Comparing Consequence Relations

Peter A. Flach
Dept. of Computer Science
University of Bristol
Bristol BS8 1UB, UK
Peter.Flach@cs.bris.ac.uk
http://www.cs.bris.ac.uk/~flach/

Abstract

The technical problem addressed in this paper is, given two rule systems for consequence relations **X** and **Y**, how to construct **Y**-approximations of a given **X**-relation. While an upper **Y**-approximation can be easily constructed if all **Y**-rules are Horn, the construction of lower **Y**-approximations is less straightforward. We address the problem by defining the notion of *co-closure under co-Horn rules*, that can be used to remedy violation of certain rules by *removing* arguments. In particular, we show how the co-closure under Monotonicity can be used to construct the monotonic *restriction* of a preferential relation. Unlike the more usual closure under the rules of **M**, this co-closure operator supports the intuition that preferential reasoning is more liberal than monotonic reasoning. The approach is embedded in a general framework for comparing rule systems for consequence relations. A salient feature of this framework is that it is also possible to compare rule systems that are not related by metalevel entailment.

1. INTRODUCTION

1.1 MOTIVATION AND SCOPE

Nonmonotonic reasoning is the process of 'tentatively inferring from given information rather more than is deductively implied' (Makinson, 1994). Nonmonotonic reasoning can thus be said to be more *liberal* than monotonic reasoning. Correspondingly, the set of arguments accepted by a nonmonotonic reasoning agent (also called a *consequence relation*, and defined as a subset of $L \times L$, where L is the language) can be divided into a deductive or monotonic part and a nonmonotonic part. Let us call the function which maps an arbitrary consequence relation to its monotonic core the *monotonic restriction*.

Although the notion of monotonic core has been considered before (Stachniak, 1993), it does not seem to occupy a central place in the theory of nonmonotonic consequence relations, and operators to construct the monotonic core of a given relation have not been defined before.[1] Kraus *et al*. (1990) define a monotonic closure operator, which however maps a consequence relation to a monotonic *superset* (and may therefore be called the *monotonic extension*). The operator seems to be inspired by the Horn form of the rules they consider. However, as we show in this paper even with Horn rules it is possible to apply them in the reverse direction to remove arguments from the consequence relation.

Another aspect we clarify in this paper is the role of metalevel entailment between rule systems. For instance, we have that all the rules of **P** are rules of **M**, hence all monotonic consequence relations are preferential. In our view this is a special case of a more general phenomenon, namely that **P**-semantics encodes more information than **M**-semantics, because it has to distinguish more consequence relations. However, the presence of metalevel entailment does not, by itself, indicate whether this extra information is used to establish a more liberal or rather a less liberal form of reasoning.

Moreover, metalevel entailment is not even a necessary condition for one rule system to be more liberal than another. This will be demonstrated by defining a variant of **P** that is incomparable to it wrt. metalevel entailment (each system includes a rule that is not a rule of the other system), yet clearly and unambiguously axiomatises a less liberal form of reasoning than **P**. In fact, failure to relate these rule systems by an existing comparison criterion was the original motivation for this paper.

1.2 AN EXAMPLE

Consider two reasoning agents NM and CNM, which differ only in the way they handle contradictory information: while NM infers everything from contradictory premisses, CNM refuses to draw any conclusions from them. For all other premisses they agree on the consequences. It follows that the set of CNM-

[1] I.e. operators that work directly on the consequence relation (rather than on its semantic characterisation).

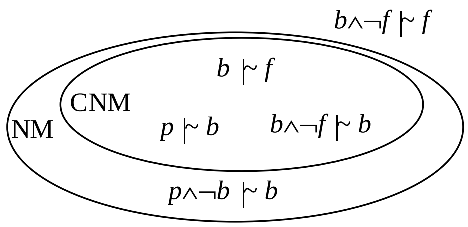

Figure 1. NM is a more liberal reasoner than CNM.

arguments is a subset of the set of NM-arguments (Figure 1). For instance, both NM and CNM infer b from p, but while NM infers anything (including b) from $p\wedge\neg b$, CNM considers those premisses to have no consequences.

Notice that NM and CNM can predict each other's behaviour and hence, in a sense, employ the same information in their reasoning. Specifically, CNM can reconstruct X's behaviour by the rule 'if I don't infer anything from given premisses, NM will infer everything from them; if on the other hand I do infer some consequences, NM will infer exactly the same'. In other words, CNM drops *conclusions* without dropping *information*.

As indicated in Figure 1 NM does not conclude f from $b\wedge\neg f$, i.e. NM considers $b\wedge\neg f$ to be contradictory. Since NM does conclude f from b it follows that NM is a nonmonotonic reasoner. Now consider two other reasoning agents M1 and M2, neither of which accepts an inference from α to β without treating $\alpha\wedge\neg\beta$ as contradictory premisses (from which they, like NM, infer everything). This means that neither M1 nor M2 can reason exactly like NM: if they want to keep the inference from b to f they should, unlike NM, consider $b\wedge\neg f$ contradictory, while if they follow NM in not considering $b\wedge\neg f$ contradictory they should drop the inference from b to f. As it turns out, M1 takes the first option and hence infers everything from $b\wedge\neg f$, while M2 drops the inference from b to f (Figure 2).

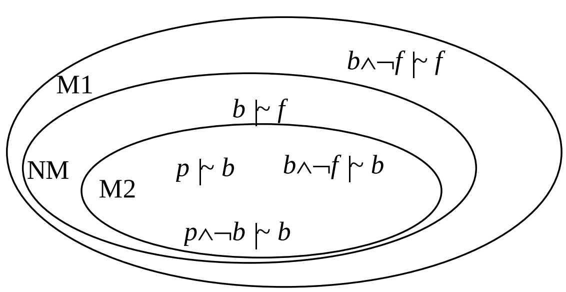

Figure 2. Upper and lower approximations of NM.

Clearly, M1 is strictly more liberal than NM and M2 is strictly less liberal than NM. Furthermore, NM is perfectly able to predict the behaviour of both M1 and M2, but neither M1 nor M2 can exactly reproduce NM's behaviour. M1 cannot predict NM, because NM treats the arguments 'from b infer f' and 'from p infer b' differently, while M1 treats them in the same way. Although M2 and NM agree on what they consider contradictory, M2 does not know for what non-contradictory $\alpha\wedge\neg\beta$ NM accepts the argument 'from α infer β'. In both cases, information has been dropped that cannot be reconstructed. Notice that — compared to NM — M1 drops information to infer *more* conclusions, while M2 drops information to infer *less* conclusions. Also, note that M1 and M2 cannot reproduce each other's behaviour.

In these examples NM embodies the prototypical nonmonotonic reasoner, who is willing to infer f from b by default, at the same time accepting $b\wedge\neg f$ as an exceptional but not contradictory circumstance (the reader may want to read 'it is a bird' for b, 'it flies' for f, and 'it is a penguin' for p — note that the inference from p to b is treated as a deductive inference by all reasoners). In contrast, M1 and M2 are classical monotonic reasoners, who are unable to deal with such default inferences: they either accept the exception $b\wedge\neg f$ as being non-contradictory and drop the default inference (M2), or else reconstruct the default inference as a deductive inference, turning the exception into a contradiction (M1).

It is easy enough to define a closure operator constructing M1 from NM. In this paper we define a co-closure operator constructing M2 from NM. As M2 represents the monotonic core of NM, this operator stays close to the intuition that NM 'jumps to conclusions'. We will also explain why CNM may be considered a more conservative form of reasoning than NM, even though there is no closure operator to map NM to CNM or *vice versa*.

1.3 APPROACH

In this paper we will address the issues mentioned above by introducing a concept of *reduction* that is similar to its counterpart in computational complexity theory. If **X** and **Y** are rule systems, we define a reduction of **X** to **Y** as a function f mapping consequence relations to consequence relations, such that x satisfies the rules of **X** iff $f(x)$ satisfies the rules of **Y**. A reduction establishes a correspondence between **X**-reasoners and **Y**-reasoners, such that any **X**-reasoner can predict the behaviour of the corresponding **Y**-reasoner. This correspondence then establishes a relation between **X** and **Y**; for instance, it may map any **X**-relation to a **Y**-relation that is a subset or superset. It can also be used to investigate the relation between rule systems that are incomparable by metalevel entailment.

The rest of the paper is organised as follows. The formal preliminaries are given in Section 2. Section 3 introduces reductions, and the derived notions of extension and restriction, and applies these to various rule systems. In Section 4 we discuss the main implications of this work.

2. PRELIMINARIES

The formal background of this paper is rooted in the work on abstract consequence relations that are axiomatised by metalevel rules (Gabbay, 1985; Makinson, 1989; Kraus *et al*., 1990). Kraus, Lehmann and Magidor have characterised several sets of such metalevel rules in their seminal paper (Kraus *et al*., 1990), the most important of which are **M** for monotonic or deductive reasoning, **P** for preferential reasoning, and **C** for cumulative reasoning. These rule systems are related by metalevel entailment: an axiomatisation of **P** is obtained by adding the rule of Or to **C**, and therefore all rules of **C** are entailed by **P** (see Definition 1 below). Similarly, **M** is axiomatised b[redacted]**P** augmented with the rule of Monotonicity.

2.1 THE [redacted]ETALANGUAGE

In this se[redacted]on we define the metalanguage u[redacted] formulating[redacted]rule systems. We mostly foll[redacted]he terminolog[redacted]and notation of (Kraus *et al*., 1990);[redacted]familiar wi[redacted]o this sect[redacted]

Throughou[redacted]he paper L is a propositional language[redacted]a countable s[redacted] of proposition symbols, closed un[redacted] usual logic[redacted]more giv[redacted]t of propositi[redacted]al models U, and a satisfaction relati[redacted] $U\times L$ that [redacted] well-behaved with respect to the [redacted] connectives[redacted]and com[redacted]act. As usual, we write [redacted]r $\forall m\in U$: m[redacted], for ar[redacted]trary $\alpha\in L$. Note that U m[redacted]a proper sub[redacted]t of th[redacted]set of all truth-assignments[redacted]proposition symbols[redacted]n L, which would reflect p[redacted]r knowledge or *backg*[redacted]*und knowledge* of the reason[redacted]g agent. Equivalently, [redacted] set[redacted]f models of an implici[redacted]background theory T, and let [redacted] stand for 'α is a logic[redacted]consequence of T'[redacted]

2.1.1 Syntax

The metalanguage fo[redacted]reasoning about c[redacted]sequence rel[redacted]tions is a restricted [redacted]edicate language b[redacted] unary metapredicate [redacted] in prefix notatio[redacted] validity with respect [redacted]U in L) and a bina[redacted] metapredica[redacted] $\mid\!\sim$ in infix notation (standing for an unspe[redacted]ied relation [redacted] consequence). In referring to object-level f[redacted] we employ a countable set of metavariab[redacted]s α, β, γ, [redacted] …, and the logical connectives from L [redacted]t as functio[redacted] symbols on the metalevel. Metalevel lit[redacted]als are atom[redacted] formulae or their negation; instead of ¬([redacted]) we write [redacted] α, and instead of $\neg(\alpha \mid\!\sim \beta)$ we write α [redacted]. Formulae [redacted] the metalanguage, often referred to as *rules* or *properties*, are of the form $P_1,\ldots,P_n$ / Q for $n\geq 0$ (usually written in an expanded Gentzen-style notation), where $P_1,\ldots,P_n$ and Q are literals. Intuitively, such a rule should be interpreted as an implication with *antecedent* $P_1,\ldots,P_n$ (interpreted conjunctively) and *consequent* Q, in which all variables are implicitly universally quantified. A *rule system* is a set of such metalevel rules, denoted by abbreviations in boldface capitals.

2.1.2 Semantics

Consequence relations provide the semantics for this metalanguage, by fixing the meaning of the metapredicate $\mid\!\sim$. Formally, a consequence relation is a subset of $L\times L$. They will be used to model part or all of the reasoning behaviour of a partic[redacted]r reasoning agent, by listing a number of *arguments*[redacted]pairs of *premiss* and *conclusion*) the agent is prepared [redacted] accept. A consequence relation satisfies a rule whene[redacted] it satisfies all instances of the rule, and violates it oth[redacted]wise, where an instance of a rule is obtained by replac[redacted] with formulae from L. A [redacted]onsequence relation satisfies an instance of a rule if, [redacted]whenever it satisfies the ground literals in the antecede[redacted]es the consequent. A conse[redacted]nce relation satisfies a negated ground literal if it doe[redacted]not satisfy the unnegated ground literal. Finally:

- [redacted] is satisfied whenever the propositional for[redacted]la from L denoted by α is [redacted]model in U;
- [redacted]$\alpha \mid\!\sim \beta$ is satisfied whenever the pair of propositional formulae from L denoted by [redacted] an element of the consequence relation.

[redacted] ignore the distinction between the metalanguage and its semantics by referring to a particular consequence relation as $\mid\!\sim$ and writing $p \mid\!\sim q$ instead of [redacted]ule system, a consequence relation satisfying the rules of **X** is called an **X**-relation. Rule system **X** *entails* rule system **Y** if every **X**-relation is a **Y**-relation.

2.2 RULE SYSTEMS

In this section we introduce the rule systems considered in this paper.

[redacted] C, P and [redacted]

Among the rule systems studied by [redacted]aus *et al*. (1990) are the following.

[redacted] (Rule systems [redacted] rule system **C** (for cumu[redacted] consists of the following rules:

Reflexivity: [redacted]

Left Logical Equivalence[redacted]$\leftrightarrow\beta\,,\ \alpha \mid\!\sim \gamma$ / $\beta \mid\!\sim \gamma$

Right Weakening: [redacted]$\rightarrow\beta\,,\ \gamma \mid\!\sim \alpha$ / $\gamma \mid\!\sim \beta$

Cut: $$\frac{\alpha \mid\!\sim \beta\,,\ \alpha\wedge\beta \mid\!\sim \gamma}{\alpha \mid\!\sim \gamma}$$

Cautious Monotonicity: $\dfrac{\alpha \mid\!\sim \beta \, , \, \alpha \mid\!\sim \gamma}{\alpha\wedge\beta \mid\!\sim \gamma}$

The rule system **P** (for prefe[redacted]ntial reasoning) consists of the rules of **C** pl[redacted] the following rule:

Or: [redacted] $\mid\!\sim \gamma \, , \, \beta \mid\!\sim \gamma$ / $\alpha\vee\beta \mid\!\sim \gamma$

The rule system **M** (for mon[redacted]) consists of the rules of **P** pl[redacted] the following rule:

Monotonicity: [redacted] $\rightarrow\beta \, , \, \beta \mid\!\sim \gamma$ / $\alpha \mid\!\sim \gamma$

The axiomatisations of **C**, **P** and **M** have been chosen such that they can be obtained from one another by adding or deleting rules. Consequently, **M** entails **P** and **P** entails **C**. Note that Cautious Monotonicity and Left Logical Equivalence are redundant in **M**, since the[redacted] are implied by Monotonicity.

The main result of (Kraus *et al*., 1990) [redacted] a characterisation of these rule systems in terms [redacted] the following semantics (with slight changes of termin[redacted]

DEFINITION 2 (Cumulative, pre[redacted]ential and monotonic structures). A *cumulativ*[redacted]structure is a triple $W = \langle S,l,<\rangle$, where S is a set[redacted] $S\rightarrow 2^U$ is a function that labels every[redacted]ate with a nonempty set of models, and < [redacted] a binary relation[2] on S. A state $s\in S$ satisfi[redacted] a formula $\alpha\in L$ iff for every model $m\in l(s)$, m [redacted] α; the set of states satisfying α is denoted [redacted] $[\alpha]$. The consequence relation defined by W is denoted by $\mid\!\sim_W$ and is defined by: $\alpha \mid\!\sim_W \beta$ iff every state minimal (wrt. <) in $[\alpha]$ satisfies β.
A *preferential* structure is a cumulative structure $\langle S,l,<\rangle$ where every label $l(s)$ is a singleton, and < is a strict partial order (i.e., < is irreflexive and transitive).
A *monotonic* structure is a preferential structure $\langle S,l,\emptyset\rangle$, i.e. the preference relation is empty.

The intermediate level of states allows the same model to appear at several points in the ordering.

2.2.2 The system CP

In order to capture the behaviour of the reasoning agent CNM from the introduction of this paper, who refuses to draw any conclusion from contradictory premisses, we introduce the following rule system.

DEFINITION 3 (Consistent prefer[redacted] ing). The rule system **CP** consists [redacted] the rules of **P** with the exception of Ref[redacted]xivity, and additionally the following two rules[redacted]

Consistent Reflexivity: $\alpha \mid\!\sim$ [redacted] $, \, \alpha \not\mid\!\sim \neg\beta$ / [redacted] $\mid\!\sim \beta$

Consistency: [redacted] $\rightarrow\neg\beta$ / $\alpha \not\mid\!\sim \beta$

A *consistent preferential* structure is a preferential structure $W = \langle S,l,<\rangle$. The consequence relation defined by W is denoted by $\mid\!\sim_W$ and is defined by: $\alpha \mid\!\sim_W \beta$ iff (*i*) $[\alpha]=\emptyset$, and (*ii*) every state minimal in $[\alpha]$ satisfies β.

Similar forms of reasoning have been considered in the literature before (e.g. Benferhat *et al*., 1992). The system **CP** is included here mainly for the sake of argument; however, we will briefly pause to comment on one of its possible applications.

Consistent preferential reasoning was studied in (Flach, 1995) as a model for a certain kind of induction called *confirmatory induction*, which is[redacted] form of closed-world reasoning based on the assumpti[redacted] 'objects that I haven't [redacted] have like objects I have[redacted]een'. In this form of reasoning $\alpha \mid\!\sim \beta$ is interpreted as[redacted]bservations α *confirm* inductive hypothesis β', and t[redacted] means that contradictory observa[redacted]ons do not confirm any [redacted]ses. Apart from being in[redacted]itively justifiable, this enables the unification of con[redacted]matory induction with *explanatory induction*, where a [redacted] entail the observations unless the[redacted]are contradictory.

Clearly, in the presence of Consi[redacted]ency, various properties such as Supraclassicality (from [redacted] $\rightarrow\beta$ derive $\alpha \mid\!\sim \beta$) are too strong; this is remedied by re[redacted]lacing Reflexivity with the weaker rule Consistent Reflexivity. Consequently, **P** and **CP** do not entail each other. Notice that consistent preferential structures consist of the same information as preferential structures, but this information is used in a different way by the addition of condition (*i*). For a proof of the completeness of **CP** with respect to consistent preferential structures see (Flach, 1995; 1996).

2.3 CLOSURES AND CO-CLOSURES

We introduce some new terminology, drawing upon an analogy with logic programming. This analogy is revealed by viewing the formulae from the object language L as ground terms in a Herbrand universe. Consequence relations then correspond to Herbrand interpretations (restricted to the metapredicate $\mid\!\sim$) of the metalanguage, whose rules can be easily transformed to clausal notation.

---

[2] < is not necessarily a partial order, but it should satisfy a certain 'smoothness condition', which is for instance satisfied if < does not have infinite descending chains.

2.3.1 Closure under Horn rules

DEFINITION 4 (Definite rules, indefinite rules, and denials). A rule $P_1,\ldots,P_n \,/\, Q$ is called

1. *definite* if all of $P_1,\ldots,P_n$ and Q are positive literals;
2. *indefinite* if at least one of $P_1,\ldots,P_n$ is a negative literal and Q is a positive literal;
3. a *denial* if all of $P_1,\ldots,P_n$ are positive literals and Q is a negative literal.[3]

This exhausts all the possibilities: the case that at least one of $P_1,\ldots,P_n$ is a negative literal and Q is a negative literal can be rewritten to case 1 or case 2.

EXAMPLE 1. All of the above rules are definite, except the added **CP**-rules: Consistent Reflexivity is an indefinite rule, and Consistency is a denial.

As is well-known, with a set of definite rules **D** one can associate an immediate consequence operator, which maps a set of arguments A to its immediate consequences under **D**, as follows (ground$_L$(**D**) stands for the set of ground instances of rules in **D** over the Herbrand universe L):

$T_{\mathbf{D}}(A) = \{Q \mid P_1,\ldots,P_n \,/\, Q$ is a definite rule in ground$_L$(**D**) and $P_1,\ldots,P_n$ is satisfied by [redacted]

We will make use of the following proposition, [redacted] known from logic programming theory.

PROPOSITION 1 (Horn closure). *Let* ***D*** *be a set of definite rules. The intersection of any set of* ***D****-relations is also a* ***D****-relation. The smallest* ***D****-relation containing a given set of arguments is unique and equal to the intersection of all* ***D****-relations containing the given arguments, and also to the least fixpoint of the immediate consequence operator* $T_{\mathbf{D}}$*, starting from the given arguments.*

The latter construction is called the ***D****-closure* of the original set of arguments. As a denial does not produce positive consequences, Proposition 1 also holds for sets of definite rules and denials, jointly called *Horn rules*.

Although they don't use the a[redacted]ove[redacted]erminology, Kraus *et al*. define, for each rule [redacted]yste[redacted] they consider, a corresponding closure oper[redacted]or. [redacted]hus, these closure operators use the metalevel r[redacted] derive further arguments. Fo[redacted] ins[redacted]nce, the **M**-closure operator will turn a preferentia[redacted]con[redacted]quence relation into a monotonic superset. Intuitiv[redacted] from the assumption that the default rules employed by the preferential reasoner are actually without exceptions.

EXAMPLE 2. Consider Figure 2. NM violates Monotonicity because $b \mathrel{|\!\sim} f$ while $b\wedge\neg f \not\mathrel{|\!\sim} f$. The **M**-closure operator will add $b\wedge\neg f \mathrel{|\!\sim} f$ by virtue of Monotonicity. Furthermore, assuming that NM is a **P**-reasoner we already have $b\wedge\neg f \mathrel{|\!\sim} \neg f$ by Reflexivity and Right Weakening. In a next iteration the **M**-closure operator will therefore add $b\wedge\neg f \mathrel{|\!\sim} f\wedge\neg f$ by virtue of Right And (a derived rule of **M**). Finally, we obtain $b\wedge\neg f \mathrel{|\!\sim} \delta$ for all $\delta\in L$ because of Right Weakening, i.e. $b\wedge\neg f$ is contradictory. Notice that in general it is insufficient to close off under Monotonicity only (see Example 5 for a counter-example).

2.3.2 Co-closure under co-Horn rules

A less common but in the context of this paper very useful dual of the above is obtained if we consider complements of consequence relations, and view $\alpha \not\mathrel{|\!\sim} \beta$ as a 'co-positive' literal and $\alpha \mathrel{|\!\sim} \beta$ as a 'co-negative' literal.

DEFINITION 5 (Co-definite rules, co-indefinite rules, co-denials, and co-Horn rules). A rule $P_1,\ldots,P_n \,/\, Q$ is called

1. *co-definite* if exactly one of $P_1,\ldots,P_n$ is a [redacted] a positive literal;
[redacted] two of $P_1,\ldots,P_n$ are [redacted] a positive literal;
[redacted] $P_1,\ldots,P_n$ are negative [redacted]ve literal;
[redacted] co-definite or a co-

[redacted]al Equivalence, Right [redacted], Consistent Reflex- [redacted]are co-definite; Cut, [redacted]d Or are co-indefinite; [redacted]ial.

We can thus define an immediate co-consequence operator given a set of co-definite rules **CD**, which operates on a set of arguments A and computes the set of immediate co-consequences of A (arguments to be removed from A) under **CD**.

$CT_{\mathbf{CD}}(A) = \{P \mid \neg P_1,\ldots,\neg P_n, P \,/\, Q$ is a co-definite rule in ground$_L$(**CD**) and $\neg P_1,\ldots,\neg P_n$ and $\neg Q$ are satisfied by $A\}$

The following proposition is the dual of Proposition 1:

PROPOSITION 2 (Co-Horn co-closure). *Let* ***CD*** *be a set of co-definite rules. The largest* ***CD****-relation contained in a given set of arguments is*

[3] For determining whether a rule is d[redacted]nite, [redacted]definite or a denial, literals with the 'built-in' predicates [redacted]and [redacted]can be ignored.

*unique and equal to the union of all **CD**-relations contained in those arguments, and also to the complement of the least fixpoint of the immediate co-consequence operator $CT_{\mathbf{CD}}$, starting from the complement of the given arguments.*

The latter construction is called the ***CD**-co-closure* of the original set of arguments. It is the main technical tool for obtaining the results in the next section.

EXAMPLE 4. Consider again Figure 2. The co-closure of NM under Monotonicity will remove $b \mathrel{|\!\sim} f$, since $b \wedge \neg f \not\mathrel{|\!\sim} f$ is satisfied by NM.

3. COMPARING RULE SYSTEMS

We now come to the main part of the paper. Section 3.1 defines reductions between rule systems, and the conditions under which these may establish extensions or restrictions. The relations between **P** and **M** and between **P** and **CP** are studied in Sections 3.2 and 3.3.

3.1 REDUCTIONS, EXTENSIONS AND RESTRICTIONS

We want to characterise the difference in information encoded in rule systems **X** and **Y**, or equivalently in the semantics characterising them. Generally speaking, a semantics for **X** has two purposes:

1. to distinguish between different **X**-relations, and
2. to distinguish between **X**-relations and non-**X**-relations, i.e. to answer the decision problem ‘is *x* an **X**-relation?’

The idea of a reduction is to find a rule system **Y** and a mapping *f* such that this latter decision problem is equivalent to the decision problem ‘is *f*(*x*) a **Y**-relation?’ (Figure 3). We may lose the distinction between some of the **X**-relations in the process, in which case the **X**-

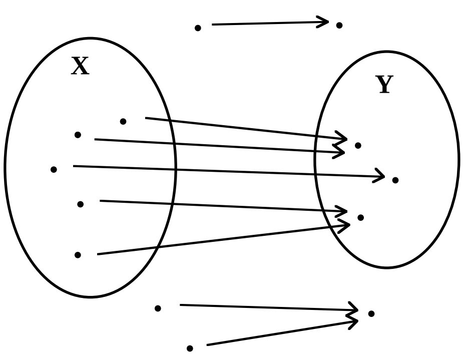

Figure 3. A reduction of **X** to **Y**.

[redacted] formation than the **Y**-[redacted]-relations are images under [redacted]s would need additional [redacted]se **Y**-relations not in the [redacted]at *f* encodes the difference [redacted]d **Y**.

DEFINITION 6 (Reduction). Given two rule systems **X** and **Y**, a *reduction* of **X** to **Y** is a function *f* mapping consequence relations to consequence relations, such that (*i*) *x* is an **X**-relation iff *f*(*x*) is a **Y**-relation; (*ii*) every **Y**-relation is the *f*-image of an **X**-relation. If such a mapping exists we say that **X** *reduces to* **Y**. If in addition **Y** reduces to **X**, we say that **X** and **Y** are *reduction-equivalent*, otherwise **X** *properly* reduces to **Y**.

Notice that the relation ‘reduces to’ is a pre-order (it is reflexive and transitive).

3.1.1 Reductions between Horn systems

The **M**-closure as defined by Kraus *et al*. is not a reduction of anything else than the empty set of rules to **M**, since it maps *any* consequence relation into a monotonic superset. In general, a reduction of **X** to **Y** must be strong enough to transform **X**-relations into **Y**-relations, but not so strong that it transforms non-**X**-relations into **Y**-relations. Clearly, the **M**-closure is too strong in this sense. There is, however, a way out by taking the *difference* between the **P**-closure and the **M**-closure.

THEOREM 3 (Horn reduction). *Let **X** and **Y** be two rule systems, such that **Y** entails **X**. If every rule of **X** and **Y** is Horn, then **X** reduces to **Y**; if in addition **X** does not entail **Y** the reduction is proper.*

Proof. If **X** and **Y** are Horn, then the closure of $\mathrel{|\!\sim}$ under **X** and **Y** is well-defined and denoted by $\mathrel{|\!\sim}_{\mathbf{X}}$ and $\mathrel{|\!\sim}_{\mathbf{Y}}$, respectively. Consider the following function:

$$f(\mathrel{|\!\sim}) = \mathrel{|\!\sim}_{\mathbf{Y}} - (\mathrel{|\!\sim}_{\mathbf{X}} - \mathrel{|\!\sim})$$

We will prove that *f* establishes a reduction of **X** to **Y**. If $\mathrel{|\!\sim}$ is an **X**-relation then $\mathrel{|\!\sim} = \mathrel{|\!\sim}_{\mathbf{X}}$ and thus $f(\mathrel{|\!\sim}) = \mathrel{|\!\sim}_{\mathbf{Y}}$. On the other hand, if $\mathrel{|\!\sim}$ violates a rule of **X** because $\alpha \not\mathrel{|\!\sim} \beta$ and $\alpha \mathrel{|\!\sim}_{\mathbf{X}} \beta$, then the inference is removed from $\mathrel{|\!\sim}_{\mathbf{Y}}$ and thus $f(\mathrel{|\!\sim})$ violates the same rule of **X**. Clearly *f* maps onto the whole of **Y** because **Y**-relations are mapped onto itself.

If **X** does not entail **Y** there are more **X**-relations than **Y**-relations, hence there is no reduction of **Y** to **X**.

Notice that both **X** and **Y** are required to be Horn — we cannot use the construction in the proof of Theorem 3 to reduce a non-Horn system to a Horn system.

3.1.2 Extensions and restrictions

Once we have established a reduction of **X** to **Y**, we may want to investigate the relation between **X**-relations and the **Y**-relations they are mapped to.

DEFINITION 7 (Extension and restriction). Given a reduction of **X** to **Y**, its restriction to the set of **X**-relations is called a *semi-reduction* of **X** to **Y**. A semi-reduction is an *extension* (*restriction*) if it maps every consequence relation to a superset (subset); we say that **X** *extends to* (*restricts to*) **Y**. The extension (restriction) is *proper* if in addition **X** properly reduces to **Y**.

The relations 'extends to' and 'restricts to' are partial orders, but **X** may both extend and restrict to **Y** (we will see below that this is the case for **P** and **M**).

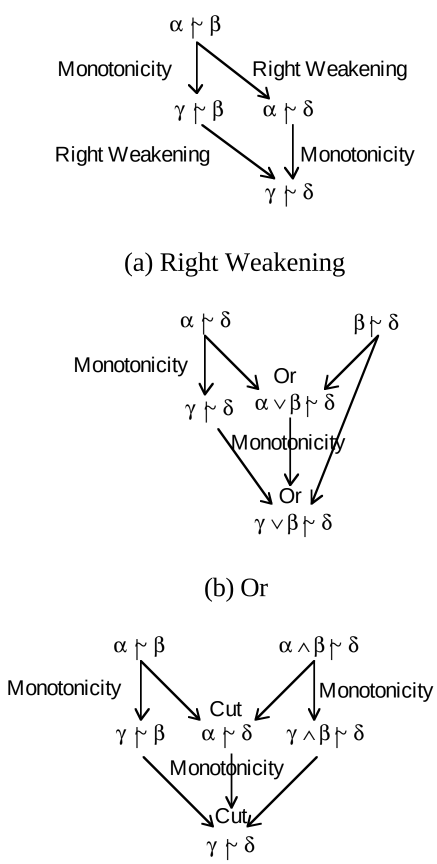

(a) Right Weakening

(b) Or

(c) Cut

Figure 4. Confluence of Monotonicity with rules of **P**.

COROLLARY 4. ***C*** *properly extends to* ***P****, and* ***P*** *properly extends to* ***M****.*
Proof. We can use Kraus *et al.*'s **P**-closure as an extension of **C** to **P**, and their **M**-closure as an extension of **P** to **M**.

As we have argued before, this closure approach establishes a relation between **P** and **M** which is intuitively unsatisfactory because it deems preferential reasoning more conservative than deductive reasoning. In the next section we define a reduction of **P** to **M** that is intuitively more appealing.

3.2 COMPARING P AND M

It is straightforward to obtain a dual to Theorem 3 which relates co-Horn rule systems by means of their co-closures. Such a result would however have limited practical importance, since none of the rule systems considered in this paper are co-Horn. However, note that Monotonicity is a co-definite rule; we will show that the co-closure under Monotonicity yields a restriction of **P** to **M**, without further help of the rules of **P**.

3.2.[redacted] The restriction of P to M

The [redacted]llowing Lemma provides the key insight.

[redacted]MMA 5. *The co-closure under Mon*[redacted]*onicity of* [redacted]
[redacted]*oof*. Clearly the co-closure under M[redacted]otonicity [redacted] any consequence relation is a [redacted]bset that [redacted] that the [redacted] not [redacted]olations of rules of **P**.
[redacted]r Reflexivity, $\alpha \mathrel{|\!\sim} \alpha$ would be [redacted]moved if [redacted]ne γ, [redacted]possible if $\mathrel{|\!\sim}$ is preferential (use [redacted]eflexivity [redacted]d Right Weakening).
[redacted]r Right Weakening, suppose α[redacted]$\sim \beta$ and [redacted]$\rightarrow\delta$. $\alpha \mathrel{|\!\sim} \delta$ would be removed if [redacted]$\rightarrow\alpha$ and [redacted]$\not\mathrel{|\!\sim} \delta$ for some γ, but then we woul[redacted]also have $\gamma \not\mathrel{|\!\sim} \beta$ by Right Weakening, hence $\alpha \mathrel{|\!\sim} \beta$ would be removed, preventing the violation of Right Weakening.
For Or and Cut an analogous argument holds (see Figure 4).
Finally, Left Logical Equivalence and Cautious Monotonicity are implied by Monotonicity.

It should be noted that the dual of Lemma 5 does not hold: if we would close off a preferential relation under Monotonicity only, the resulting relation may violate some rule of **P**. Monotonicity by itself does not fully characterise the difference between a **P**-relation and its **M**-extension.

EXAMPLE 5. C[redacted]sider the preferential structure with states $s<t<$[redacted]$\neg$c and $\neg$d, t satis[redacted]es a, $\neg$b, c and $\neg$d, and u satisfies a, b, c[redacted]nd d. We thus have a $|\!\sim$ b and c$\wedge$b $|\!\sim$ d, but [redacted] $|\!\not\sim$ d, c $|\!\not\sim$ b, and c $|\!\not\sim$ d. Suppose now [redacted]$\rightarrow$a, then closing off under Monotonicity a[redacted]s c $|\!\sim$ b but not c $|\!\sim$ d. The resulting relation violates Cut and is therefore not an **M**-relation.

We will now show that there is a reduction of **P** to **M** of which the co-closure under Monotonicity establishes the semi-reduction.

THEOREM 6. ***P** properly restricts to **M**.*
Proof. Let $|\!\sim$ be an arbitrary consequence relation, let $|\!\sim_{\mathbf{P}}$ denote its **P**-closure, and let $|\!\sim_{\mathbf{M}'}$ denote the co-closure under Monotonicity of $|\!\sim$. Consider the function g defined as follows:

$$g(|\!\sim) = \begin{cases} |\!\sim_{\mathbf{M}'} & \text{if } |\!\sim_{\mathbf{P}} = |\!\sim \\ |\!\sim & \text{otherwise} \end{cases}$$

We will prove that g establishes a reduction of **P** to **M**. If $|\!\sim$ satisfies **P** then $|\!\sim = |\!\sim_{\mathbf{P}}$ and therefore $g(|\!\sim) = |\!\sim_{\mathbf{M}'}$ which satisfies **M** by Lemma 5. On the other hand, if $|\!\sim$ violates a rule of **P** then $|\!\sim_{\mathbf{P}} \neq |\!\sim$, hence $g(|\!\sim) = |\!\sim$ violates **M**.
Since **P** does not entail **M** there are more **P**-relations than **M**-relations, hence there is no reduction of **M** to **P**.
Finally, we have that **P**-relations are mapped to subsets, which means that the semi-reduction is a restriction.

The reduction g in the proof of Theorem 6 is admittedly not very elegant — note however that $g(|\!\sim) = |\!\sim_{\mathbf{M}'}$ — $(|\!\sim_{\mathbf{P}}$ — $|\!\sim)$ doesn't work because co-closure under Monotonicity may remove violations of **P**. In any case, the importance of this result is that we have obtained an alternative way of relating **P** and **M**, by defining the **M**-restriction of a **P**-relation as its co-closure under Monotonicity.

3.2.2 Semantic characterisations

For completeness we also give semantic characterisations of the above semi-reductions of **P** to **M**. The **M**-extension of a **P**-relation is obtained by throwing out every state which represents an exception to a preferential argument (the preference order becoming obsolete in the process).[4]

[4] Similar results have been obtained by (Stachniak, 1993; Benferhat *et al.*, 1996).

THEOREM 7 (Semantic characterisation of extension of **P** to **M**). *Let $|\!\sim$ be a preferential relation characterised by the preferential structure $\langle S,l,<\rangle$, and let $|\!\sim'$ be characterised by the monotonic structure $\langle S',l,\emptyset\rangle$ with*

$$S' = S - \{s\in S \mid s \text{ satisfies } \alpha\wedge\neg\beta \text{ for some } \alpha \,|\!\sim\, \beta\}$$

*$|\!\sim'$ is the **M**-extension of $|\!\sim$.*
Proof. For every argument $\alpha \,|\!\sim\, \beta$ we have that every state in S' satisfies $\alpha\rightarrow\beta$, hence $|\!\sim'$ is a superset of $|\!\sim$. Since $|\!\sim'$ is monotonic and $|\!\sim_{\mathbf{M}}$ is the smallest monotonic superset of $|\!\sim$, we have $|\!\sim' \supseteq |\!\sim_{\mathbf{M}}$; we will prove that $|\!\sim' \subseteq |\!\sim_{\mathbf{M}}$.
Suppose therefore $\alpha \,|\!\sim'\, \beta$; we will prove that $\alpha \,|\!\sim_{\mathbf{M}}\, \beta$. If $\alpha \,|\!\sim\, \beta$ then clearly $\alpha \,|\!\sim_{\mathbf{M}}\, \beta$; so suppose $\alpha \,|\!\not\sim\, \beta$, i.e. there exist states in S satisfying $\alpha\wedge\neg\beta$. Since $\alpha \,|\!\sim'\, \beta$ all such states have been removed from S — that is, for every state $s\in S$ satisfying $\alpha\wedge\neg\beta$ there are $\delta,\varepsilon\in L$ such that $\delta \,|\!\sim\, \varepsilon$ and s satisfies $\delta\wedge\neg\varepsilon$. Let Δ denote the [redacted] these $\delta,\varepsilon\in L$, then we [redacted]y a valid **P**-derivation [redacted]nd we have $\delta \,|\!\sim\, \varepsilon$ for [redacted]erefore **true** $|\!\sim$ $\delta\rightarrow\varepsilon$ [redacted] we have $\alpha \,|\!\sim_{\mathbf{M}}\, \Delta\rightarrow\beta$ [redacted] β.

[redacted]on is obtained by ignoring [redacted] characterisation of [redacted]et $|\!\sim$ be a preferential [redacted]e preferential structure [redacted] characterised by the monotonic [redacted]ucture $\langle S,l,\emptyset\rangle$. $|\!\sim'$ is the **M**-restriction of[redacted]
Proof. As be[redacted]re we denote the **M**-restriction of $|\!\sim$ by $|\!\sim_{\mathbf{M}'}$. Cl[redacted]rly $|\!\sim'$ is a subset of $|\!\sim$. Since $|\!\sim'$ is monotoni[redacted]otonic subset of $|\!\sim$, [redacted] prove that $|\!\sim' \supseteq |\!\sim_{\mathbf{M}'}$[redacted]
Suppose the[redacted]fore that $\alpha \,|\!\not\sim'\, \beta$; we will prove $\alpha \,|\!\not\sim_{\mathbf{M}'}\, \beta$. If [redacted] β; so suppose α $|\!\sim$[redacted] Since $\alpha \,|\!\not\sim'\, \beta$ there exists a state in S (non-m[redacted]imal in $[\alpha]$) satisfying $\alpha\wedge\neg\beta$. It follows that [redacted]$\wedge\neg\beta \,|\!\not\sim\, \beta$, hence $\gamma \,|\!\not\sim_{\mathbf{M}'}\, \beta$ for every γ such that [redacted]$\wedge\neg\beta\rightarrow\gamma$ due to the construction of the co-c[redacted]sure under Monotonicity — in particular $\alpha \,|\!\not\sim_{\mathbf{M}'}\, \beta$.

As a general conclusion we may say that the relation between **P** and **M** is ambiguous (at least on purely formal grounds), since **P** both extends and restricts to **M**. Our intuition that **P** establishes a logic of 'jumping to conclusions' must therefore be rooted in pragmatics. We will return to the issue in Section 4 below.

3.3 COMPARING P AND CP

The relation between **P** and **CP** is of interest, because neither of these rule systems entails the other. We show that they are still comparable within our framework.

THEOREM 9. ***P** restricts to **CP**, and **CP** extends to **P**.*

Proof. A bijection between the set of **P**-relations and the set of **CP**-relations is established by the fact that their semantic structures take the same form. So let $\langle S,l,<\rangle$ be a (consistent) preferential structure defining the **P**-relation $|\!\sim$ and the **CP**-relation $|\!\sim'$. These two consequence relations only differ in arguments with premisses α that are unsatisfiable in S: such premisses are uniquely defined by $\alpha \;|\!\sim\; \neg\alpha$ or alternatively $\alpha \;\not|\!\sim'\; \alpha$. We can then define the following functions:

$$h_1(\,|\!\sim\,) = |\!\sim \;—\; \{\,\langle\alpha,\beta\rangle \mid \alpha,\beta\in L \text{ and } \alpha \;|\!\sim\; \neg\alpha\}$$

$$h_2(\,|\!\sim'\,) = |\!\sim' \cup\; \{\,\langle\alpha,\beta\rangle \mid \alpha,\beta\in L \text{ and } \alpha \;\not|\!\sim'\; \alpha\}$$

It is easy to show that these functions define reductions from **P** to **CP** and from **CP** to **P**, respectively. The corresponding semi-reductions establish a restriction of **P** to **CP** and and extension of **CP** to **P**, respectively.

This result unequivocally establishes **P** as a more liberal form of reasoning than **CP**.

4. DISCUSSION

In this paper we have proposed the notion of reducibility between rule systems in order to characterise their difference. A reduction of **X** to **Y**, if it exists, shows that **X**-semantics has more degrees of freedom than **Y**-semantics. It also constructs a '**Y**-approximation' for a given **X**-relation. We have demonstrated that this notion is more general then metalevel entailment or closure operators by applying it to rule systems that don't entail each other.

In our framework the relation between **M** and **P** is inherently ambiguous: by throwing away exceptions to defaults we construct an **M**-extension of a preferential relation, by throwing away the defaults themselves we construct an **M**-restriction. While the latter reduction is the reason for saying that preferential reasoning jumps to conclusions that are not deductively justified, our framework provides no formal reason for preferring the **M**-restriction over the **M**-extension as the canonical reduction of **P** to **M**. This can of course be seen as a shortcoming of our framework, but it seems to be very hard to explain, in a semantics-independent way, why it is more natural to construct a monotonic structure from a preferential one by throwing away the preference order rather than removing exceptional states.

In the literature the emphasis has been on extensions through closure operators. This suggests a tendency to view metalevel rules as uni-directional inference rules, used to expand a given set of arguments (*cf*. the question 'What does a conditional knowledge base entail?' (Kraus *et al*., 1990; Lehmann & Magidor, 1992)). However, we have shown that, even if a rule like Monotonicity is a definite rule, it may be sometimes more natural to apply its contrapositive. In other words, such rules are not primarily inference rules, but rather rationality postulates constraining reasoning behaviours. Any consequence relation satisfying a particular rule system is considered rational with respect to the reasoning form axiomatised by [redacted] stronger rule system puts [redacted]asoning behaviours — but [redacted]whether the extra rules lead [redacted]orm of reasoning.

[redacted]y the situation with respect [redacted]notonicity, as studied by [redacted]y: $\dfrac{\alpha \;\not|\!\sim\; \neg\beta \;,\; \alpha \;|\!\sim\; \gamma}{\alpha\wedge\beta \;|\!\sim\; \gamma}$

[redacted] extended with Rational [redacted]al Monotonicity is an [redacted]closure operator (there may be several smallest **R**-relations containing a given consequence relation $|\!\sim$). From the metalevel viewpoint this is a perfectly natural situation: our metarules are rationality postulates, which may be simply too weak to fully prescribe the behaviour of a reasoning agent. On the other hand, from the connective viewpoint such indefiniteness is clearly unsatisfying, and Lehmann and Magidor go at great lengths to define the notion of rational closure (a preferred superset of $|\!\sim$ satisfying **R**). However, notice that Rational Monotonicity is a co-definite rule, hence we may investigate its co-closure. Now, if Rational Monotonicity were independent of the rules of **P** in the same way as Monotonicity is independent of the rules of **P** (Lemma 5), it would follow that **P** actually *restricts* to **R**, and we could define rational 'closure' of an arbitrary consequence relation as closure under **P** followed by co-closure under Rational Monotonicity. We leave the investigation of this conjecture as future work.

Acknowledgements

When writing this paper I have profited from discussions with Joe Halpern, Daniel Lehmann, John-Jules Meyer, Mark Ryan, and participants of the *Third Dutch/German Workshop on Nonmonotonic Reasoning Techniques and their Applications*, where an earlier version of this paper was presented. Remarks of anonymous reviewers have also been helpful. Part of this work was supported by Esprit IV Long Term Research Project 20237 ILP2.

References

Salem Benferhat, Didier Dubois and Henri Prade (1992). Representing default rules in possibilistic logic. *Proc. 3d Int. Conf. on Principles of Knowledge Representation and Reasoning*, pp.673–684. Morgan Kaufmann.

Salem Benferhat, Didier Dubois and Henri Prade (1996). Beyond counter-examples to nonmonotonic formalisms: a possibility-theoretic analysis. *Proc. 12th Int. Eur. Conf. on Artificial Intelligence*, pp.652–656. John Wiley.

Peter Flach (1995). Conjectures: an inquiry concerning the logic of induction. PhD thesis, Tilburg University.

Peter Flach (1996). Rationality postulates for induction. *Proc. 6th Int. Conf. on Theoretical Aspects of Rationality and Knowledge*, pp.267–281. Morgan Kaufmann.

Dov Gabbay (1985). Theoretical foundations for non-monotonic reasoning in expert systems. In *Logics and Models of Concurrent Systems*, pp.439–457. Springer.

Sarit Kraus, Daniel Lehmann and Menagem Magidor (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207.

Daniel Lehmann and Menagem Magidor (1992). What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60.

David Makinson (1989). General theory of cumulative inference. *Proc. 2nd Int. Workshop on Non-Monotonic Reasoning*, pp.1–18. Lecture Notes in Artificial Intelligence 346, Springer.

David Makinson (1994). General patterns in nonmonotonic reasoning. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol.3, pp.35–110. Clarendon Press.

Zbigniew Stachniak (1993). Algebraic semantics for cumulative inference operations. *Proc. 11th Nat. Conf. on Artificial Intelligence AAAI-93*, pp.444–449. MIT Press.