Cumulative Inference Relations for JTMS and Logic Programming

Gerhard Brewka
GMD, Postfach 1240
D-5205 Sankt Augustin, Germany

David Makinson
Les Etangs B2, La Ronce
92410 Ville d'Avray, France

Karl Schlechta
IBM Germany, Wissenschaftliches Zentrum, IWBS
P.O. Box 80 08 80, D-7000 Stuttgart 80, Germany
and University of Hamburg, Germany

13 June 1991

Abstract

This paper makes three main points. We observe first that the inference relation induced by a set of JTMS justification rules under the grounded model semantics (or equivalently, by a logic program with negation under the Gelfond-Lifschitz semantics) is not in general cumulative: the addition to a set of assumptions of some of the derivable conclusions may lead to a loss of others.

We then show how cumulativity may be restored by adapting a technique recently applied by Brewka to default logic. The basic idea is to upgrade the universe of discourse: replace the elementary propositions, between which inference customarily takes place, by more complex items consisting of elementary propositions indexed by certain of the "reasons" that lead to their acceptance.

However, as we finally show, the indexed JTMS still has a shortcoming: it does not give an adequate treatment of the phenomenon of "floating conclusions". The problem of finding an alternative aproach that handles floating conclusions adequately without losing cumulativity again, remains open.

1 Background

Research in the area of nonmonotonic reasoning has produced many different formalizations: circumscription of McCarthy [11], [12], default logic of Reiter [15], and autoepistemic logic of Moore [13], to name just a few of the most popular. The inference relations that these formalizations generate are not monotonic, but nevertheless they sometimes have interesting and useful formal properties. These properties can be used to compare the behaviour of the different formalisms, and can also serve as considerations when weighing the "quality" of the inference relations induced. This field has been investigated by Makinson [10], Kraus, Lehmann and Magidor [8] and others.

One particularly desirable property of inference relations is cumulativity. Intuitively, cumulativity says that adding a conclusion to a set of premises does not change the available conclusions. It is thus a kind of stability. More formally, *cumulativity* in its finitary form can be expressed as the condition:

$$\text{If } A \mid\!\sim x \text{ then } A \mid\!\sim y \text{ iff } A \cup \{x\} \mid\!\sim y,$$

where x, y are any propositions, A is any set of propositions, and $\mid\!\sim$ is the inference relation in question.

Each of the two halves of this notion is of interest. The condition:

$$\text{If } A \mid\!\sim x \text{ and } A \cup \{x\} \mid\!\sim y \text{ then } A \mid\!\sim y$$

is known as *cut* (or cumulative transitivity) and is generally easy to satisfy. The other half:

$$\text{If } A \mid\!\sim x \text{ and } A \mid\!\sim y \text{ then } A \cup \{x\} \mid\!\sim y$$

is known as *cautious monotony* (or cumulative monotony) and is much more difficult to satisfy.

For example, it turns out that in Reiter's default logic, the skeptical notion of inference defined by the intersection of extensions satisfies cut, but not in general cautious monotony. Indeed this failure happens in two distinct ways. In the limiting case that a set A of propositions has no extensions under default rules D, so that vacuously $A \mid\!\sim x, y$ for all propositions x, y, it may nevertheless happen that the superset $A' = A \cup \{x\}$ of A does have an extension, not containing some proposition y. For example: let D contain the unique rule $\emptyset; \neg x/x$ and put $A = \emptyset$, $A' = \{x\}$. In the principal case that A does have extensions, cumulativity may fail for another reason. The following example is from Makinson [9]: D contains two rules $\emptyset; a/a$ and $a \vee b; \neg a/\neg a$ whilst $A = \emptyset$ and $A' = \{a \vee b\}$. Then A has $Cn(a)$ - i.e. the set of all classical consequences of a - as unique extension under D, whereas A' has two extensions: $Cn(a)$ and $Cn(\neg a, b)$, so that $A \mid\!\sim a$ but not $A' \mid\!\sim a$.

In the limiting case that A has no extensions under D cumulativity can be restored by the simple trick of redefining $\mid\!\sim$ for that case, putting there $A \mid\!\sim x$ iff $x \epsilon Cn(A)$. The principal case is however a more delicate matter. It has been shown by Brewka [3] that the cumulativity of default logic can be restored by upgrading the basic elements of discourse. Instead of plain propositions, Brewka uses more complex items, consisting of propositions indexed by "reasons" for why they may be believed. This shift in ontology, accompanied by a suitable modification of the definition of an extension taking the indices into account, yields a cumulative version of default logic which Brewka calls CDL.

It is known that there are close links between default logic and the justification-based truth maintenance systems of Doyle [5], often known as JTMS. Indeed, as shown by several authors, eg Reinfrank, Dressler and Brewka [14], Brewka [2] and Elkan [6], every JTMS can be seen as a restriction to an impoverished language of a propositional default or autoepistemic logic. It is thus natural to ask whether the skeptical inference relations induced by the JTMS concepts may also fail cumulativity. Note that the counterexample described above in the context of default logic is not appropriate, as it uses disjunction and so cannot be expressed in the language of JTMS. A different kind of counterexample is thus needed. That is the subject of the following section.

We recall that Elkan [6] was able to establish an equivalence between grounded models of JTMS's and another computationally oriented focus of research, stable sets of purely propositional logic programs with negation under the semantics of Gelfond and Lifschitz [7]. Our results in this paper will thus also apply to the latter.

2 The Failure of Cumulativity in JTMS and LPN

There are several different but equivalent ways of defining the behaviour of a JTMS formally - see for example Reinfrank, Dressler and Brewka [14] for one of them. The presentation below is trivially equivalent to the elegant and simple one of Elkan [6]. We use lower case letters from anywhere in the alphabet for propositions, and upper case Roman letters for sets of propositions. Recall that a *justification rule* is a directed propositional clause $a \leftarrow b_1, \ldots, b_h, \neg c_1, \ldots, \neg c_k$ where a, the b's and the c's are all elementary propositions (atoms). It is usually assumed that the language has a finite set of elementary propositions; however such an assumption is not needed for any of the results of this paper.

Definition 2.1 *A* model *is just a set of elementary propositions. A model M is said to be* closed *under a justification rule $a \leftarrow b_1, \ldots, b_h, \neg c_1, \ldots, \neg c_k$ iff whenever all $b_i \epsilon M$ and no $c_i \epsilon M$, then $a \epsilon M$. A model is said to be closed under a set of justification rules iff it is closed under all of the rules in the set.*

Definition 2.2 *A model M is said to be* grounded *in a set A of propositions under a set Σ of justification rules iff M can be ordered as $(a_1, a_2, \ldots,)$ such that for each i, either $a_i \epsilon A$ or there is a justification rule $a \leftarrow b_1, \ldots, b_h, \neg c_1, \ldots, \neg c_k$ in Σ such that $a_i = a$, all $b_j \epsilon \{a_1, \ldots, a_{i-1}\}$ and all $c_j \notin M$.*

Strictly speaking, Elkan does not separate A from Σ, treating the elements of A as justification rules with an empty body. When dealing with the generated inference operations, with argument A, it is however convenient to separate out the two clearly. Elkan and some others also define groundedness only for models M that already satisfy A and are closed under Σ. That, indeed, will remain our ultimate area of application. But clearly the concept of groundedness makes sense quite generally, and it turns out that the availability of the unrestricted formulation simplifies proofs.

Definition 2.3 *We say that M is a* grounded model *of A under Σ, iff*
(1) $A \subseteq M$,
(2) M is closed under Σ, and
(3) M is grounded in A under Σ.
We write $gr_\Sigma(A)$ for the set of all grounded models of A under Σ, or more briefly $gr(A)$ when Σ is understood as fixed.

The task of a JTMS can be described as one of computing one, or all, of the elements of $gr_\Sigma(A)$ for any given Σ, A.

Note that conditions (1) and (2) of the definition of $gr_\Sigma(A)$ are preserved downwards along the A and Σ axes respectively, in the sense that if $A \subseteq M$ and $A' \subseteq A$ then $A' \subseteq M$, and if M is closed under Σ and $\Sigma' \subseteq \Sigma$ then M is closed under Σ'. On the other hand, condition (3) is preserved upwards along the A and Σ axes, in the sense that if M is grounded in A under Σ and $A \subseteq A'$, $\Sigma \subseteq \Sigma'$ then M is grounded in A' under Σ'. It is the interplay of these two preservation properties that gives $gr_\Sigma(A)$ much of its characteristic behaviour. In particular, when M, Σ are both fixed we immediately have the useful:

Lemma 2.1 *Whenever $M \epsilon gr_\Sigma(A)$ then $M \epsilon gr_\Sigma(B)$ for every B with $A \subseteq B \subseteq M$.*

Definition 2.4 *Let Σ be a set of justification rules. If A is a set of elementary propositions and x is an elementary proposition, we say that $A \mid\sim_\Sigma x$ iff either:*
(1) $gr_\Sigma(A)$ is non-empty and $x \epsilon \bigcap gr_\Sigma(A)$, or
(2) $gr_\Sigma(A)$ is empty and $x \epsilon A$.

In other words, in the principal case that $gr_\Sigma(A)$ is non-empty, $A \mid\sim_\Sigma x$ iff x is in every model $M \supseteq A$ satisfying Σ that is grounded in A under Σ. In the limiting case that $gr_\Sigma(A)$ is empty, and thus $\bigcap gr_\Sigma(A)$ is the set of all elementary propositions, we put $A \mid\sim_\Sigma x$ iff x is in A rather than iff it is in $\bigcap gr_\Sigma(A)$, in order to avoid trivial failure of cumulativity in that case. Clearly the inference relation $\mid\sim_\Sigma$ satisfies inclusion, i.e. $A \mid\sim_\Sigma x$ for every $x \epsilon A$. It is also straightforward to verify that the inference relation $\mid\sim_\Sigma$ satisfies cut, i.e.:

Observation 2.2 *Let Σ be a set of justification rules, A a set of elementary propositions, and x,y elementary propositions. If $A \mid\sim_\Sigma x$ and $A \cup \{x\} \mid\sim_\Sigma y$ then $A \mid\sim_\Sigma y$.*

Proof: As Σ is held fixed, we drop the subscript from $\mid\sim_\Sigma$ for ease of reading. Suppose $A \mid\sim x$ and $A \cup \{x\} \mid\sim y$; we want to show $A \mid\sim y$. For the pricipal case that $gr(A)$ is non-empty, let M be any grounded model of A under Σ; we need to show that $y \epsilon M$. By the first supposition $x \epsilon M$ and so $A \subseteq A \cup \{x\} \subseteq M$, so by lemma 2.1 M is a grounded model of $A \cup \{x\}$ under Σ, so by our second supposition $y \epsilon M$ as desired. In the limiting case that $gr(A)$ is empty, our first supposition gives $x \epsilon A$ so $A \cup \{x\} = A$ and we conclude trivially from our second supposition. □

Clearly the same verification goes through for the general version of cut, i.e. the rule that whenever $A \subseteq B \subseteq C(A)$ then $C(B) \subseteq C(A)$. Here C is the inference operation corresponding to the relation $\mid\sim$, and is defined by the equation $C(A) = \{x : A \mid\sim x\}$.

On the other hand, although cautious monotony is guaranteed in the limiting case that A has no grounded models, this is not so for the principal case:

Observation 2.3 *The JTMS relation $\mid\sim_\Sigma$ may fail cautious monotony, and hence cumulativity.*

Example of Failure of Cumulativity: Put Σ to be made up of the following justification rules: $a \leftarrow \neg b$; $c \leftarrow a$; $b \leftarrow c, \neg a$. Put $A = \emptyset$. Then A has just one grounded model under Σ, namely $\{a, c\}$, so that $A \mid\sim a$ and $A \mid\sim c$. On the other hand, $A \cup \{c\}$ has two grounded models under Σ, namely $\{a, c\}$ and $\{c, b\}$, so that we do not have $A \cup \{c\} \mid\sim a$.

A similar example giving the same pattern of grounded models, due to Jürgen Dix (personal communication) is as follows: $a \leftarrow \neg b$; $b \leftarrow \neg a$; $c \leftarrow \neg c$; $c \leftarrow \neg b$.

For the record, we also translate these examples into the language of logic programs with negation. According to the semantics of Gelfond and Lifschitz [7], in each case

$\{a, c\}$ is the unique stable set for $\Sigma \cup A$, but the addition of the derivable proposition c gives us two distinct stable sets $\{a, c\}$ and $\{c, b\}$ coinciding with the two grounded models.

Neither of the above two examples is stratified. This is no accident, for when Σ is stratified then $|\!\sim_\Sigma$ is cumulative, as we now show.

Recall that a set Σ of justification rules is called *stratified* in the sense of Apt, Blair and Walker [1] iff it is possible to assign non-negative integer ranks to all elementary propositions in such a way that for every rule $a \leftarrow b_1, \ldots, b_h, \neg c_1, \ldots, \neg c_k$ in Σ, $rank(a) \geq rank(b_i)$ and $rank(a) > rank(c_j)$ for all $i \leq h$ and $j \leq k$. Call a set Σ of justification rules *functional* (resp. *partially functional*) for grounded models iff every set A of elementary propositions has exactly one (resp. at most one) grounded model under Σ. Now it is well known from work of Gelfond and Lifschitz [7] and others that (i) every stratified set Σ of justification rules is functional for grounded models. Using definition 2.4 and lemma 2.1 above, it is also easy to verify that (ii) whenever Σ is partially functional for grounded models then its induced inference relation $|\!\sim_\Sigma$ is cumulative. For suppose $A |\!\sim x$ and $A |\!\sim y$; we need to show $A \cup \{x\} |\!\sim y$. In the principal case that $gr(A)$ is non-empty there is an $M \epsilon gr(A)$, so $A \subseteq A \cup \{x\} \subseteq M$ so by lemma 2.1, $M \epsilon gr(A \cup \{x\})$, so by partial functionality, $M = \bigcap gr(A \cup \{x\})$ so since $A |\!\sim y$ we have $A \cup \{x\} |\!\sim y$. In the limiting case that $gr(A)$ is empty, definition 2.4 tells us that $x \epsilon A$ so $A \cup \{x\} = A$ and thus since $A |\!\sim y$ we trivially have $A \cup \{x\} |\!\sim y$.

Finally we note that in the light of the above results Jürgen Dix [4] investigates the fortunes of cumulativity under certain more restricted semantics for JTMS and logic programs, notably the so-called "well founded" semantics that is intermediate between the stratified semantics and the grounded model/stable set semantics considered here.

3 Cumulative JTMS

One way of restoring cumulativity would be to use a translation into default logic, together with the restoration of cumulativity to the latter as carried out by Brewka [3]. However that procedure is rather indirect, unnecessarily complex (as default logic is more complex than JTMS) and makes it difficult to adapt the current labelling diagrams for ordinary non-cumulative JTMS to the cumulative version. For these reasons we proceed directly, borrowing from Brewka [3] the fundamental idea of upgrading the ontology, but with an independent and rather simpler technical development.

We refine the units of discourse between which inference takes place, from simple propositions to more complex items, called indexed propositions. These are propositions indexed by suitable traces of the reasons for believing them. It turns out that in the case of JTMS it is sufficient for the index to keep track of the outsupports that are involved in the "derivation" of the proposition. The indexing is rather simpler than that used in Brewka's reconstruction of default logic, which also keeps track of intermediate conclusions.

An *indexed proposition* is defined to be a pair (a, A) where a is an elementary proposition and A is a set of elementary propositions. We continue to use lower case letters from anywhere in the alphabet for propositions, and upper case Roman letters for sets of propositions. We introduce upper case script letters for sets of indexed propositions.

If $\mathcal{A}$ is a set of indexed propositions, we write *prop*$(\mathcal{A})$ for $\{a : (a, A) \epsilon \mathcal{A}$ for some $A\}$, and similarly $index(\mathcal{A})$ for $\bigcup\{A : (a, A) \epsilon \mathcal{A}$ for some $a\}$. The following definition parallels, in the indexed context, definition 2.1.

Definition 3.1 *An* indexed model *is a set $\mathcal{M}$ of indexed propositions such that prop($\mathcal{M}$) $\cap$ index($\mathcal{M}$) $= \emptyset$. We say that an indexed model $\mathcal{M}$ is* closed *under a justification rule $a \leftarrow b_1, \ldots, b_h, \neg c_1, \ldots, \neg c_k$ iff whenever $B_1, \ldots, B_h$ are sets of elementary propositions and all of $(b_1, B_1), \ldots, (b_h, B_h)$ are in $\mathcal{M}$ and no c_j is in prop($\mathcal{M}$) then (a,A) is in $\mathcal{M}$, where $A = B_1 \cup \ldots . \cup B_h \cup \{c_1, \ldots, c_k\}$.*

In order to simplify proofs, it is also convenient to decompose further the notion of groundedness. Roughly speaking, we replace the single linear ordering that occurs in definition 2.2 by a family of trees (one for each element of the model) serving as "potential proofs" of elementary propositions. These trees show explicitly the dependency relationships that can be associated with linear orderings, and thereby facilitate manipulation, notably by a form of composition.

Definition 3.2 *Let Σ be a set of justification rules and $\mathcal{A}$ a set of indexed propositions. By a* potential proof *of an indexed proposition (y, Y) from $\mathcal{A}$ using Σ we mean a finite tree with the following properties:*
(1) Each node of the tree is labelled with an indexed proposition;
(2) The root of the tree is labelled with the indexed proposition (y, Y);
(3) For every node n of the tree, labelled by an indexed proposition (x,X), either n has no nodes above it and $(x, X)\epsilon\mathcal{A}$, or there is a justification rule $a \leftarrow b_1, \ldots, b_h, \neg c_1, \ldots, \neg c_k$ in Σ such that:
(3.1) a=x,
(3.2) The node n has just h nodes immediately above it in the tree, labelled by indexed propositions $(b_1, B_1), \ldots, (b_h, B_h)$ for some $B_1, \ldots, B_h$,
(3.3) $X = \{c_1, \ldots, c_k\} \cup B_1 \cup \ldots . \cup B_h$.

Intuitively, a potential proof of (y, Y) from $\mathcal{A}$ using Σ is a tree that records a "derivation" of y in which all steps are well justified by rules in Σ in so far as their insupports are concerned, but are allowed to appeal to any outsupports whatsoever in order to fire rules. Nevertheless, the outsupports used are carefully recorded at each stage of the "derivation" for subsequent analysis, and moreover the recording, as described in clause (3.3), is cumulative in character.

Clearly, when π is a potential proof of (y, Y) then for every indexed proposition $(x, X)\epsilon\pi$, $X \subseteq Y$.

It is possible to compose potential proofs, in particular in the manner described in the following lemma, which will assist rapid proof of the cumulativity and loop properties for the inference relation shortly to be defined.

Lemma 3.1 *Let $\mathcal{A}$, $\mathcal{B}$ be sets of indexed propositions. Suppose that:*
(1) There is a potential proof π of (y, Y) from $\mathcal{B}$ using Σ,
(2) For each $(b_i, B_i)\epsilon\mathcal{B}$, there is a potential proof σ_i of (b_i, B_i) from $\mathcal{A}$ using Σ.
Then there is a potential proof of (y, Y) from $\mathcal{A}$ using Σ.

Proof: Take the finite tree π, and for each leaf node of π whose label (b_i, B_i) is not already in $\mathcal{A}$, continue the tree with a copy of σ_i. Clearly the resulting tree is still finite and is a potential proof of (y, Y) from $\mathcal{A}$ using Σ. $\Box$

Definition 3.3 *We say that an indexed model $\mathcal{M}$ is* grounded *in $\mathcal{A}$ under Σ iff for every indexed proposition $(y, Y)\epsilon\mathcal{M}$ there is a potential proof of (y,Y) from $\mathcal{A}$ using Σ.*

Note that all of the indices of such a potential proof π will be compatible with the propositions of $\mathcal{M}$, in the sense that for all $(x, X)\epsilon\pi$, $X \cap prop(\mathcal{M}) = \emptyset$. For as already remarked after definition 3.2 we have $X \subseteq Y$, and since $(y, Y)\epsilon\mathcal{M}$ we also have $Y \subseteq index(\mathcal{M})$. Putting these two together gives $X \cap prop(\mathcal{M}) \subseteq Y \cap prop(\mathcal{M}) \subseteq index(\mathcal{M}) \cap prop(\mathcal{M}) = \emptyset$ by definition 3.1 of an indexed model. If in addition $\mathcal{A} \subseteq \mathcal{M}$ and $\mathcal{M}$ is closed under Σ, then also all the propositions of such a π will be propositions of $\mathcal{M}$, in the sense that for all $(x, X)\epsilon\pi$, $x\epsilon prop(\mathcal{M})$. This is easily seen by induction from the leaves of π to its root, using the hypotheses that $\mathcal{A} \subseteq \mathcal{M}$ and $\mathcal{M}$ is closed under Σ together with the already noted $X \cap prop(\mathcal{M}) = \emptyset$.

Thus if $\mathcal{A} \subseteq \mathcal{M}$, $\mathcal{M}$ is closed under Σ, and $\mathcal{M}$ is grounded in $\mathcal{A}$ under Σ, then every $(y, Y)\epsilon\mathcal{M}$ has a potential proof π that is authorized by $\mathcal{M}$ in the sense that for all $(x, X)\epsilon\pi$ we have $X \cap prop(\mathcal{M}) = \emptyset$ and $x\epsilon prop(\mathcal{M})$. This provides the rationale for the following definition.

Definition 3.4 *We say that $\mathcal{M}$ is a* grounded (indexed) model *of $\mathcal{A}$ under Σ iff*
(1) $\mathcal{A} \subseteq \mathcal{M}$,
(2) $\mathcal{M}$ is closed under Σ, and
(3) $\mathcal{M}$ is grounded in $\mathcal{A}$ under Σ.
We write the set of grounded (indexed) models of $\mathcal{A}$ as $gr_\Sigma(\mathcal{A})$, or more briefly as $gr(\mathcal{A})$ when Σ is understood as fixed.

Noting that groundedness is again preserved upwards along the $\mathcal{A}$ axis, it is again immediate from the definitions that:

Lemma 3.2 *If $\mathcal{M}\epsilon gr_\Sigma(\mathcal{A})$ then $\mathcal{M}\epsilon gr_\Sigma(\mathcal{B})$ for every $\mathcal{B}$ with $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{M}$.*

It is not difficult to show that this definition of a grounded model of $\mathcal{A}$ agrees with that of section 2 in the case that all index sets in elements of $\mathcal{A}$ are empty. Put formally, we have the following, whose verification is straightforward but tedious and so omitted.

Observation 3.3 *Let Σ be a set of justification rules, and $\mathcal{A}$ a set of indexed propositions. Then the function $\mathcal{M} \to prop(\mathcal{M})$ is a surjective mapping from the set of (indexed) grounded models of $\mathcal{A}$ under Σ to the set of all (unindexed) grounded models M of $prop(\mathcal{A})$ under Σ that satisfy the condition that $M \cap index(\mathcal{A}) = \emptyset$.*

Corollary 3.4 *Let Σ be a set of justification rules, and $\mathcal{A}$ a set of indexed propositions. If $index(\mathcal{A})$ is empty then the function $\mathcal{M} \to prop(\mathcal{M})$ is a surjective mapping from the set of (indexed) grounded models of $\mathcal{A}$ under Σ to the set of all (unindexed) grounded models of $prop(\mathcal{A})$ under Σ.*

With these concepts, we are now in a position to define a skeptical inference relation $\mid\!\sim_\Sigma$ between sets of indexed propositions (on the left) and indexed propositions (on the right) in a natural manner.

Definition 3.5 *Let Σ be any set of justification rules. If (x,X) is an indexed proposition and $\mathcal{A}$ is a set of indexed propositions, we put $\mathcal{A} \mid\!\sim_\Sigma (x, X)$ iff either:*
(1) $gr_\Sigma(\mathcal{A})$ is non-empty and $(x, X)\epsilon \bigcap gr_\Sigma(\mathcal{A})$, or
(2) $gr_\Sigma(\mathcal{A})$ is empty and $(x, X)\epsilon\mathcal{A}$.

As usual, we shall drop the subscript Σ whenever it is understood as fixed. To express the proof of cumulativity briefly, it will also be convenient to see $\mid\!\sim$ as an inference operation C, defined by putting $C(\mathcal{A}) = \{(x, X) : \mathcal{A} \mid\!\sim (x, X)\}$. In this guise, definition 3.5 says $C(\mathcal{A}) = \bigcap gr(\mathcal{A})$ when $gr(\mathcal{A}) \neq \emptyset$ and $C(\mathcal{A}) = \mathcal{A}$ when $gr(\mathcal{A}) = \emptyset$. It will help visualize the proof of loop if we also write $\mathcal{A} \mid\!\sim \mathcal{B}$ for $\mathcal{A} \mid\!\sim b$ for all $b \epsilon \mathcal{B}$, in other words, for $\mathcal{B} \subseteq C(\mathcal{A})$.

Lemma 3.5 *If $\mathcal{A} \mid\!\sim \mathcal{B}$ then every indexed model grounded in $\mathcal{B}$ is grounded in $\mathcal{A}$.*

Proof: Suppose $\mathcal{A} \mid\!\sim \mathcal{B}$. In the case that $gr(\mathcal{A})$ is empty we have $\mathcal{B} \subseteq \mathcal{A}$ so trivially the lemma holds. In the case that $gr(\mathcal{A})$ is not empty, there is an $\mathcal{N} \epsilon gr(\mathcal{A})$ and so $\mathcal{B} \subseteq \mathcal{N}$. Now suppose $\mathcal{M}$ is an indexed model that is grounded in $\mathcal{B}$; we want to show that $\mathcal{M}$ is grounded in $\mathcal{A}$. Let $(y, Y) \epsilon \mathcal{M}$. Since $\mathcal{M}$ is grounded in $\mathcal{B}$, there is a potential proof π of (y, Y) from $\mathcal{B}$ using Σ. Since on the other hand $\mathcal{N} \epsilon gr(\mathcal{A})$, $\mathcal{N}$ is grounded in $\mathcal{A}$ so for each $(b_i, B_i) \epsilon \mathcal{B} \subseteq \mathcal{N}$ there is a potential proof σ_i of (b_i, B_i) from $\mathcal{A}$ using Σ. Hence by lemma 3.1 there is a potential proof of (y, Y) from $\mathcal{A}$ using Σ. Thus $\mathcal{M}$ is grounded in $\mathcal{A}$ and the proof is complete. $\Box$

Lemma 3.5 is a key observation, on which our method of proof of cumulativity and loop rests. It illustrates the usefulness, as remarked after definition 2.2, of ensuring that the concept of groundedness of a model in a set of premises is well-defined even in the case that the model does not satisfy those premises. Evidently it is not true to say that whenever $\mathcal{A} \mid\!\sim \mathcal{B}$ then every grounded indexed model of $\mathcal{B}$ is also a model of $\mathcal{A}$.

We shall show cumulativity in its general "infinitary" form: $\mathcal{A} \subseteq \mathcal{B} \subseteq C(\mathcal{A})$ implies $C(\mathcal{A}) = C(\mathcal{B})$. Clearly this implies the finitary form presented in section 1, taking $\mathcal{B}$ to be the union of $\mathcal{A}$ with a singleton.

Theorem 3.6 *Let Σ be any set of justification rules. Then the inference relation $\mid\!\sim_\Sigma$ defined above satisfies inclusion and cumulativity.*

Proof: For inclusion, suppose $(a, A) \epsilon \mathcal{A}$. If $gr(\mathcal{A}) = \emptyset$ then immediately $\mathcal{A} \mid\!\sim (a, A)$. Suppose $\mathcal{M} \epsilon gr(\mathcal{A})$. Then by clause (1) of the definition of $gr(\mathcal{A})$, we have $\mathcal{A} \subseteq \mathcal{M}$ so $(a, A) \epsilon \mathcal{M}$ and we are done.

For cumulativity, suppose $\mathcal{A} \subseteq \mathcal{B} \subseteq C(\mathcal{A})$. We need to show that $C(\mathcal{A}) = C(\mathcal{B})$. In the limiting case that $gr(\mathcal{A})$ is empty, we have $C(\mathcal{A}) = \mathcal{A}$ and we are done. So suppose that $gr(\mathcal{A})$ is not empty.

For the easy half of cumulativity, i.e. cut, we need to show that $C(\mathcal{B}) \subseteq C(\mathcal{A})$. Clearly it will suffice to show that $gr(\mathcal{B})$ is also non-empty and $gr(\mathcal{A}) \subseteq gr(\mathcal{B})$. Clearly, the latter suffices for the former. Now if $\mathcal{M} \epsilon gr(\mathcal{A})$ we have $\mathcal{A} \subseteq \mathcal{B} \subseteq C(\mathcal{A}) = \bigcap gr(\mathcal{A}) \subseteq \mathcal{M}$, and so by lemma 3.2, $\mathcal{M} \epsilon gr(\mathcal{B})$ as desired.

For the tricky half of cumulativity, i.e. cautious monotony, we need to show $C(\mathcal{A}) \subseteq C(\mathcal{B})$. Since $gr(\mathcal{A})$ and $gr(\mathcal{B})$ are non-empty, it suffices to show $gr(\mathcal{B}) \subseteq gr(\mathcal{A})$. Let $\mathcal{M} \epsilon gr(\mathcal{B})$. Then $\mathcal{B} \subseteq \mathcal{M}$, $\mathcal{M}$ is closed under Σ, and $\mathcal{M}$ is grounded in $\mathcal{B}$ under Σ. Since $\mathcal{A} \subseteq \mathcal{B}$ we have $\mathcal{A} \subseteq \mathcal{M}$, and since $\mathcal{B} \subseteq C(\mathcal{A})$ lemma 3.5 tells us that $\mathcal{M}$ is grounded in $\mathcal{A}$ under Σ, so that $\mathcal{M} \epsilon gr(\mathcal{A})$ as desired. $\Box$

Example of Restoration of Cumulativity: We consider the example of non-cumulativity from section 2, to see why it no longer works in the indexed context. Recall that Σ consists of the three justification rules: $a \leftarrow \neg b$; $c \leftarrow a$; $b \leftarrow c, \neg a$. Let $\mathcal{A} = \emptyset$. Under our definitions, there is a unique grounded model of $\mathcal{A}$ under Σ, namely $\mathcal{M} = \{(a, \{b\}), (c, \{b\})\}$, corresponding to the model $\{a, c\}$ in the unindexed case. Now let $\mathcal{A}' = \{(c, \{b\})\}$. This continues to have $\mathcal{M}$ as its unique grounded model under Σ. In the unindexed case, the set $A' = \{c\}$ had a second grounded model giving trouble, namely $\{c, b\}$. And in the indexed case, of course the indexed set $\mathcal{A}'' = \{(c, \emptyset)\}$ with empty index likewise has a second grounded model, namely $\{(c, \emptyset), (b, \{a\})\}$. But the index b in our indexed set $\mathcal{A}' = \{(c, \{b\})\}$ prevents $\mathcal{A}'$ from having a second grounded model. For example, the "model" $\mathcal{N} = \{(c, \{b\}), (b, \{b, a\})\}$ will not do, for we have $prop(\mathcal{N}) \cap index(\mathcal{N}) \neq \emptyset$, so it is not a model at all.

The indexed JTMS inference relations given by definition 3.5 also have another desirable property, known as *loop*. In its general "infinitary" form, this is the condition that for all $\mathcal{A}_1, \ldots, \mathcal{A}_n$, if $\mathcal{A}_1 \mid\!\sim \mathcal{A}_2 \mid\!\sim \ldots . \mid\!\sim \mathcal{A}_n \mid\!\sim \mathcal{A}_1$ then $C(\mathcal{A}_i) = C(\mathcal{A}_j)$ for all $i, j \leq n$.

Theorem 3.7 *Let Σ be any set of justification rules. Then the inference relation $\mid\!\sim\; := \;\mid\!\sim_\Sigma$ given by definition 3.5 satisfies loop.*

Proof: Suppose $\mathcal{A}_1 \mid\!\sim \mathcal{A}_2 \mid\!\sim \ldots . \mid\!\sim \mathcal{A}_n \mid\!\sim \mathcal{A}_1$. We want to show $C(\mathcal{A}_i) = C(\mathcal{A}_j)$ for all $i, j \leq n$. Clearly, it will suffice to show that for all $i, j \leq n$, $gr(\mathcal{A}_i) = gr(\mathcal{A}_j)$, for which we need only show that for all $i \leq n$, $gr(\mathcal{A}_i) \subseteq gr(\mathcal{A}_{i+1})$. Let $\mathcal{M} \epsilon gr(\mathcal{A}_i)$. Since $\mathcal{A}_i \mid\!\sim \mathcal{A}_{i+1}$ we have $\mathcal{A}_{i+1} \subseteq \mathcal{M}$ and also $\mathcal{M}$ satisfies Σ. So we need only show that $\mathcal{M}$ is grounded in $\mathcal{A}_{i+1}$. But for this we need only go backwards in the circle from $\mathcal{A}_i$ through $\mathcal{A}_1$ to $\mathcal{A}_{i+1}$ applying lemma 3.5 at each step. □

Corollary 3.8 *Let Σ be any set of justification rules. Then the inference relation $\mid\!\sim\; := \;\mid\!\sim_\Sigma$ given by definition 3.5 (restricted to finite sets of indexed propositions on the left) is determined by some stoppered preferential model structure whose preference relation is both reflexive and transitive.*

Proof: The terminology is taken from Makinson [9]. It was shown by Kraus, Lehmann and Magidor [8] that every inference relation satisfying inclusion, cumulativity, and loop is so determined. □

4 A Shortcoming

We have gained cumulativity, but there is a cost, which appears difficult to avoid. As it stands, the above indexed version of JTMS fails to give an adequate treatment of the phenomenon of "floating conclusions". We give a simple example : Consider the set Σ of justification rules consisting of the following: $a \leftarrow \neg b$; $b \leftarrow \neg a$; $c \leftarrow a$; $c \leftarrow b$.

Under ordinary JTMS the empty set A has just two grounded models under Σ. One is $\{a, c\}$ and the other is $\{b, c\}$. The proposition c is in both of them, so we have that for ordinary JTMS, $A \mid\!\sim_\Sigma c$. Yet it is clear that the "derivation" of c is different in the two grounded models: in one it comes via a with outsupport b, and in the other it comes via b with outsupport a.

This has repercussions when we reconsider the example in indexed terms. Again we get two grounded models of the empty set. One of them is $\{(a,\{b\}),\ (c,\{b\})\}$ in which the outsupport b serves as index for a and c; the other one is $\{(b,\{a\}),\ (c,\{a\})\}$ in which the outsupport a indexes b and c. So we have $(c,\{b\})$ in one grounded model and $(c,\{a\})$ in the other, but there is no indexed proposition (c,X) that is common to both, so that under skeptical indexed JTMS we do not have $\mathcal{A} \mid\!\sim_\Sigma (c,X)$ for any X, i.e. no conclusion whatsoever can be drawn about c. This falls short of the conclusion which, as we have just seen, is provided by unindexed JTMS. Moreover, it is rather counterintuitive, for we would tend to say that the empty set of premises does indeed yield c, although with variable outsupports.

Evidently the same example arises for the cumulative version of default reasoning as developed in Brewka [3]. It too fails to handle floating conclusions.

Our example of a floating conclusion uses an unstratified set Σ of justification rules. In fact, there are no stratified examples. For it is immediate from the notion of a floating conclusion that if Σ is partially functional for indexed grounded models (in the sense that each set $\mathcal{A}$ of indexed propositions has at most one grounded model under Σ; cf section 2) then we can never have floating conclusions under Σ. And it is also possible to show, in the indexed case as in the unindexed one, that if Σ is stratified then it is partially functional for indexed grounded models.

A word of warning should however be made here. In contrast to the unindexed case, stratified sets Σ of justification rules are not always *fully* functional for indexed grounded models. Consider the example where Σ consists of the single rule $a \leftarrow$ and $\mathcal{A}$ consists of the single indexed proposition $(b,\{a\})$. Then $\mathcal{A}$ has no indexed models under Σ, grounded or ungrounded. For if $\mathcal{M}$ is a model of $\mathcal{A}$ then $\mathcal{A} \subseteq \mathcal{M}$ so $(b,\{a\})\epsilon\mathcal{M}$ so by definition 3.1, $a \notin prop(\mathcal{M})$. On the other hand, since $\mathcal{M}$ is closed under Σ we have by definition 3.1 again that $(a,\emptyset)\epsilon\mathcal{M}$ and thus $a\epsilon prop(\mathcal{M})$ giving us a contradiction.

If in this example we replace the rule $a \leftarrow$ by the rule $a \leftarrow \neg c$ with non-empty body, then a similar argument shows that although $\mathcal{A}$ now has an indexed model under Σ, it still does not have any grounded indexed models.

These examples should not be surprising, for after all the indices of an indexed proposition (x,X) in the premise set $\mathcal{A}$ are meant to record the outsupports of an *already accomplished* potential proof of x. Thus intuitively a stratification should assign them a rank less than that of x, whilst the latter should be no greater than the least rank of any head of any rule in Σ.

It thus seems that in the indexed case an adequate notion of stratification should take as its objects not merely sets Σ of justification rules but rather pairs $(\Sigma,\mathcal{A})$ where $\mathcal{A}$ is a set of indexed propositions. However we shall not attempt to develop here such an extended theory of stratification, as it is not necessary for our purposes.

Is there a way of refining the indexed JTMS so as to handle floating conclusions satisfactorily, without losing cumulativity again ? What happens, for example, if we redefine $\mathcal{A} \mid\!\sim_\Sigma (x,X)$ to hold (in the principal case that $gr_\Sigma(\mathcal{A})$ is not empty) iff for every grounded model $\mathcal{M}$ of $\mathcal{A}$, there is an index Y with $(x,Y)\epsilon\mathcal{M}$, and $X = \bigcap\{Y : (x,Y)\epsilon\mathcal{M}$ for some $\mathcal{M}\epsilon gr_\Sigma(\mathcal{A})\}$? Or, alternatively, that $X = \bigcup\{Y : (x,Y)\epsilon\mathcal{M}$ for some $\mathcal{M}\epsilon gr_\Sigma(\mathcal{A})\}$?

It appears to be difficult, if not impossible, to achieve the desired result by any such adjustment of the definition of the inference relation $\mid\!\sim_\Sigma$ in the context of the given system of indices. To see this, consider a slightly more complex version of the above example, putting Σ' as follows:

$$a \leftarrow \neg b;\ b \leftarrow \neg a;\ c \leftarrow a, \neg b;\ c \leftarrow b, \neg a;\ d \leftarrow a, \neg b;\ d \leftarrow b, \neg a;$$
$$a \leftarrow c, \neg d;\ a \leftarrow d, \neg c;\ b \leftarrow c, \neg d;\ b \leftarrow d, \neg c;\ f \leftarrow \neg e;$$

Here too we get two grounded models of the empty set, one containing among other things $(c, \{a\})$ and the other containing $(c, \{b\})$, but with no (c, X) common to both. Suppose we want to come into line with intuition by treating c as an acceptable conclusion from $\mathcal{A} = \emptyset$ under Σ'. As we are working with indexed propositions, c must be entered with an index, and the only two natural candidates for the index are $\emptyset$ and $\{a, b\}$. The empty set would be the index, for example, under the first idea mentioned above for refining the definition of $\mid\!\sim$, whereas $\{a, b\}$ would be the index under the second one. But each of these options leads to trouble.

Suppose, first, that we consider $(c, \{a, b\})$ as a conclusion of $\mathcal{A} = \emptyset$ under Σ', i.e. we define $\mid\!\sim$ in such a way that we have $\mathcal{A} \mid\!\sim (c, \{a, b\})$. Now $\mathcal{A}' = \{(c, \{a, b\})\}$ has no models, grounded or otherwise, under Σ', for in a model $\mathcal{M}$ we must have $index(\mathcal{M}) \cap prop(\mathcal{M})$ empty, so if $\mathcal{A}' \subseteq \mathcal{M}$ we must have $a, b \notin prop(\mathcal{M})$ so that $\mathcal{M}$ cannot satisfy either of the rules $b \leftarrow \neg a$, $a \leftarrow \neg b$. Since $\mathcal{A}'$ has no models under Σ', $\mathcal{A}' \mid\!\sim (x, X)$ iff $(x, X) \epsilon \mathcal{A}'$. Thus $\mathcal{A}' \not\mid\!\sim (f, \{e\})$ although clearly $(f, \{e\})$ is in every grounded model of $\mathcal{A}$ so that $\mathcal{A} \mid\!\sim (f, \{e\})$ and cumulativity thus fails.

Suppose, on the other hand, that we consider $(c, \emptyset)$ as an acceptable conclusion from $\mathcal{A}$. Then by symmetry we should also consider $(d, \emptyset)$ as an acceptable conclusion from $\mathcal{A}$. Put $\mathcal{A}" = \{(c, \emptyset)\}$. This set has three grounded models under Σ, namely:

$$\{(c, \emptyset), (a, \{b\}), (c, \{b\}), (d, \{b\}), (f, \{e\})\},$$
$$\{(c, \emptyset), (b, \{a\}), (c, \{a\}), (d, \{a\}), (f, \{e\})\},$$
$$\{(c, \emptyset), (a, \{d\}), (b, \{d\}), (f, \{e\})\}.$$

Note that the third grounded model of $\mathcal{A}"$ contains nothing of the form (d, X), so that $\mathcal{A}" \not\mid\!\sim (d, \emptyset)$. Thus again cumulativity fails.

This example illustrates the difficulty of handling floating conclusions adequately without losing cumulativity again, on a strategy of making minor adjustments to the definition of $\mid\!\sim$ whilst leaving unchanged the notion of a grounded indexed model and the general system of indexing. The authors are currently investigating the possibility of obtaining the desired properties by means of more radical reconstructions.

References

[1] Apt, K., H.Blair and A.Walker: 1988, "Towards a theory of declarative knowledge", in J.Minker ed., Foundations of Deductive Databases (Los Altos: Morgan Kaufmann) 89-148.

[2] Brewka, G.: 1989, Nonmonotonic Reasoning - From Theoretical Foundation Towards Efficient Computation, Ph.D thesis, University of Hamburg.

[3] Brewka, G.: to appear, "Cumulative default logic - in defense of nonmonotonic inference rules", Artificial Intelligence 50 (1991)

[4] Dix, J.: "Cumulativity in the semantics of logic programs", this volume.

[5] Doyle, J.: 1979, "A truth maintenance system", Artificial Intelligence 12 231-272.

[6] Elkan,C.: 1990, "A rational reconstruction of nonmonotonic truth maintenance systems", Artificial Intelligence 43 219-234.

[7] Gelfond, M. and V.Lifschitz: 1988, "The stable model semantics for logic programming", in K.Bowen and R.Kowalski eds, Fifth International Conference Symposium on Logic Programming 2 (Cambridge Mass., MIT Press) 1070-1080.

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

[9] Makinson, D.: 1989, "General theory of cumulative inference", in M.Reinfrank and others eds, Non-Monotonic Reasoning (Berlin: Springer-Verlag) 1-18.

[10] Makinson, D.: to appear, "General patterns in nonmonotonic reasoning", in D.Gabbay and others eds, Handbook of Logic in Artificial Intelligence and Logic Programming, vol. II: Nonmonotonic and Uncertain Reasoning (Oxford University Press).

[11] McCarthy, J.: 1980, "Circumscription - a form of nonmonotonic reasoning", Artificial Intelligence 13 27-39.

[12] McCarthy, J.: 1984, "Applications of circumscription to formalizing common sense knowledge", Artificial Intelligence 28 89-118.

[13] Moore, R.: 1985, "Semantical considerations on nonmonotonic logic", Artificial Intelligence 25 75-94.

[14] Reinfrank, M., O.Dressler and G.Brewka: 1989, "On the relation between truth maintenance and autoepistemic logic", Proceedings of the IJCAI-89, 1206-1212.

[15] Reiter, R.: 1980, "A logic for default reasoning", Artificial Intelligence 13 81-132.