A in an epistemic state if $A \to B$ holds in all preferred states from $\langle A \rangle$.*

A set of conditionals that are valid in an epistemic state $\mathbb{E}$ forms a *skeptical inference relation* determined by $\mathbb{E}$.

A *credulous inference* (or explanation) is obtained by inferring conclusions supported by at least one preferred belief state consistent with the facts. In other words, B is a credulous conclusion from A if at least one preferred admissible belief state in $\langle A \rangle$, taken together with A, implies B.

Definition 3. *B is a* credulous consequence *of A in an epistemic state if $A \to B$ is supported by at least one preferred belief state in $\langle A \rangle$.*

The set of conditionals that are credulously valid in an epistemic state $\mathbb{E}$ forms a *credulous inference relation* determined by $\mathbb{E}$. The above definition generalizes the notion of explanation from (Poole 1988). Credulous inference is only one, though important, instance of a broad range of non-skeptical inference relations (see (Bochman 2003)).

Axiomatic characterization

A common ground for axiomatization of both skeptical and credulous inference is provided by a logic of conditionals suggested in (van Benthem 1984). The *basic inference relation* is determined by the following postulates:

Reflexivity $\quad A \mathrel{|\!\sim} A$

Left Equivalence $\quad$ If $\vDash A \leftrightarrow B$ and $A \mathrel{|\!\sim} C$, then $B \mathrel{|\!\sim} C$

Right Weakening $\quad$ If $A \mathrel{|\!\sim} B$ and $B \vDash C$, then $A \mathrel{|\!\sim} C$

Antecedence $\quad$ If $A \mathrel{|\!\sim} B$, then $A \mathrel{|\!\sim} A \wedge B$

Deduction $\quad$ If $A \wedge B \mathrel{|\!\sim} C$, then $A \mathrel{|\!\sim} B \to C$

Cautious Monotony $\quad$ If $A \mathrel{|\!\sim} B \wedge C$, then $A \wedge B \mathrel{|\!\sim} C$

Since all the above postulates involve at most one conditional premise, the basic entailment boils down to derivability among single conditionals. The following theorem provides a direct description of this derivability relation.

Theorem 1. *$A \mathrel{|\!\sim} B \Vdash_{\mathcal{B}} C \mathrel{|\!\sim} D$ if and only if either $C \vDash D$, or $A \to B \vDash C \to D$ and $C \to \neg D \vDash A \to \neg B$.*

Basic inference does not allow to combine different conditionals, but it is complete for derivability among individual conditionals; it captures exactly the one-premise derivability of both skeptical and credulous inference relations.

The following postulate was singled out in (Gabbay 1985) as a characteristic feature of sceptical inference:

(And) $\quad$ If $A \mathrel{|\!\sim} B$ and $A \mathrel{|\!\sim} C$, then $A \mathrel{|\!\sim} B \wedge C$.

Indeed, in the framework of basic inference, And is all we need for capturing precisely the preferential inference relations from (Kraus, Lehmann, & Magidor 1990). Such inference relations provide a complete axiomatization of skeptical inference with respect to epistemic states. An important special case of preferential inference, rational inference relations, are determined by linearly ordered epistemic states; they are obtained by adding further

(Rat. Monotony) If $A \mathrel{|\!\sim} B$ and $A \not\mathrel{|\!\sim} \neg C$, then $A \wedge C \mathrel{|\!\sim} B$.

In contrast, credulous inference relations do not satisfy And. Still, they are axiomatized as basic inference relations satisfying Rational Monotony.

Rational Monotony is not a 'Horn' rule, so it does not allow us to derive new conditionals from given ones. In fact, credulous inference relations do not derive much more conditionals than what can be derived already by basic inference (see (Bochman 2001)). This indicates that there should be no hope to capture credulous nonmonotonic reasoning by derivability in some nonmonotonic logic. Though less evident, the same holds for skeptical inference. Both these kinds of inference need to be augmented with an appropriate globally nonmonotonic semantics that would provide a basis for the associated systems of defeasible entailment, as described in the next section.

Defeasible entailment

Practically all problems of reasoning with default conditionals are reducible to the question what conditionals can be derived from a conditional default base. The latter problem constitutes therefore the main task of a theory of default conditionals (cf. (Lehmann & Magidor 1992)).

For a skeptical reasoning, a most plausible understanding of default conditionals is obtained by viewing them as skeptical inference rules in epistemic states. Accordingly, preferential inference can be considered as a *logic* behind skeptical nonmonotonic reasoning. This does not mean, however, that nonmonotonic reasoning about default conditionals is reducible to preferential derivability. Preferential inference is severely sub-classical and does not allow us, for example, to infer "Red birds fly" from "Birds fly". In fact, this is precisely the reason why such inference relations have been called nonmonotonic. Clearly, there are good reasons for not accepting such a derivation as a *logical* rule; otherwise "Birds fly" would imply also "Birds with broken wings fly" and even "Penguins fly". Still, we can accept "Red birds fly" as a *nonmonotonic* (or defeasible) conclusion in the absence of information against it. By doing this, we would just follow the general strategy of nonmonotonic reasoning of making reasonable assumptions on the basis of available information. Thus, the logical core of skeptical inference, preferential inference relations, should be augmented with a mechanism of making nonmonotonic conclusions. Speaking generally, we would like to keep reasoning classically about default conditionals insofar as this does not conflict with the default base and evidence. In contrast to preferential inference, which is locally nonmonotonic, this reasoning will be *globally* nonmonotonic, since addition of new conditionals can block some of the conclusions made earlier.

On the semantic side, default conditionals are constraints on epistemic states, but usually there is a large number of epistemic states that satisfy a given set of conditionals, so we have an opportunity to choose among them. Our guiding principle in this choice can still be that the epistemic states should be as normal as is permitted by the constraints. By choosing particular such states, we will adopt conditionals that will not be derivable by preferential inference alone.

The above considerations lead to a seemingly inevitable conclusion that default conditionals possess a clear logical meaning and associated logical semantics based on epistemic states (or possible worlds models), but they still need a globally nonmonotonic semantics that would provide an interpretation for the associated defeasible entailment.

Actually, the literature on nonmonotonic reasoning is abundant with theories of defeasible entailment. Initial formal systems, rational closure (Lehmann 1989) and Pearl's Z (Pearl 1990), have turned out to be insufficient for representing defeasible entailment, so they have been refined to systems such as lexicographic inference (Lehmann 1995), and similar modifications of Z (e.g., (Tan & Pearl 1995)). Unfortunately, the refined systems have encountered an opposite problem, namely, together with some desirable properties, they produced unwanted conclusions. All these systems have been based on rational inference. A more general approach based on preferential inference has been suggested in (Geffner 1992). Finally, a more syntactic approach has been pursued in inheritance hierarchies (see (Horty 1994)). Though inheritance reasoning deals with a quite restricted class of conditionals constructed from literals, it has achieved a remarkably close correspondence with intuition.

Despite the diversity, most of the systems of defeasible entailment presuppose that classical implications corresponding to default conditionals should serve as defaults in the nonmonotonic reasoning sanctioned by a default base. This idea can be made precise by requiring that the epistemic states for a default base $\mathfrak{B}$ should be base-generated by the corresponding set $\vec{\mathfrak{B}}$ of material implications (see above). Already this constraint on intended models allows us to derive "Red birds fly" from "Birds fly" for default bases that do not contain conflicting information about redness. It also sanctions defeasible entailment across exception classes, unlike Z and rational closure that cannot make such a derivation. Still, the constraint is insufficient for some important reasoning patterns. What is lacking here is a principled way of constructing a preference order on default sets. As for now, Geffner's conditional entailment and inheritance reasoning constitute two most plausible solutions.

Geffner's theory provides a very plausible interpretation of defeasible entailment. Still, it does not capture inheritance reasoning. The main difference between the two is that inheritance hierarchies determine priorities in a context-dependent way, namely in presence of other defaults that provide a (preemption) link between two defaults. Indeed, it has been shown in (Bochman 2001) that inheritance reasoning is representable by epistemic states that are base-generated by default conditionals ordered by certain conditional priority orders (see above). Still, the corresponding construction could hardly be called simple or natural.

A more natural representation of inheritance reasoning has been given in (Dung & Son 2001) as an instantiation of an argumentation theory that belongs already to explanatory nonmonotonic formalisms. Furthermore, Geffner himself has shown in (Geffner 1992) that conditional entailment still does not capture some important derivations, and it should be augmented with an explicit representation of causal reasoning. In fact, the causal generalization suggested by Geffner in the last chapters of his book has served as one of the inspirations for a causal theory of reasoning about actions and change (see (Turner 1999)). This theory will be described later as an essential part of the explanatory approach to nonmonotonic reasoning.

Finally, a most glaring omission of the above picture of defeasible entailment is that it does not cover *credulous*, or explanatory, nonmonotonic reasoning. Furthermore, for now it is even unclear whether the above approach in terms of epistemic states is capable of representing such a reasoning, though the representation of inheritance reasoning in this framework suggests that it might. Anyway, explanatory reasoning is a well-established theory in its own right, so our next task will consist in singling out its basic principles.

Explanatory Nonmonotonic Reasoning

Explanation is the basic ingredient of explanatory nonmonotonic reasoning. Propositions may be not only true or false in a model, but some of them are explainable (justified) by other accepted facts and rules. In the epistemic setting, some of the propositions are *derivable* from other by rules that are admissible in the situation. In the objective setting, some facts are *caused* by other facts and causal rules acting in the domain. Furthermore, explanatory nonmonotonic reasoning is based on very strong principles of *Explanation Closure* or *Causal Completeness* (Reiter 2001), according to which any fact holding in a model should be explained, or caused, by the rules of the domain. Incidentally, it is these principles that make explanatory reasoning nonmonotonic.

By the above description, *abduction* is an integral part of explanatory nonmonotonic reasoning. Abducibles correspond not to normality defaults, but to conjectures representing base causes or facts that do not require explanation; we assume the latter only for explaining evidence.

Explanatory formalisms often adopt simplifying assumptions that exempt, in effect, certain propositions from the burden of explanation. *Closed World Assumption* (Reiter 1978) is the most important assumption of this kind. According to it, negative assertions do not require explanation. It is important to note that the *minimization principle* is actually a result of combining Explanation Closure with the Closed World Assumption. Consequently, the minimization principle need not be viewed as a principle of scaled preference of negative information; rather, it is a by-product of the stipulation that negated propositions can be accepted without any further explanation, while positive assertions always require explanation. This understanding explains why McCarthy's circumscription, that is based on the principle of minimization, is subsumed also by explanatory formalisms.

The above principles form an ultimate basis for all formal systems of explanatory nonmonotonic reasoning. They presuppose, however, a richer picture of what is in the world than what is usually captured in logical models. The world is not a mere assemblage of unrelated facts, it has a structure that forms a basis for our explanatory and causal claims. It is this structure that makes the world intelligible and controllable. By this picture, explanatory and causal relations form an integral part of understanding of and acting in the world. Such relations should form an integral part of knowledge representation, at least in Artificial Intelligence.

Explanatory nonmonotonic reasoning and logic

Default and autoepistemic logics, and semantics of logic programming, are usually described in a shortcut way in accordance with the following identity:

Nonmonotonic Logic = Syntax + Nonmonotonic Semantics.

The very name 'Nonmonotonic *Logic*' conveys here similarity with ordinary logical systems, for which the equality is appropriate. The analogy is clear, but unfortunately misleading. The nonmonotonic semantics does not determine the meaning of the propositions and rules of the formalism, so the above 'shortcut' definition leaves us without an exact or even clear meaning of the source syntax. Default rules do not bear on their heads information about when and how they can be applied. This is the main reason why the knowledge representation in such systems is essentially an art based on accumulated experience.

In order to determine the logical meaning, or informational content, of a nonmonotonic theory, we should consider its underlying logic. Fortunately, such a logic can often be restored from the nonmonotonic semantics. Given a syntactic formalism $\mathbb{F}$ with a nonmonotonic semantics $\mathbb{S}$, the syntactic formalism determines the basic informational units that we will call theories, for which the semantics provides a nonmonotonic interpretation (a set of models). Let $\mathbb{S}(\Delta)$ denote the nonmonotonic semantics of a theory Δ. Then theories Γ and Δ can be called *(nonmonotonically) equivalent* if $\mathbb{S}(\Delta) = \mathbb{S}(\Gamma)$. This equivalence does not determine, however, the logical meaning of theories; note that due to the nonmonotonicity of $\mathbb{S}$, we may also have $\mathbb{S}(\Delta \cup \Phi) \neq \mathbb{S}(\Gamma \cup \Phi)$, for some theory Φ. A standard definition of meaning in logic says, however, that two notions have the same meaning if they determine the same extension in all contexts. In our case, a context can be seen as a larger theory including a given one, which leads us to the following

Definition 4. *Theories* Γ *and* Δ *are* strongly $\mathbb{S}$-equivalent, *if* $\mathbb{S}(\Delta \cup \Phi) = \mathbb{S}(\Gamma \cup \Phi)$, *for any theory* Φ.

This notion of strong equivalence has actually been suggested in logic programming (see (Lifschitz, Pearce, & Valverde 2001)), but it has general significance. It is already a logical notion, since strongly equivalent theories are interchangeable in any larger theory without changing the associated nonmonotonic semantics. This suggests that there may exist a logic $\mathbb{L}$ formulated in the syntax $\mathbb{F}$ such that theories are strongly equivalent if and only if they are logically equivalent in $\mathbb{L}$. In this case, the logic $\mathbb{L}$ can be viewed as the underlying logic of the formalism that will determine the logical meaning of theories and, in particular, of the rules and propositions of the syntactic framework $\mathbb{F}$.

The attention to the underlying logics behind nonmonotonic reasoning is rewarded with a better understanding of the range of such logics that are appropriate for explanatory nonmonotonic reasoning. It reveals, in particular, that the traditional map of such a reasoning is patently incomplete, and should be completed with a number of important formalisms, creating a continuous range from logic programming to modal nonmonotonic logics.

A distinctive feature of default and modal nonmonotonic logics is that they are inherently *epistemic* formalisms. Namely, they are essentially based on beliefs and knowledge, so the semantic models represent possibly incomplete sets of beliefs, while their rules allow us to make inferences based on absence of belief, or consistency. Due to its epistemic character, default logic is a logically weak formalism that does not support many classical inferences (such as reasoning by cases). It also has other well-known shortcomings, and numerous variants of default logic have been suggested in attempts to make it more in accord with intuitions. However, relatively modest success of these attempts has shown that it is impossible to radically improve default logic without abandoning its epistemic interpretation.

On the other side of the map, logic programs can be embedded into the above epistemic formalisms, *though not vice versa*. This means that logic programming is a more specific nonmonotonic formalism with a richer logic. Furthermore, in between these extreme cases, there is a room for intermediate systems, in particular, for causal reasoning.

Causal reasoning is now a dominant approach for solving the frame problem in representing actions and change (see (Giunchiglia *et al.* 2004)). It employs the already mentioned distinction between facts that hold in a situation versus facts that are caused by other facts and the rules. All facts that hold in a situation should be either caused by other occurrent facts, or else preserve their truth-values in time (by the *inertia assumption*). Causal reasoning constitutes an important turning point in the development of explanatory nonmonotonic reasoning, since from its very beginning it was designed as a formalism that should provide an objective description of factual and causal information about action domains. It has shown that an epistemic view of explanatory nonmonotonic reasoning is not the only possibility.

Biconsequence Relations

Biconsequence relation is a consequence relation for reasoning with respect to a pair of contexts. Taken in an abstract setting, the two contexts will be termed, respectively, the contexts of *truth* and *falsity*. In the truth context propositions are evaluated as being true or non-true, while in the falsity context they can be false or non-false. Then a bi-context reasoning can be interpreted as a reasoning with possibly inconsistent and incomplete information, and also as a four-valued reasoning (see (Belnap 1977)).

On the interpretation suitable for nonmonotonic reasoning, the truth context is the main (objective) one, while the falsity context provides assumptions, or explanations, that justify inferences in the main context.

A *bisequent* is an inference rule of the form $a : b \Vdash c : d$, where a, b, c, d are sets of propositions. According to the abstract, four-valued interpretation, it says

> *'If all propositions from a are true and all propositions from b are false, then either one of the propositions from c is true or one of the propositions from d is false'.*

According to the *explanatory interpretation*, it says

> *'If no proposition from b is assumed, and all propositions from d are assumed, then all propositions from a hold only if one of the propositions from c holds'.*

A *biconsequence relation* is a set of bisequents satisfying the usual rules of Reflexivity, Monotonicity and Cut with respect to each of the two contexts. It can be seen as a fusion, or fibring, of two Scott consequence relations.

$\overline{u}$ will denote the complement of a set u of propositions. A pair (u, v) of sets of propositions is a *bitheory* of a biconsequence relation if $u : \overline{v} \nVdash \overline{u} : v$. Bitheories are closed with respect to the bisequents of a biconsequence relation. A set u is a *theory* of $\Vdash$, if (u, u) is a bitheory of $\Vdash$.

By a *bimodel* we will mean a pair of sets of propositions. A set of bimodels will be called a *binary semantics*.

Definition 5. *A bisequent $a : b \Vdash c : d$ is* valid *in a binary semantics $\mathcal{B}$, if, for any $(u, v) \in \mathcal{B}$, if $a \subseteq u$ and $b \subseteq \overline{v}$, then either $c \cap u \neq \emptyset$, or $d \cap \overline{v} \neq \emptyset$.*

The set of bisequents valid in a binary semantics forms a biconsequence relation. Moreover, any biconsequence relation is determined in this sense by its canonical semantics defined as the set of its bitheories. Now, any bimodel (u, v) can be viewed as a four-valued interpretation, where u is the set of true propositions, while v is the set of propositions that are not false. Biconsequence relations provide in this sense a syntactic formalism for four-valued reasoning.

A *bisequent theory* is an arbitrary set of bisequents. For any bisequent theory Δ there is a least biconsequence relation $\Vdash_\Delta$ containing it that describes the logical content of Δ. This allows us to extend the notions of a bitheory and propositional theory to arbitrary bisequent theories.

Structural rules Some additional structural rules for biconsequence relations play an important role in what follows. A biconsequence relation is *consistent*, if it satisfies

Consistency $\quad A : A \Vdash$

Consistency says that no proposition can be taken to hold without assuming it. This amounts to restricting the binary semantics to *consistent* bimodels, that is, bimodels (u, v) such that $u \subseteq v$. On the four-valued interpretation, no proposition can be both true and false in such models.

A biconsequence relation is *regular* if it satisfies

Regularity $\quad$ If $b : a \Vdash a : b$, then $: a \Vdash : b$.

Regularity asserts that an admissible set of assumptions should be compatible with taking these assumptions as actually holding. It holds for a *quasi-reflexive* binary semantics in which, for any bimodel (u, v), (v, v) is also a bimodel.

Local four-valued connectives Any four-valued connective is definable in biconsequence relations via introduction and elimination rules. For this study, however, we will restrict ourselves to the *locally classical* connectives that behave as ordinary classical connectives with respect to each of the two contexts. A functionally complete basis for such connectives is provided by the following two connectives. The first is a four-valued *conjunction*:

$A \wedge B$ is true iff A is true and B is true.
$A \wedge B$ is false iff A is false or B is false.

The second is the *local negation* $\neg$:

$\neg A$ is true iff A is not true $\quad$ $\neg A$ is false iff A is not false

As usual, the disjunction $A \vee B$ is defined as $\neg(\neg A \wedge \neg B)$.

Let $\neg u$ denote the set $\{\neg A \mid A \in u\}$. Then any bisequent $a : b \Vdash c : d$ is equivalent to $\Vdash \bigvee(\neg a \cup c) : \bigwedge(d \cup \neg b)$. Consequently, the local connectives allow us to reduce bisequents to that of the form $\Vdash A : B$, where A and B are classical propositions. The latter bisequents will correspond to production rules $B \Rightarrow A$ of production and causal inference relations, discussed later.

Nonmonotonic Semantics

Nonmonotonic semantics of a biconsequence relation is defined as a set of its *explanatory closed* theories, namely theories for which presence and absence of propositions in the main context is explained (i.e., derived) when the theory itself is taken as the assumption context.

Exact semantics A most general kind of nonmonotonic reasoning is obtained by requiring that the assumption context should determine itself as a unique objective state.

Definition 6. *A theory u of a biconsequence relation $\Vdash$ is* exact, *if there is no $v \neq u$ such that (v, u) is a bitheory of $\Vdash$. The set of exact theories forms an* exact nonmonotonic semantics *of* $\Vdash$.

This semantics is nonmonotonic, since the set of exact theories does not change monotonically with the growth of the set of bisequents. Regular biconsequence relations constitute the underlying logic for this nonmonotonic semantics.

Default semantics A more familiar class of nonmonotonic models, extensions, correspond to extensions of default logic and stable models of logic programs.

Definition 7. *A theory u is an* extension *of a biconsequence relation, if there is no bitheory (u', u) such that $u' \subset u$. A* default nonmonotonic semantics *is the set of extensions.*

Any exact theory is an extension, though not vice versa. Extensions explain only why they have the propositions they have. In other words, we are relieved from the necessity of explaining why propositions do *not* belong to an extension. Now, Consistency ($A : A \Vdash$) amounts to refutation of any proposition that is assumed not to hold, so the default nonmonotonic semantics is precisely an exact semantics under the stronger logic of consistent biconsequence relations.

Interpretation of logic programs Biconsequence relations provide a logical basis of logic programming. General program rules $\mathbf{not}\, d, c \leftarrow a, \mathbf{not}\, b$ can be directly interpreted as bisequents $a : b \Vdash c : d$. Then we have

Theorem 2. *Stable models of a general logic program coincide with the extensions of the corresponding bisequent theory.*

Moreover, this correspondence is bidirectional, since any bisequent in a four-valued language is reducible to a set of bisequents without connectives, and hence to program rules.

Production and Causal Inference

Now we introduce a primary *logical* system for explanatory reasoning. The system of production inference can be viewed as a generalization of classical logic obtained by dropping the Reflexivity postulate of classical inference. It originates in input/output logics of (Makinson & van der Torre 2000). Biconsequence relations turn out to constitute a structural counterpart of this logical formalism.

Production inference relations are based on conditionals of the form $A \Rightarrow B$ saying '*A produces, or explains, B*'.

Definition 8. *A* (regular) production inference relation *is a binary relation $\Rightarrow$ on the set of classical propositions satisfying the following postulates:*

(Strengthening) *If $A \vDash B$ and $B \Rightarrow C$, then $A \Rightarrow C$;*

(Weakening) *If $A \Rightarrow B$ and $B \vDash C$, then $A \Rightarrow C$;*

(And) *If $A \Rightarrow B$ and $A \Rightarrow C$, then $A \Rightarrow B \wedge C$;*

(Cut) *If $A \Rightarrow B$ and $A \wedge B \Rightarrow C$, then $A \Rightarrow C$;*

(Truth) $\mathbf{t} \Rightarrow \mathbf{t}$;

(Falsity) $\mathbf{f} \Rightarrow \mathbf{f}$.

The most significant 'omission' of the above set of postulates is the absence of reflexivity $A \Rightarrow A$. It is this feature that creates a possibility of nonmonotonic reasoning.

Production rules are extended to rules with sets of propositions in premises using the familiar compactness recipe:

$$u \Rightarrow A \;\equiv\; \bigwedge a \Rightarrow A, \text{ for some finite } a \subseteq u.$$

Let $\mathcal{C}(u)$ denote the set $\{A \mid u \Rightarrow A\}$. The production operator $\mathcal{C}$ is monotonic and continuous, and it plays the role of a derivability operator. A *theory* of a production relation is a deductively closed set u such that $\mathcal{C}(u) \subseteq u$.

A semantics for production relations can be given in terms of pairs of deductive theories called, as before, bimodels.

Definition 9. *A* classical bimodel *is a pair of consistent deductively closed sets. A* classical binary semantics *is a set of classical bimodels. A classical binary semantics is* consistent, *if $u \subseteq v$, for any bimodel (u, v).*

The validity of production rules is defined as follows.

Definition 10. *A production rule $A \Rightarrow B$ is* valid *in a classical binary semantics $\mathcal{B}$ if, for any bimodel (u, v) from $\mathcal{B}$, $A \in v$ only if $B \in u$.*

Consistent classical binary semantics provides an adequate representation of regular production inference relations (see (Bochman 2004a)).

A set of production rules will be called a *causal theory*. Semantics of a causal theory will be identified with the semantics of the least production relation including it.

General nonmonotonic semantics Production inference determines also a natural nonmonotonic semantics. Namely, the fact that the operator $\mathcal{C}$ is not reflexive creates an important distinction among theories of a production relation.

Definition 11. *A* nonmonotonic production semantics *of a production inference relation is the set of all its* exact theories, *namely sets u of propositions such that $u = \mathcal{C}(u)$.*

An exact theory describes an informational state in which every proposition is *explained* by other propositions accepted in this state, so it complies with the explanatory closure assumption. The nonmonotonic semantics is globally nonmonotonic, since adding new production rules may lead to a nonmonotonic change of the associated semantics, even though production rules themselves are monotonic, since they satisfy Strengthening (the Antecedent). It can be shown that regular production inference relations provide a logic of reasoning with exact theories.

Causal Inference A regular production inference relation is *causal* if it satisfies

(Or) If $A \Rightarrow C$ and $B \Rightarrow C$, then $A \vee B \Rightarrow C$.

Causal inference sanctions reasoning by cases, so the production rules can already be seen as causal rules. The semantics of causal production relations is provided by usual possible worlds models (W, R, V), where W is a set of worlds, R an accessibility relation, and V a valuation function. Causal inference relations require that R should be *quasi-reflexive*, that is, $\alpha R \beta$ holds only if $\alpha R \alpha$.

Definition 12. *A rule $A \Rightarrow B$ is* valid *in a possible worlds model (W, R, V) if, for any $\alpha, \beta \in W$ such that $\alpha R \beta$, if A holds in α, then B holds in β.*

Causal inference relations constitute a logical counterpart of biconsequence relations for the language of local classical connectives. By this correspondence, a production rule $A \Rightarrow B$ can be seen as a bisequent $\Vdash B : A$ saying that if A is assumed, then B should hold.

Causal nonmonotonic semantics For the causal reading of production rules, we restrict the nonmonotonic semantics to exact theories that are worlds and obtain the *causal nonmonotonic semantics* that coincides with the semantics of causal theories from (McCain & Turner 1997).

A world α is an exact world of a production inference relation if and only if, for any propositional atom p,

$$p \in \alpha \text{ iff } \alpha \Rightarrow p \qquad \text{and} \qquad \neg p \in \alpha \text{ iff } \alpha \Rightarrow \neg p.$$

Causal inference relations provide an adequate logic for the causal nonmonotonic semantics. Moreover, in the general correspondence between causal and biconsequence relations, the causal semantics corresponds to the exact nonmonotonic semantics of biconsequence relations.

By the above description, the exact worlds are determined ultimately by *determinate* rules $A \Rightarrow l$, where l is a literal. The causal nonmonotonic semantics of a determinate causal theory Δ coincides with the classical semantics of its *completion*, $comp(\Delta)$, defined as the set of classical formulas

$$p \leftrightarrow \bigvee\{A \mid A \Rightarrow p \in \Delta\} \quad \neg p \leftrightarrow \bigvee\{A \mid A \Rightarrow \neg p \in \Delta\},$$

for any atom p, plus the set $\{\neg A \mid A \Rightarrow \mathbf{f} \in \Delta\}$.

As for biconsequence relations, the default nonmonotonic semantics of causal theories can be obtained by imposing a causal postulate corresponding to Consistency postulate for biconsequence relations: for any propositional atom p,

(Default Negation) $\neg p \Rightarrow \neg p$.

Default Negation stipulates that negations of atomic propositions are self-explanatory, and hence it provides a simple causal expression for Reiter's Closed World Assumption. This kind of causal inference can also be used as a logical basis for logic programming. Namely, a program rule $\mathbf{not}\, d, c \leftarrow a, \mathbf{not}\, b$ can be faithfully interpreted as the causal rule $d, \neg b \Rightarrow \bigwedge a \rightarrow \bigvee c$ (see (Bochman 2004b)).

Epistemic Explanatory Reasoning

Epistemic formalisms of default and modal nonmonotonic logics find their natural place in the framework of supraclassical biconsequence relations, defined below. On epistemic understanding of biconsequence relations, the main and assumption contexts are treated, respectively, as contexts of knowledge and belief: propositions in the main context are viewed as known, while the assumption context forms the associated set of beliefs. Accordingly, both contexts correspond to incomplete deductive theories.

Definition 13. *A biconsequence relation in a classical language is* supraclassical, *if it satisfies*

Supraclassicality *If $a \vDash A$, then $a : \Vdash A :$ and $: A \Vdash : a$.*

Falsity $\mathbf{f} : \Vdash$ *and* $\Vdash : \mathbf{f}$.

Due to Supraclassicality, both contexts respect the classical entailment. In addition, sets of positive premises and negative conclusions can be replaced by their conjunctions, but positive conclusion sets and negative premise sets are not replaceable in this way by classical disjunctions.

A semantics of supraclassical biconsequence relations is obtained from the general binary semantics by requiring that bimodels are pairs of consistent deductively closed sets (see Definition 9). Structural rules for biconsequence relations are also extended to the supraclassical case.

A supraclassical biconsequence relation is *saturated*, if it is consistent, regular, and satisfies the following postulate:

Saturation $\Vdash A \vee B, \neg A \vee B : B$.

For a deductive theory u, let $u\bot$ denote the set of all maximal sub-theories of u, plus u itself. Then a classical bimodel (u, v) is *saturated*, if $u \in v\bot$. A classical binary semantics $\mathcal{B}$ is *saturated* if it is regular, and all its bimodels are saturated. Such a semantics provides an adequate interpretation for saturated biconsequence relations.

Classical nonmonotonic semantics The notions of an exact theory and extension can be directly extended to supraclassical consequence relations, but they form now deductively closed sets. Supraclassical biconsequence relations that are consistent and regular constitute a maximal logic adequate for extensions. For such biconsequence relations, an extension is described as a set of propositions that are provable if taken as the set of assumptions:

$$u = \{A \mid : \overline{u} \Vdash A : u\}.$$

Under this nonmonotonic semantics, bisequent theories having only rules of the form $a{:}b \Vdash c{:}$ provide an exact representation for the *disjunctive default logic* (Gelfond *et al.* 1991). For singular rules $a{:}b \Vdash C{:}$, it reduces to the original default logic of (Reiter 1980).

The following nonmonotonic semantics constitutes an exact non-modal counterpart of Moore's autoepistemic logic.

Definition 14. *A theory u of a supraclassical biconsequence relation $\Vdash$ is an* expansion *of $\Vdash$, if, for any $v \in u\bot$ such that $v \neq u$, the pair (v, u) is not a bitheory of $\Vdash$. The set of expansions determines the* autoepistemic semantics *of $\Vdash$.*

Any extension is an expansion, though not vice versa. In fact, expansions can be precisely characterized as extensions of saturated biconsequence relations.

The next result states an important sufficient condition for coincidence of expansions and extensions. A bisequent theory is *positively simple*, if positive premises and positive conclusions of any bisequent are sets of classical literals.

Theorem 3. *If a bisequent theory is positively simple, then its expansions coincide with classical extensions.*

Bisequents with only classical literals in premises and conclusions correspond to program rules of *extended* logic programs with classical negation. The semantics of such programs is determined by answer sets that coincide with extensions of respective bisequent theories. Moreover, such bisequent theories are positively simple, so by the above theorem extended logic programs obliterate the distinction between extensions and expansions. This is the logical basis for a possibility of representing extended logic programs also in autoepistemic logic (Lifschitz & Schwarz 1993).

Modal Nonmonotonic Logics

A representation of modal nonmonotonic logics can be given in the framework of modal biconsequence relations. The modal operator L in this setting reflects assumptions (or beliefs) as propositions in the main context.

Definition 15. *A supraclassical biconsequence relation in a modal language will be called* modal *if it satisfies*

Positive Reflection $A : \Vdash LA :$,

Negative Reflection $: LA \Vdash : A$,

Negative Introspection $: A \Vdash \neg LA :$.

Any theory of a modal biconsequence relation is a *modal stable set* in the sense of (Moore 1985), and hence modal extensions and expansions will always be stable theories.

For a modal logic $\mathcal{M}$, a modal biconsequence relation $\Vdash$ is an *$\mathcal{M}$-biconsequence relation*, if $\Vdash A:$, for every modal axiom A of $\mathcal{M}$. A modal biconsequence relation is an *F-biconsequence relation*, if it is regular and satisfies

F $\Vdash A, LA \to B : B$.

F-biconsequence relations provide a concise representation for the modal logic S4F. Any bisequent $a{:}b \Vdash c{:}d$ of an F-biconsequence relation is already reducible to a modal formula $\bigwedge(La \cup L\neg Lb) \to \bigvee(Lc \cup L\neg Ld)$. Consequently, any bisequent theory Δ in such a logic is reducible to an ordinary modal theory that we will denote by $\tilde{\Delta}$.

Modal nonmonotonic semantics By varying the modal logic, we obtain a range of modal nonmonotonic semantics. A set of propositions is an *$\mathcal{M}$-extension* (*$\mathcal{M}$-expansion*) of a bisequent theory Δ, if it is an extension (resp. expansion) of the least $\mathcal{M}$-biconsequence relation containing Δ. $\mathcal{M}$-extensions of plain modal theories Δ coincide with *$\mathcal{S}$-expansions* from (Marek, Schwarz, & Truszchinski 1993).

Since modal extensions and expansions are modal stable theories, they are determined by their non-modal subsets. This suggests a possibility of reducing modal nonmonotonic reasoning to a nonmodal one, and vice versa.

Let u_o denote the set of all non-modal propositions in u. Similarly, $_o\Vdash$ will denote the restriction of a modal biconsequence relation $\Vdash$ to the non-modal language.

Theorem 4. *If u is a modal stable theory, then u is an extension of $\Vdash$ if and only if u_o is an extension of $_o\Vdash$.*

By this result, non-modal supraclassical biconsequence relations are sufficiently expressive to capture modal nonmonotonic reasoning. In the other direction, in modal F-biconsequence relations any bisequent theory Δ is reducible to a usual modal theory $\tilde{\Delta}$. This allows us to use ordinary modal logical formalisms for representing non-modal nonmonotonic reasoning. Thus, the following result generalizes the result of (Truszczyński 1991) about a modal embedding of default theories.

Theorem 5. *If Δ is an objective bisequent theory, then classical extensions of Δ are precisely objective parts of S4F-extensions of $\tilde{\Delta}$.*

We end with considering modal expansions. Two kinds of expansions are important for a general description. The first is *stable expansions* of Moore's autoepistemic logic. They coincide with $\mathcal{M}$-expansions for any modal logic $\mathcal{M}$ in the range $5 \subseteq \mathcal{M} \subseteq$ KD45. The second kind of expansions is *reflexive expansions* of Schwarz' reflexive autoepistemic logic (Schwarz 1992). They coincide with $\mathcal{M}$-expansions for any modal logic in the range $KT \subseteq \mathcal{M} \subseteq$ SW5. 'Normal' expansions in general can be viewed as a combination of these two kinds of expansions:

Theorem 6. *A set of propositions is a K-expansion of a modal bisequent theory Δ iff it is both a stable and reflexive expansion of Δ.*

Conclusions

Despite clear success, twenty five years of nonmonotonic reasoning research have shown that we need a deep breath and long term objectives in order to make nonmonotonic reasoning a viable tool for the challenges posed by AI. There is still much to be done in order to meet the actual complexity of reasoning tasks required by the latter. In particular, the relation between the two principal paradigms of nonmonotonic reasoning has emerged as the main theoretical problem for a future development of the field.

By a presumably vague, but inspiring analogy, in nonmonotonic reasoning we have both a global relativity theory of preferential reasoning and a local quantum mechanics of explanatory reasoning. So, what we need is a unified theory of nonmonotonic reality. As in physics, however, this unified theory is not going to emerge as a straightforward juxtaposition of these components.

References

Andreka, H.; Ryan, M.; and Schobbens, P.-Y. 2002. Operators and laws for combining preference relations. *Journal of Logic and Computation* 12:13–53.

Belnap, Jr, N. D. 1977. A useful four-valued logic. In Dunn, M., and Epstein, G., eds., *Modern Uses of Multiple-Valued Logic*. D. Reidel. 8–41.

Bochman, A. 2001. *A Logical Theory of Nonomonotonic Inference and Belief Change*. Springer.

Bochman, A. 2003. Brave nonmonotonic inference and its kinds. *Annals of Mathematics and Artificial Intelligence* 39:101–121.

Bochman, A. 2004a. A causal approach to nonmonotonic reasoning. *Artificial Intelligence* 160:105–143.

Bochman, A. 2004b. A causal logic of logic programming. In Dubois, D.; Welty, C.; and Williams, M.-A., eds., *Proc. Ninth Conference on Principles of Knowledge Representation and Reasoning, KR'04*, 427–437.

Bochman, A. 2005. *Explanatory Nonmonotonic Reasoning*. World Scientific.

Dung, P. M., and Son, T. C. 2001. An argument-based approach to reasoning with specificity. *Artificial Intelligence* 133:35–85.

Gabbay, D. M. 1985. Theoretical foundations for nonmonotonic reasoning in expert systems. In Apt, K. R., ed., *Logics and Models of Concurrent Systems*. Springer.

Geffner, H. 1992. *Default Reasoning. Causal and Conditional Theories*. MIT Press.

Gelfond, M.; Lifschitz, V.; Przymusińska, H.; and Truszczyński, M. 1991. Disjunctive defaults. In *Proc. Second Int. Conf. on Principles of Knowledge Representation and Reasoning, KR'91*, 230–237.

Giunchiglia, E.; Lee, J.; Lifschitz, V.; McCain, N.; and Turner, H. 2004. Nonmonotonic causal theories. *Artificial Intelligence* 153:49–104.

Horty, J. F. 1994. Some direct theories of nonmonotonic inheritance. In Gabbay, D. M.; Hogger, C. J.; and Robinson, J. A., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming* **3***: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford: Oxford University Press.

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

Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55:1–60.

Lehmann, D. 1989. What does a conditional knowledge base entail? In Brachman, R., and Levesque, H. J., eds., *Proc. 2nd Int. Conf. on Principles of Knowledge Representation and Reasoning, KR'89*, 212–222. Morgan Kaufmann.

Lehmann, D. 1995. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence* 15:61–82.

Lifschitz, V., and Schwarz, G. 1993. Extended logic programs as autoepistemic theories. In Pereira, L. M., and Nerode, A., eds., *Proc. Second Int. Workshop on Logic Programming and Nonmonotonic Reasoning*, 101–114. MIT Press.

Lifschitz, V.; Pearce, D.; and Valverde, A. 2001. Strongly equivalent logic programs. *ACM Transactions on Computational Logic* 2:526–541.

Lifschitz, V. 1985. Computing circumscription. In *Proc. 9th Int. Joint Conf. on Artificial Intelligence, IJCAI-85*, 121–127. Morgan Kaufmann.

Makinson, D., and van der Torre, L. 2000. Input/Output logics. *Journal of Philosophical Logic* 29:383–408.

Makinson, D. 1989. General theory of cumulative inference. In Reinfrank, M., ed., *Nonmonotonic Reasoning*, volume 346 of *LNAI*, 1–18. Springer.

Marek, W., and Truszczyński, M. 1993. *Nonmonotonic Logic, Context-Dependent Reasoning*. Springer.

Marek, V. W.; Schwarz, G. F.; and Truszchinski, M. 1993. Modal nonmonotonic logics: ranges, characterization, computation. *Journal of ACM* 40:963–990.

McCain, N., and Turner, H. 1997. Causal theories of action and change. In *Proceedings AAAI-97*, 460–465.

Moore, R. C. 1985. Semantical considerations on nonmonotonic logic. *Artificial Intelligence* 25:75–94.

Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings of the Third Conference on Theoretical Aspects of Reasoning About Knowledge (TARK'90)*, 121–135. San Mateo, CA: Morgan Kaufmann.

Poole, D. 1988. A logical framework for default reasoning. *Artificial Intelligence* 36:27–47.

Poole, D. 1989. Explanation and prediction: An architecture for default and abductive reasoning. *Computational Intelligence* 5:97–110.

Reiter, R. 1978. On closed world data bases. In Gallaire, H., and Minker, J., eds., *Logic and Data Bases*. Plenum Press. 119–140.

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

Reiter, R. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamic Systems*. MIT Press.

Schlechta, K. 2004. *Coherent Systems*, volume 2 of *Studies in Logic and Practical Reasoning*. Amsterdam: Elsevier.

Schwarz, G. 1992. Reflexive autoepistemic logic. *Fundamenta Informaticae* 17:157–173.

Shoham, Y. 1988. *Reasoning about Change*. Cambridge University Press.

Tan, S.-W., and Pearl, J. 1995. Specificity and inheritance in default reasoning. In *Proceedings Int. Joint Conf. on Artificial Intelligence, IJCAI-95*, 1480–1486.

Truszczyński, M. 1991. Modal interpretations of default logic. In Myopoulos, J., and Reiter, R., eds., *Proceedings Int. Joint Conf. on Artificial Intelligence, IJCAI'91*, 393–398. San Mateo, Calif.: Morgan Kaufmann.

Turner, H. 1999. A logic of universal causation. *Artificial Intelligence* 113:87–123.

van Benthem, J. 1984. Foundations of conditional logic. *Journal of Philosophical Logic* 13:303–349.