Published in:
Logic Journal of the IGPL 11 (2003) 69-96.
http://www3.oup.co.uk/igpl/Volume_11/Issue_01/
The present text corrects an error in the published version: see note at end

Bridges between Classical and Nonmonotonic Logic

David Makinson, *Department of Computer Science, King's College London, Strand Campus, WC2R 2LS London, UK. E-mail: makinson@dcs.kcl.ac.uk*

Abstract

The purpose of this paper is to take some of the mystery out of what is known as nonmonotonic logic, by showing that it is not as unfamiliar as may at first sight appear. In fact, it is easily accessible to anybody with a background in classical propositional logic, provided that certain misunderstandings are avoided and a tenacious habit is put aside. In effect, there are logics that act as natural bridges between classical consequence and the principal kinds of nonmonotonic logic to be found in the literature. Like classical logic, they are perfectly monotonic, but they already display some of the distinctive features of the nonmonotonic systems. As well as providing easy conceptual passage to the nonmonotonic case these logics, which we call paraclassical, have an interest of their own.

Keywords: Nonmonotonic logic, Defaults, Assumptions, Valuations, Rules

1 Introduction

1.1 Recalling Classical Consequence

We assume familiarity with classical propositional consequence, but refresh the memory with some points that are essential for what is to follow.

Classical logic uses a formal language whose propositions (or formulae – we will use the two terms interchangeably) are made up from an infinite list of elementary letters by means of the two-place connectives $\wedge,\vee$ and the one-place connective $\neg$, understood in terms of their usual truth-tables, respectively for conjunction, disjunction, and negation. The set of all these formulae is called L.

An *assignment* is a function on the set of all elementary letters into the two-element set {1,0}. Each assignment may be extended in a unique way to a *valuation*, that is, a function v on the set of all formulae into the two-element set {1,0} that agrees with the assignment on elementary letters and behaves in accord with the standard truth-tables for the compound formulae made up using $\wedge,\vee,\neg$. When A is a *set* of formulae, one writes $v(A) = 1$ as shorthand for $v(a) = 1$ for all $a \in A$.

Let A be any set of formulae, and let x be an individual formula. One says that x is a *classical consequence* of A iff there is no valuation v such that $v(A) = 1$ whilst $v(x) = 0$. The standard notation is A |- x, and the sign |- is called 'gate' or 'turnstile'. When dealing with individual formulae on the left, the notation is simplified a little by dropping parentheses, writing a |- x in place of $\{a\}$ |- x and $x \in Cn(\{a\})$.

Thus classical consequence is a *relation* between propositions, or more generally between *sets* A of such propositions and individual propositions x. It may also be seen as an *operation* acting on sets A of propositions to give larger sets $Cn(A)$. These two representations of classical consequence are trivially interchangeable. Given a relation |-, we may define the operation Cn by setting $Cn(A) = \{x: A$ |- $x\}$; and conversely we may define |- from Cn by the rule A |- x iff $x \in Cn(A)$.

Both of the representations are useful. Sometimes one is more convenient than another. For example, it is often easier to visualize things in terms of the relation, but more concise to formulate and prove them using the operation. The same will be true when we come to non-classical consequence. For this reason, in this paper we will constantly be hopping from one notation to the other as two ways of saying the same thing, and we encourage the reader to do the same.

Classical consequence has a number of useful properties. To begin with, it is a *closure relation*, in the sense that it satisfies the following three conditions for all formulae a,x and all sets A,B of formulae:

Reflexivity alias *Inclusion*	A |- a whenever $a \in A$
Cumulative Transitivity, CT (alias *Cut*)	Whenever A |- b for all $b \in B$ and $A \cup B$ |- x then A |- x
Monotony	Whenever A |- x and $A \subseteq B$ then B |- x

Expressed in the language of *Cn*, this means that classical consequence is a *closure operation* in the sense that it satisfies the following conditions for all sets *A*,*B* of formulae.

Reflexivity alias *Inclusion*	$A \subseteq Cn(A)$
Cumulative Transitivity, CT (alias *Cut*)	$A \subseteq B \subseteq Cn(A)$ implies $Cn(B) \subseteq Cn(A)$
Monotony	$A \subseteq B$ implies $Cn(A) \subseteq Cn(B)$

Those already familiar with the concept of a closure operation in abstract algebra or topology may be a little surprised to see the condition of cumulative transitivity where they are accustomed to seeing idempotence, i.e. the condition that $Cn(A) = Cn(Cn(A))$. The two are equivalent, not as individual conditions, but as components of the trio of conditions defining the notion of a closure operation. The left-in-right half of idempotence already follows from inclusion, while its converse half is covered by a limiting instance of CT, namely the one where *B* is chosen to be $Cn(A)$. Conversely, CT is immediate from monotony and idempotence. However, when considering relations that fail monotony, it turns out best to focus on CT rather than idempotence.

The three conditions defining the notion of a closure relation are examples of what are known as *Horn rules*. Roughly speaking, a Horn rule tells us that if such-and-such and so-and-so (any number of times) are all elements of the relation, then so is something else. None of the suppositions of a Horn rule may be negative – none of them can require that something is *not* an element of the relation. Nor is the conclusion allowed to be disjunctive – it cannot say that given the suppositions, *either* this *or* that is in the relation. Horn rules have very useful properties, most notably that whenever a Horn rule is satisfied by every relation in a family, then the relation formed by taking the intersection of the entire family also satisfies the rule.

Finally, we recall that classical consequence is *compact*, in the sense that whenever A |- x then there is a finite subset $A' \subseteq A$ with A' |- x. In the language of operations, whenever $x \in Cn(A)$ then there is a finite subset $A' \subseteq A$ with $x \in Cn(A')$.

These are *abstract* properties of classical consequence, in the sense that they make no reference to any of the connectives $\wedge,\vee,\neg$. Evidently, the relation also has a number of properties concerning each of these connectives arising from their respective truth-tables, for example the property that $a \vee b \in Cn(a)$. We shall not enumerate these, but recall just one that will play a very important role in what follows: the property of *disjunction in the premises, alias* OR. It says that whenever $A \cup \{a\}$ |- x and $A \cup \{b\}$ |- x then $A \cup \{a \vee b\}$ |- x. In the language of operations, $Cn(A \cup \{a\}) \cap Cn(A \cup \{b\}) \subseteq Cn(A \cup \{a \vee b\})$.

1.2 Some Misunderstandings

For a person coming to nonmonotonic reasoning for the first time, it can be rather difficult to get a clear grip on what is going on. This is partly due to some misunderstandings which, while very natural, distort understanding from the beginning. We warn the reader in advance, even if some of the technical details can become fully clear only as the story develops. The important thing is to begin with the right *gestalt*.

Weaker or Stronger?

The first thing that one hears about nonmonotonic logic is, evidently, that it is not monotonic. In other words, it fails the principle that whenever x follows from a set A of propositions then it also follows from every set B such that $B \supseteq A$. By contrast, classical logic satisfies this principle, as is immediate from its definition in the preceding section.

Given the failure of this classical principle, it is natural to imagine that nonmonotonic logic is weaker than classical logic. And indeed, in one sense it is. Typically, the set of Horn rules that a nonmonotonic consequence relation satisfies is a proper subset of those satisfied by classical consequence. For example, for preferential consequence relations (to be explained below), the rules of reflexivity and cumulative transitivity always hold while monotony need not.

But in another, and much more basic sense, nonmonotonic logics are *stronger* than their classical counterpart. Recall that classical consequence is a relation, i.e. under the usual understanding of relations in set-theoretical terms, a set of ordered pairs. Specifically, |- is a set of ordered pairs (A,x), where A is a set of propositions and x is an individual proposition. If written as an operation Cn, it is a set of ordered pairs (A,X), where both components are sets of propositions. It is at this level that the most basic comparison of classical with nonmonotonic consequence arises.

Suppose we take a nonmonotonic consequence relation |~ (usually called 'snake'). It too is a set of ordered pairs (A,x). And as we will see, for the typical instances in the literature, it is a *superset* of the classical consequence relation. In other words, |- $\subseteq$ |~, where $\subseteq$ is set inclusion. Likewise, suppose that we take a nonmonotonic consequence operation, usually referred to by the letter C. Then we have $Cn \leq C$, where $\leq$ is not quite set inclusion between the two operations, but set inclusion between their values, i.e. $Cn(A) \leq C(A)$ for all sets A of propositions.

It is in this sense that nonmonotonic consequences relations are stronger than classical consequence, and for this reason they are referred to as *supraclassical* relations. Although they are stronger, they are rather less regular in their behaviour. Certain Horn rules, such as monotony and transitivity, which hold for the (smaller) classical relation, can fail for the (larger) nonmonotonic ones. In technical language, we have |- $\subseteq$ |~ and $Cn \leq C$, even when H(|~) $\subseteq$ H(|-) and $H(C) \subseteq H(Cn)$, where $H(.)$ is the set of Horn rules that the relation or operation satisfies.

In what follows, this relationship should be borne in mind. The relations that we will be considering, both monotonic and nonmonotonic, are all supraclassical in the sense of including classical consequence, even when failing certain of the classical Horn principles.

Classical or Non-Classical?

In so far monotony fails, the logic of nonmonotonic consequence relations certainly differs from classical logic. But it would be quite misleading to refer to it as a kind of non-classical logic as that term is used when referring to, say, intuitionistic logic. For in contrast to that case, we do not reject classical consequence as incorrect – as already remarked, it is included in the nonmonotonic relations under study. Nor do we say that there is anything wrong with monotony itself. We are showing how the 'good old relation of classical consequence' may be deployed in certain ways, to define further, stronger relations that are of practical value, but which happen to fail monotony.

One Logic or Many?

There is a third common misunderstanding of what nonmonotonic logic is all about. From the classical context, we are familiar with the idea that there is just one core logic, up to notational differences and matters like choice of primitives. That core is classical logic, and it is also the logic that we use when reasoning ourselves in the metalanguage.

Even intuitionists and relevantists, who do not accept all of classical logic, feel the same way, but about their own systems, which are subsystems of the classical one. They have some difficulties, one might add, in reconciling this view with their own practice in the metalanguage, where they usually use classical logic.

Given the unicity of classical inference, it is natural for the student, puzzled by seeing several different kinds of nonmonotonic consequence, to ask: what is *real* nonmonotonic inference? Which is the *correct* nonmonotonic consequence relation? What is the one that we use in practice, even if we can study others?

The answer is that *there is none*. There is no unique nonmonotonic consequence relation, but indefinitely many of them. We have all those relations that can be generated from certain kinds of structure, whose ingredients are allowed to vary freely within the boundaries of suitable formal conditions. Likewise, there are all those relations satisfying certain syntactic conditions such as conjunction of conclusions, disjunction of premises, cumulative transitivity, or cautious monotony. Moreover, if one tries to get away from non-uniqueness by intersecting all the many relations or operations, the result is just classical logic.

Leaving aside technical details, the essential message is as follows. Don't expect to find *the* nonmonotonic consequence relation that will always, in all contexts, be the right one to use. Rather, expect to find several families of such relations, some interesting conditions that they sometimes satisfy, sometimes fail, and some interesting ways of generating them mathematically from underlying structures.

This intrinsic non-uniqueness was, in the historical development of logic, a rather new feature of the theory of nonmonotonic inference and, we should add, of the parallel development of AGM belief revision. It appears to be responsible for some of the initial difficulties that the enterprise experienced in being assimilated by the broader community of logicians.

1.3 A Habit to Suspend

We will be showing that there are systems that act as natural bridges between classical consequence and nonmonotonic logics. These bridge systems are also supraclassical, but they are perfectly monotonic, indeed they are closure operations.

'But if these bridge systems are supraclassical', one may ask, 'how are they possible? Classical consequence is *already* maximal, in the sense that there is no stronger closure operation in the same underlying language, other than the trivial one by which every proposition of the language implies every other one. So how can these 'bridge logics' be closure operations and at the same time supraclassical?'

Indeed, this result has been part of the folklore since the early twentieth century. But the formulation above omits a vital element, which is not always made as explicit as it should be. That is the condition of being *closed under substitution*, which we now explain.

By *substitution* we mean what is often called, at greater length, *uniform substitution of arbitrary formulae for the elementary letters in a formula*. For example, when a is the formula $p\wedge(q\vee\neg r)$, where p,q,r are three distinct elementary letters, then one substitution σ may replace all occurrences of p by, say, r, all occurrences of q by $\neg p$, and (simultaneously, not subsequently) all occurrences of r by $\neg(p\wedge s)$. That will give us $\sigma(a) = r\wedge(\neg p\vee\neg\neg(p\wedge s))$. Simplifications, such as the elimination of the double negation, are not part of the substitution, but possible later operations. Uniform substitution for elementary letters is a kind of construction, not an inference. It should not be confused with another operation, rather confusingly also sometimes called substitution, which replaces one of more occurrences of a formula (not just an elementary letter) by another formula to which it is classically equivalent (not by an arbitrary formula).

When A is a *set* of formulae, substitution on it is understood pointwise; that is, $\sigma(A)$ is understood to be $\{\sigma(a) : a \in A\}$.

Classical consequence is *closed under substitution*, in the sense that whenever A |- x then $\sigma(A)$ |- $\sigma(x)$. In the language of operations, whenever $x \in Cn(A)$ then $\sigma(x) \in Cn(\sigma(A))$, or more concisely $\sigma(Cn(A)) \subseteq Cn(\sigma(A))$.

The maximality of classical logic may now be expressed as follows: *there is no supraclassical closure relation in the same language as classical* |-, *that is closed under substitution, except for* |- *itself and the total relation*. Likewise for operations. Here, the total relation is the one that relates every premise (or set of premises) to every conclusion; as an operation it sends any set of formulae to the set of all formulae. The proof of the observation is indeed a simple one, and we recall it here.

- *Proof of the maximality of classical logic*. Let |-′ be any closure relation that is closed under substitution and also properly supraclassical, i.e. |- ⊂ |-′. By the last hypothesis, there are A,x with A |-′ x but A |-/ x. From the latter, there is a classical valuation v with $v(A) = 1$, $v(x) = 0$. Substituting tautologies for elementary letters that are true under v, and contradictions for letters that are false under v we see, by an easy induction on depth of formulae, that $\sigma(A)$ is a set of tautologies and $\sigma(x)$ is a contradiction. Since by hypothesis $|\!\!-'$ is closed under substitution, $A \,|\!\!-' x$ implies $\sigma(A) \,|\!\!-' \sigma(x)$. But since $\sigma(A)$ is a set of tautologies we have by classical logic that for arbitrary B, $B \,|\!\!- \sigma(a)$ for all $\sigma(a) \in \sigma(A)$, and likewise since $\sigma(x)$ is a contradiction we have $\sigma(x) \,|\!\!- y$ for every formula y. Thus since $|\!\!- \subseteq |\!\!-'$ we have $B \,|\!\!-' \sigma(a)$ for all $\sigma(a) \in \sigma(A)$, and $\sigma(A) \,|\!\!-' \sigma(x)$, and $\sigma(x) \,|\!\!-' y$. Putting these three together with cumulative transitivity and monotony of $|\!\!-'$, we get $B \,|\!\!-' y$ and the proof is complete.

The moral of this story is that the supraclassical closure relations that we shall be offering as bridges between classical consequence and nonmonotonic consequence relations *are not closed under substitution*. Nor, for that matter, are the nonmonotonic relations that issue from them. This runs against ingrained habit. Students of logic are brought up with the idea that any decent consequence relation should be purely formal, or structural, and hence satisfy substitution. Indeed, those terms are often used in the texts as synonyms for closure under substitution. To understand nonmonotonic logic, this is a habit to suspend.

1.4 Three Ways of Getting More Conclusions

In this paper we will describe three different ways of getting out of a set of premises more than is authorized by straightforward application of classical consequence, without amplifying the language in which these premises are stated, which remains that of classical logic. Roughly speaking, the first method uses additional background *assumptions*. The second restricts the set of *valuations* that are considered possible. And the third uses additional background *rules*.

Each of these procedures gives rise to a corresponding kind of monotonic consequence operation, which we call *pivotal-assumption*, *pivotal-valuation*, and *pivotal-rule* consequence respectively. The term 'pivotal' is used because the additional assumptions, the restriction in the set of valuations, and the additional rules, are held fixed irrespective of variation in the premises of the consequence operation. As we will see, the three kinds of monotonic operation are not entirely equivalent to each other. But they are all supraclassical closure operations, i.e. they include classical consequence and also satisfy both cut and monotony. Such consequence relations will be called *paraclassical*.

The three kinds of paraclassical consequence serve as conceptual bridges to corresponding families of nonmonotonic consequence, formed essentially by allowing the additional background assumptions, valuation restriction, and additional rules to vary in a principled manner with the premises under consideration. In this way we obtain three satellite kinds of nonmonotonic consequence operation: *default-assumption*, *default-valuation*, and *default-rule* consequence.

We emphasize that all of these concepts are developed over a purely boolean propositional language, without adding further connectives to the usual truth-functional ones. In particular, we will not be considering autoepistemic logics, in which one adds to the boolean connectives a special kind of modal operator whose self-referential characteristics engender a nonmonotonic behaviour.

2 First Bridge - Using Additional Background Assumptions

2.1 From Classical Consequence to Pivotal Assumptions

We begin by examining the simplest kind of paraclassical consequence and its transformation into a form of nonmonotonic reasoning, namely inference with additional background assumptions.

In daily life, the assumptions that we make when reasoning are not all of the same level. Generally, there will be a few that we display explicitly, because they are special to the situation under consideration or in some other way deserving particular attention. There will usually be many others that we do not bother even to mention, because we take them to be part of shared common knowledge, or in some other way trivial. This phenomenon was already known to the ancient Greeks, who used the term *enthymeme* to refer to an argument in which one or more premises are left implicit. That is the idea that we develop in this section.

We work with the same propositional language as in classical logic, with the set of all its formulae called L. Let $K \subseteq L$ be a set of formulae. Intuitively K will be playing the role of a set of background assumptions or, as they are called in Gärdenfors and Makinson (1994), 'expectations'. Let A be any set of formulae, and let x be an individual formula.

We say that x is a *consequence of A modulo the assumption set K*, and write $A \mid\!\!-_K x$ *alias* $x \in Cn_K(A)$ iff there is no valuation v such that $v(K \cup A) = 1$ whilst $v(x) = 0$. Equivalently, iff $K \cup A \mid\!\!- x$. As simple as that! And we call a relation or operation a *pivotal-assumption consequence* iff it is identical with $\mid\!\!-_K$ (resp. Cn_K) for some set K of formulae. Note that there is not a unique pivotal-assumption consequence relation, but many – one for each value of K.

Since classical consequence is monotonic, pivotal-assumption consequence relations and operations are supraclassical in the sense defined earlier. That is, for every K we have $\mid\!\!- \subseteq \mid\!\!-_K$ and $Cn \leq Cn_K$. They also share a number of abstract properties with classical consequence. In particular, as is immediate from the definition, they satisfy inclusion, cumulative transitivity and monotony, and thus are closure operations. They are also compact, and have the property of disjunction in the premises.

Given that they are supraclassical closure relations, we know from the general observation in section 1.4 that the $\mid\!\!-_K$ are not, in general, closed under substitution. It is interesting to see through an example why this is so. Let $K = \{p\}$where p is an elementary letter. Choose any other elementary letter q, and put $A = \{q\}$ and $x = p \wedge q$. Then $A \mid\!\!-_K x$ since $\{p,q\} \mid\!\!- p \wedge q$. Now let σ be the substitution that replaces every elementary letter by, say, its own negation, so that $\sigma(p) = \neg p$ and $\sigma(q) = \neg q$. Then $\sigma(A) \mid\!\!-\!\!/_K\ \sigma(x)$ since $K \cup \{\neg q\} = \{p, \neg q\} \mid\!\!-\!\!/\ \neg p \wedge \neg q$. Analysing this example, we see what makes it work: the substitution is applied to the explicit premises A and the conclusion x, but not to the background assumption set K, which is held constant.

A striking feature of pivotal-assumption consequence (which separates it from the next bridge system that we will be describing) is that the above properties also suffice to characterize it. In other words, we have the following representation theorem for pivotal-assumption consequence relations (and likewise for the corresponding operations). *Let $\mid\!\!-'$ be any supraclassical closure relation that is compact and satisfies the condition of disjunction in the premises. Then there is a set K of formulae such that $\mid\!\!-' = \mid\!\!-_K$.*

This representation theorem is formulated in Rott (2001) (section 4.4 observation 5), but appears to have been part of the folklore for many decades. Its proof is abstract but short. It is most easily expressed using the operation notation. We give it here, but some readers may prefer to skip it, passing directly to section 2.2.

- *Proof of representation theorem for pivotal-assumption consequence*. Let Cn' be any supraclassical closure operation that is compact and satisfies the condition of disjunction in the premises. Put $K = Cn'(\emptyset)$, where $\emptyset$ is the empty set of formulae. We claim that $Cn' = Cn_K$. It will suffice to show both $Cn_K \leq Cn'$ and $Cn' \leq Cn_K$.

  For the inclusion $Cn_K \leq Cn'$ we need to show that for any A, $Cn(Cn'(\emptyset)\cup A) \subseteq Cn'(A)$. Since by supraclassicality $Cn \leq Cn'$, we have $Cn(Cn'(\emptyset)\cup A) \subseteq Cn'(Cn'(\emptyset)\cup A) = Cn'(A)$ by well-known properties of closure operations.

  For the converse inclusion $Cn' \leq Cn_K$ we need to show that for any A, $Cn'(A) \subseteq Cn(Cn'(\emptyset)\cup A)$. This is where we also need compactness and disjunction in the premises. Suppose $x \in Cn'(A)$. Then by the compactness of Cn' there is a finite $B \subseteq A$ with $x \in Cn'(B)$. Let b be the conjunction of all the finitely many elements of B. Since Cn' is a supraclassical closure operation, we get $x \in Cn'(b)$ and thus in turn $\neg b\vee x \in Cn'(b)$. But again by supraclassicality of Cn', $\neg b\vee x \in Cn'(\neg b)$. Applying disjunction in the premises as the *coup de grace*, we have $\neg b\vee x \in Cn'(b\vee\neg b) = Cn'(\emptyset)$ with the help of supraclassicality and closure again. To show that $x \in Cn(Cn'(\emptyset)\cup A)$ as desired, it will thus suffice to show that $x \in Cn(\{\neg b\vee x\}\cup A)$. But by the construction of b, we have by properties of classical consequence that $b \in Cn(A)$ and so again by properties of classical consequence $x \in Cn(\{\neg b\vee x\}\cup A)$, thereby completing the proof.

2.2 From Pivotal Assumptions to Default Assumptions

What has all this to do with nonmonotonic inference? The consequence relations $|\!-_K$ are, as we have seen, monotonic. But nonmonotonicity is generated if we allow the background assumption set K to vary with A, in particular, to be diminished when A is inconsistent with K.

Specifically, we obtain an interesting nonmonotonic relation if we work with the maximal subsets K' of K that are consistent with A, and accept as output only what is common to their separate outputs. We call this relation *default-assumption consequence*, to bring out its close relation to the preceding pivotal-assumption consequence.

To give the definition more explicitly, let $K \subseteq L$ be a set of formulae, which again will play the role of a set of background assumptions. Let A be any set of formulae, and let x be an individual formula.

- We say that a subset K' of K is consistent with A iff there is a classical valuation v with $v(K'\cup A) = 1$.

- A subset K' of K is called *maximally consistent* (briefly, *maxiconsistent*) *with* A iff it is consistent with A but is not a proper subset of any $K'' \subseteq K$ that is consistent with A.

- Finally, we define the relation $|\sim_K$ of *consequence modulo the default assumptions* K by putting $A \mid\sim_K x$ iff $K'\cup A \mid\!- x$ for every subset $K' \subseteq K$ that is maxiconsistent with A. Writing C_K for the corresponding operation, this puts $C_K(A) = \cap\{Cn(K'\cup A)\colon K' \subseteq K$ and K' maxiconsistent with $A\}$.

We call a relation or operation a *default-assumption consequence* iff it is identical with $|\sim_K$ (resp. C_K) for some set K of formulae. Note again that there is not a unique default-assumption consequence relation, but many – one for each value of K.

Default-assumption consequence operations/relations are nonmonotonic. That is, we may have $A \mid\sim_K x$ but not $A\cup B \mid\sim_K x$ where A,B are sets of propositions. Likewise, we may have $a \mid\sim_K x$ without $a\wedge b \mid\sim_K x$ where a,b are individual propositions.

To illustrate the failure of monotony, suppose $K = \{p\rightarrow q,\ q\rightarrow r\}$ where p,q,r are distinct elementary letters of the language and $\rightarrow$ is the truth-functional (*alias* material) conditional connective. Then $p \mid\sim_K r$ since the premise p is consistent with the whole of K and clearly $\{p\}\cup K \mid\!- r$. But $p\wedge\neg q \mid\!\sim\!/_K\, r$, for the premise $p\wedge\neg q$ is no longer consistent with the whole of K. There is a unique maximal subset $K' \subseteq K$ that is consistent with $p\wedge\neg q$, and that is $K' = \{q\rightarrow r\}$, and clearly $\{p\wedge\neg q\}\cup K' \mid\!-\!/\, r$ (witness the valuation v with $v(p) = 1$ and $v(q) = v(r) = 0$). We gained premises, but because of the consistency requirement we lost assumptions.

As monotony can fail, default-assumption consequence operations are not in general closure operations, even though they satisfy both inclusion and cumulative transitivity (and hence also idempotence). They are also supraclassical and satisfy disjunction in the premises.

We note one further property. Although they may fail monotony, they always satisfy a weakened version of it called *cautious monotony*. This says that whenever $A \mid\sim_K x$ and $A \mid\sim_K y$ then $A\cup\{x\}\mid\sim_K y$. More generally: whenever $A \mid\sim_K x$ for all $x\in B$ and $A \mid\sim_K y$ then $A\cup B \mid\sim_K y$. In the succinct notation of operations: whenever $A \subseteq B \subseteq C_K(A)$ then $C_K(A) \subseteq C_K(B)$. We omit the verification; it can be found for example in Makinson (1994), which reviews the behaviour of these operations in considerable detail.

Conceptually, monotonic pivotal-assumption consequences Cn_K provide a natural ‘half-way house’ between classical Cn and nonmonotonic default-assumption consequences C_K. The definition of Cn_K adds the component K of background assumptions. The default operation requires as well the consistency of these assumptions with the set A of explicit premises, failing which it contracts from the assumptions to ensure consistency.

On the other hand, in terms of set-inclusion the order is reversed. It is default-assumption consequence C_K that lies between classical Cn and pivotal-assumption consequence. That is, we have $Cn \leq C_K \leq Cn_K$, since $Cn(A) \subseteq Cn(K'\cup A) \subseteq Cn(K\cup A)$ whenever $K' \subseteq K$.

It is time to confess to a terminological embarrassment. We are using the term ‘consequence operation’ to cover both monotonic operations like Cn, Cn_K and nonmonotonic ones like C_K. In this, we are following the usage of some authors, such as Kraus, Lehmann and Magidor (1990), but we are running against that of others, such as Lindström (1991), Makinson (1994) and Rott (2001). They mark the monotonic/nonmonotonic divide by speaking of ‘consequence’ operations for the former and ‘inference’ operations for the latter. However, this terminological separation is rather cumbersome to maintain in practice, particularly as it is not rooted in any verbal contrast to be found in ordinary language. In this paper, therefore, we use the term 'consequence' in both the monotonic and nonmonotonic cases. However, we mark the difference notationally, using Cn with subscripts for the monotonic operations, and C with subscripts for the nonmonotonic ones.

We should also issue a warning about the kinds of assumption set that may usefully be used here. The operation of default-assumption consequence diminishes in interest as the set K of background assumptions becomes more redundant, i.e. contains more consequences of its own elements. In the limiting case that K is already closed under classical consequence, i.e. when $K = Cn(K)$, then notoriously the operation becomes devoid of interest, because it collapses into classical consequence in the principal case that the premises A are inconsistent with the background assumptions K. To be precise: when $K = Cn(K)$ then either $C_K(A) = Cn(K \cup A)$ or $C_K(A) = Cn(A)$ according as A is, or is not, consistent with K.

Default-assumption consequence has a rather long and complicated history. For singleton $A = \{a\}$ the operation was defined by Alchourrón and Makinson (1982), but as an account of the *revision* of a belief set K by a new belief a. The collapse into classical consequence when $K = Cn(K)$ was first proven there. The family of all subsets of K maximally consistent with A was also studied by Poole (1988), but as part of a study of *abduction*, i.e. the formation of hypotheses to explain data. For an overview, see e.g. Makinson (1994).

2.3 Generalizations and Variants

The concept of default-assumption consequence can be generalized in several interesting ways.

<u>Default-Assumptions with Constraints</u>

One generalization is obtained by adding to the set K of background assumptions another set J of 'constraints'. This is also required to be consistent with the explicit premises of an inference. But when this consistency fails, J is not diminished; nor does it participate actively in the generation of conclusions. In effect, it reduces the quantity of background assumptions from K that may legitimately be used alongside A.

To be precise, the definition of default-assumption consequence is re-run as follows. Let $K,J \subseteq L$ be sets of formulae, which will play the role of background assumptions and constraints respectively. We define the relation $|\!\sim_{KJ}$ of *default-assumption consequence with constraints* by putting $A \mathrel{|\!\sim_{KJ}} x$ iff $x \in Cn(K' \cup A)$ for every subset $K' \subseteq K$ with K' maximal among the subsets of K such that K' is consistent with $A \cup J$. Writing C_{KJ} for the corresponding operation, this says $C_{KJ}(A) = \cap\{Cn(K' \cup A)$: $K' \subseteq K$ and K' maxiconsistent with $A \cup J\}$. If there are no such subsets, i.e. in the limiting case that J itself is inconsistent with A, this puts $C_{KJ}(A) = L$. As one might expect, this operation is rather less regularly behaved than the one without constraints. It was introduced by Poole (1988) alongside the version without constraints.

<u>Partial Meet Default Assumptions</u>

Another possible generalization arises if one is more selective when intersecting the maximal subsets of K. Instead of intersecting *all* of the sets $Cn(K' \cup A)$, for subsets $K' \subseteq K$

maxiconsistent with A, one might want to intersect only *some* of them. The decision which to intersect may in turn be made by an arbitrary selection function (which in effect imposes no constraints on the choice) or by adding a relation over the K' and picking out the maximal (or, what amounts to the same from an abstract perspective, the minimal) ones.

This generalization corresponds closely to that effected in the 'partial meet' (*alias* AGM) approach to belief change, developed by Alchourrón, Gärdenfors and Makinson (1985). It has the advantage that it allows the set K of background assumptions to be closed under classical consequence without thereby collapsing the nonmonotonic operation into the classical one, as we saw (section 2.2) happens when one intersects *all* the sets $Cn(K'\cup A)$.

Maximally Informative Background Assumption Sets

A third possible generalization is to relax the requirement of maximality. The idea is that we should consider not just the biggest subsets, but the most *informative* subsets of K that are consistent with the premises. All of the maximally large subsets will be maximally informative, but some among the non-maximal subsets may be just as informative as their maximal supersets. In this case, it is suggested, they should also be counted in the family scheduled for intersection, thus in turn diminishing the set obtained by that intersection.

This kind of approach has been urged by Isaac Levi (for parallel constructions in the logic of belief revision) in a number of publications including his (1996). Its formal properties have been studied by Rott and Pagnucco (1999).

One could also take some measure of value that combines several competing factors, not all of which are positively correlated with size. For example, the measure might balance size, informativeness, simplicity of formulation for man or machine, simplicity of application, etc. In this case not all of the maximal subsets of K that are consistent with the premises will be of maximal value, nor conversely. As far as the author is aware, this level of generalization has not been explored in formal terms.

Background Boundaries

An interesting variant of default-assumption consequence was formulated by Forget, Risch and Siegel (2001), using an idea of Siegel and Forget (1996). Instead of working with subsets K' of the set K of background assumptions, they look at individual elements $k \in K$. At the same time, instead of requiring that $A\cup K' \mid\!- x$ for every appropriate K', they require only that $A\cup\{k\}$ is *consistent* with x for every appropriate k. This treats K like a set of boundary elements, with each of which consistency is desired, rather than as an auxiliary set of assumptions.

Stated in full, the definition as follows, writing K as a superscript rather than as a subscript to mark the difference with our principal notion of default-assumption consequence. We define $A \mid\!\sim^{K} x$ iff $A\cup\{k\}$ is consistent with x for every $k \in K$ that is consistent with A. In other words: iff $A\cup\{k\} \mid\!/\!- \neg x$ for every $k \in K$ with $A\cup \mid\!/\!- \neg k$.

In fact, Forget, Risch and Siegel (2001) defined their relation dually, in terms of a 'bad' set Z rather than a 'good' set K. Their actual definition is as follows: $A \mid\!\sim^{Z} x$ iff for every $z \in Z$, if $A\cup\{x\} \mid\!- z$ then $A \mid\!- z$. Intuitively: iff adding the potential conclusion to the set of premises does not increase the 'bad' propositions that are implied beyond what A already implies. The two definitions are easily verified to be equivalent when we put $K = \{\neg z\colon z \in Z\}$ and conversely $Z = \{\neg k\colon k \in K\}$.

It is also easy to verify that for any K we have $|\!\sim_K \subseteq |\!\sim^K$. That is, whenever $A |\!\sim_K x$ then $A |\!\sim^K x$. In other words. Poole default consequence modulo an assumption set K is a subrelation of Forget/Risch/Siegel consequence modulo the same set.

Comparative Expectation and Related Ideas

One may also elaborate the notion of background assumptions in quite different ways, giving variant accounts. We describe briefly one of the most influential.

Instead of fixing the background assumptions as a *set*, one may say that to be a background assumption is a matter of degree. Any consistent proposition of the language may to serve as an additional assumption, to a degree limited by its plausibility. Plausibility is represented by a *relation* $<$ over formulae of the language. Thus $z<y$ is read as saying that y is strictly more plausible (less unexpected, etc) than z.

Given such a relation of plausibility, we need to define an appropriate notion of consequence that allows sufficiently plausible propositions to be used as additional assumptions. One way of doing this is to define $a |\!\sim_< x$ to hold iff $x \in Cn(\{a\} \cup \{y\colon \neg a<y\})$. In words: $a |\!\sim_< x$ iff x is a classical consequence of premise a supplemented by all those propositions that are strictly more plausible than $\neg a$. The negation attached to a here is not an oversight. For with appropriate conditions imposed on the relation $<$, the y with $\neg a<y$ form a subset of the propositions y that are classically consistent with a whenever a is itself consistent. Further conditions on the relation $<$ give rise to further elegant properties of the generated relation $|\!\sim_<$.

The construction thus gives a motivated selection of a *unique* set K, depending on a, such that $Cn(\{a\} \cup K)$ is consistent when a itself is: namely $K = \{y\colon \neg a<y\}$. Under this approach we no longer need to intersect a family of sets, and this is usually regarded as a positive feature. It was introduced by Gärdenfors and Makinson (1994) under the name of 'comparative expectation inference', and corresponds closely to an approach to belief revision, called 'belief revision via epistemic entrenchment' that was developed by the same authors in an earlier paper (1988).

The definition of $|\!\sim_<$ given above makes sense only for individual propositions a as premises, as it makes use of the negation of the premise, and negation is not defined as an operation on arbitrary (infinite) sets of propositions. However, it can be reformulated so as to apply also to infinite premise sets. We will not go into further detail about this kind of relation, referring the reader to the recent book Rott (2001).

2.4 Recapitulation

Rising above details, the main point that we want to insist upon in this section is that an important family of nonmonotonic consequence operations (namely, default-assumption consequence and its generalizations, and also comparative expectation inference and its variants) arises from a perfectly monotonic one (pivotal-assumption consequence) by allowing its distinctive ingredient (the set K of background assumptions) to vary with the premise set A. The monotonic family thus serves as a natural stepping-stone towards the nonmonotonic one. Ultimately, they both reflect the ancient notion of an enthymeme.

3 Second Bridge - Restricting the Set of Valuations

3.1 From Classical Consequence to Pivotal Valuations

In the preceding section, we built a mid-station between classical and nonmonotonic reasoning. The essential idea was to augment the explicit premises A by a set K of background assumptions. We now do almost the same thing in a different way, which points towards a substantially different nonmonotonic generalization. We *restrict the set of valuations* that are considered. In other words, we fix a pivotal set W of valuations, and redefine consequence modulo it instead of the set V of all valuations.

Let $W \subseteq V$ be a set of valuations on the language L. Let A be any set of formulae, and let x be an individual formula. We say that x is a *consequence of* A *modulo the valuation set* W, and write $A \mid\!\!-_W x$ *alias* $x \in Cn_W(A)$ iff there is no valuation $v \in W$ such that $v(A) = 1$ whilst $v(x) = 0$. We call a relation or operation a *pivotal-valuation consequence* iff it coincides with $\mid\!\!-_W$ (resp. Cn_W) for some set W of valuations. Note again that there is not a unique pivotal-valuation consequence relation, but many – one for each value of W.

Immediately from this definition, pivotal-valuation consequence relations/operations are also supraclassical, i.e. $Cn \leq Cn_W$ for any choice of W. They also satisfy inclusion, cumulative transitivity and monotony, and thus are closure operations. We are thus still in the realm of paraclassical inference. Pivotal-valuation consequence relations also have the property of disjunction in the premises. But, as we would already expect, they are not closed under substitution. Moreover – and this is the new feature – they are not always compact.

To illustrate the failure of compactness, let W be the set of all those valuations v such that only finitely many elementary letters (out of our infinite supply P of elementary letters for generating formulae of the language L) are true under v. On the one hand, for no finite subset $P' \subseteq P$ do we have $P' \mid\!\!-_W (p\wedge\neg p)$, since W contains a valuation that makes all the finitely many letters in P' true, and this valuation evidently makes $(p\wedge\neg p)$ false. On the other hand, $P \mid\!\!-_W (p\wedge\neg p)$ holds vacuously, since there is no valuation in W that satisfies all the infinitely many letters in P.

This failure of compactness for some pivotal-valuation consequence operations shows that not every such operation is a pivotal-assumption one, for as we have seen, all of the latter do satisfy compactness. On the other hand, we do have the converse: every pivotal-assumption consequence operation is a pivotal-valuation one, so that the former family is strictly narrower than the latter family. The verification is almost immediate. Let Cn_K be any pivotal-assumption consequence operation with assumption set K. We need to show that $Cn_K = Cn_W$ for an appropriate choice of a set W of valuations. It suffices to find a $W \subseteq V$ such that for all A and all $v \in V$, $v(K\cup A) = 1$ iff $v \in W$ and $v(A) = 1$. We simply put W be the set of all valuations v such that $v(K) = 1$. Then for every valuation v, $v(K\cup A) = 1$ iff $v(K) = 1$ and $v(A) = 1$, i.e. iff $v \in W$ and $v(A) = 1$ and we are done.

Putting this together with points already noted, we can conclude that *the pivotal-assumption consequence operations are precisely the pivotal-valuation ones that are compact*. For we have just shown that every pivotal-assumption consequence operation is a pivotal-valuation one, and we have already remarked that the former are always compact. Conversely, we have noted that any pivotal-valuation operation is a supraclassical consequence operation satisfying disjunction in the premises; so if it is also compact the representation theorem of section 2.1 tells us that it is a pivotal-assumption relation.

In particular, in a finite language (that is, one generated by boolean connectives from a finite set of elementary letters), where compactness always holds, the families of pivotal-assumption and pivotal-valuation consequence operations coincide. For the computer scientist, who always works in a finite language, pivotal-assumption and pivotal-valuation are thus equivalent. For the logician, who takes perverse pleasure in the subtleties of the infinite, they are not. We persist with the infinite case from here to the end of the section; hard-line finitists will skip to the next one.

The question arises whether we can characterize the family of pivotal-valuation consequence operations in terms of the properties that these operations satisfy. It would be pleasant to be able to report that the family is fully characterized by the properties of being a closure operation, supraclassical, and satisfying disjunction in the premises. This would give us a representation theorem for the family, paralleling the one that we presented in section 2.1 for pivotal-assumption consequence operations.

Examination of the proof that we gave in section 2.1 for the pivotal-assumption representation theorem may make one rather sceptical about such a conjecture. For half of that proof depends applying the compactness property, which is no longer available, in order to make a single formula do the work of an infinite set A of premises in so far as any given conclusion x is concerned.

In fact, an elegant counterexample has been devised by Karl Schlechta: a supraclassical closure operation satisfying disjunction in the premises that is not a pivotal-valuation consequence operation. Schlechta's example shows even more: it is not even the intersection of any family of preferential consequence operations (to be defined in the next section); but that aspect need not concern us for the present. As the construction and especially the verification of the example are rather intricate, we omit them here. See Schlechta (1992) or the overview Makinson (1994) Observation 3.4.10 for the full story.

The question thus remains: can we characterize the family of all pivotal-valuation consequence operations in some other way, i.e. as the supraclassical closure operations satisfying disjunction in the premises plus some other property yet to be identified? As far as the author knows, this question has not been answered.

Generalizing to Ideals

We mention another, subtler, way of restricting the set of valuations. It is equivalent to pivotal-valuation consequence in the finite case, i.e. when the set of all valuations of the language is finite, but is a little more general in the infinite case. Given its abstract character, some readers may wish to skip directly to section 3.2.

Recall that we defined $A \mid\!\!-_W x$ to hold iff there is no valuation $v \in W$ such that $v(A) = 1$ whilst $v(x) = 0$. This is equivalent to saying: iff every valuation $v \in V$ such that $v(A) = 1$, $v(x) = 0$ is in $V\text{-}W$, i.e. iff $\{v \in V: v(A) = 1, v(x) = 0\} \subseteq V\text{-}W$, i.e. iff $\{v \in V: v(A) = 1, v(x) = 0\} \in 2^U$ where $U = V\text{-}W$. The set 2^U is a family of subsets of V with special properties; in particular, it contains the empty set, it contains all subsets of each of its elements, and is closed under the union of any two of its elements. We may thus generalize by introducing the abstract notion of an *ideal* over V, which is any family Δ of subsets of V such that $\emptyset \in \Delta$, $S \in \Delta$ whenever $S \subseteq T \in \Delta$, and $S \cup T \in \Delta$ whenever $S,T \in \Delta$.

An ideal is thus like a family of 'small' subsets, in so far as it contains the empty set and all of the subsets of any of its elements. But it is dissimilar in two ways: an ideal need not contain all the sets that are numerically as small as one of its elements (for example it may contain some singletons but not others), and it contains the union of any finite number of its elements, and so is at least as tall as it is wide. The notion of an ideal is dual to that of a filter, which is rather more familiar in logical studies. The definition below could be formulated equivalently, but rather less intuitively, in terms of filters.

Let Δ be any ideal over the set V of all valuations on the language L. Let A be any set of formulae, and let x be an individual formula. We say that x is a *consequence of A modulo the ideal* Δ, and write $A \mid\!\!-_\Delta x$ *alias* $x \in Cn_\Delta(A)$ iff $\{v \in V: v(A \cup \{\neg x\}) = 1\} \in \Delta$. We call a relation or operation a *pivotal-exception consequence* iff it coincides with $\mid\!\!-_\Delta$ (resp. Cn_Δ) for some ideal Δ over V. Note again that there is not a unique pivotal-exception consequence relation, but many – one for each value of Δ.

Any pivotal-valuation consequence operation is evidently a pivotal-exception one – given Cn_W, simply put $\Delta = 2^{V\text{-}W}$, and Δ will be an ideal over V with $Cn_\Delta = Cn_W$. Evidently, the converse also holds whenever the ideal Δ is *principal*, i.e. whenever there is a $U \subseteq V$ such that $\Delta = 2^U$. We can thus say that *the pivotal-valuation consequence operations are precisely the pivotal-exception ones generated by a principal ideal over V*.

Pivotal-exception consequence behaves very much like its less general sibling. For example, it satisfies disjunction in the premises, but not always compactness (by the same counterexample as for pivotal-valuation consequence in section 3.1). Indeed, the compactness of a pivotal-exception operation implies that it is a pivotal-valuation consequence. For by the representation theorem of section 2.1, such an operation will be a pivotal-assumption consequence and thus, as noted in section 3.1, a pivotal-valuation one.

3.2 From Pivotal Valuations to Default Valuations

As the reader will already have guessed from the earlier discussion of pivotal-assumption consequence, nonmonotonicity can arise from pivotal valuations if we allow the restricted set W of valuations to vary with the premise set A, giving us *default valuations*.

The best-known way of doing this was devised by Shoham (1988). His essential idea was to focus on those valuations satisfying the premise set A, that are *minimal* under some background ordering $<$ over the set of all valuations (not to be confused with the relation of plausibility or comparative expectation between formulae, which we considered when discussing variants of default-assumption consequence, and which we also wrote as $<$).

A *preferential model* is understood to be a pair $(W,<)$ where W is, as in the monotonic case, a set of valuations on the language L – not necessarily the entire set V – and $<$ is a relation over W. Given a preferential model $(W,<)$ we say that a formula x is a *preferential consequence* of a set A of formulae, and write $A \mid\!\sim_< x$ iff $v(x) = 1$ for every valuation $v \in W$ that is minimal among those in W that satisfy A. Thus the consequence relation depends on both W and $<$, and strictly speaking the snake sign should also carry W as a subscript; but we omit it to simplify notation. When preferential consequence is read as an operation rather than a relation, it is written as $C_<(A)$. Pivotal-valuation consequence $A \mid\!-_W x$ thus becomes the special case where $<$ is the empty relation so that all valuations in W are equally preferred.

It can be convenient to express the definition in a more compact manner. When W is a set of valuations and A is a set of formulae, write $|A|_W$ for the set of all valuations in W that satisfy A, i.e. $|A|_W = \{v \in W: v(A) = 1\}$. Write $min_<(|A|_W)$ for the set of all minimal elements of $|A|_W$. In this notation, $A \mid\!\sim_< x$ iff whenever $v \in min_<(|A|_W)$ then $v(x) = 1$, i.e. iff $min_<(|A|_W) \subseteq |x|_W$.

Note that in all these definitions, no constraints are placed on the relation $<$ over W. It need not be transitive or even irreflexive – much less complete. The concept of minimality makes perfect sense for an arbitrary relation. If R is any relation whatsoever over a set X, then for any $Y \subseteq X$ we take the minimal elements of Y under R to be the $y \in Y$ such that for no $y' \in Y$ do we have $(y',y) \in R$. The notion of a preferential model is thus well-defined without placing any constraints on its relation, but to guide intuitions it is useful to keep at the back of one's mind the typical case that it is both irreflexive and transitive.

Preferential consequence relations/operations are nonmonotonic. For a simple example, consider a language with at least three elementary letters p,q,r, and put $W = \{v_1, v_2\}$ where $v_1(p) = v_2(p) = 1$, $v_1(q) = 0$, $v_2(q) = 1$, $v_1(r) = 1$, $v_2(r) = 0$, and order W by putting $v_1 < v_2$. Informally, we describe this by saying: let p be true in both valuations, q true in just the top one, and r true in just the bottom one. Diagrammatically:

$$\begin{array}{l} \bullet\ v_2: p,q \\ | \\ \bullet\ v_1: p,r \end{array}$$

where we mention at each valuation point only the letters that are true there. Then $p \mid\!\sim_< r$ since the least valuation in which p is true is v_1, and r is also true there; but $p \wedge q \not\mid\!\sim_< r$, since the least valuation in which $p \wedge q$ is true is v_2, and r is false there.

Apart from failing monotony, preferential consequence relations are remarkably well behaved. They are supraclassical and satisfy disjunction in the premises. They also satisfy cumulative transitivity and inclusion (the last implied anyway by supraclassicality) and so are also idempotent. However, they lack some of the properties of their assumption-based analogues. We mention two important ones.

- They do not always satisfy cautious monotony, which is in effect, a converse of cumulative transitivity. In the language of consequence operations, cautious monotony says: $A \subseteq B \subseteq C(A)$ implies $C(A) \subseteq C(B)$. In terms of relations: whenever $A \mid\!\sim b$ for all $b \in B$ and $A \mid\!\sim x$ then $A \cup B \mid\!\sim x$. In the case that A,B are singletons it amounts to: whenever $a \mid\!\sim b$ and $a \mid\!\sim x$ then $a \wedge b \mid\!\sim x$.

Although this property may fail for preferential consequence relations, it always holds in the finite case. More generally, it holds whenever there are no infinite descending chains in the preferential model. More generally still, it holds whenever the model satisfies a condition known as *stoppering* (alias *smoothness*). This says that whenever $v \in |A|_W$ then either $v \in min_<(|A|_W)$ or there is a $u < v$ with $u \in min_<(|A|_W)$.

- Another property that preferential consequence relations may lack is *preservation of classical consistency*. This says that whenever A |~ f then A |- f (where f is a classical contradiction and |- is classical consequence). Equivalently: whenever $C(A)$ is classically inconsistent then so is A itself. The property holds for default-assumption consequences $|\!\sim_K$ (although we didn't mention it in the preceding section), but not in general for preferential consequences $|\!\sim_<$, essentially because some classical valuations may be missing from the model. When some of the valuations that make a proposition a true are missing from W, we can have $v(a) = 0$ for all $v \in W$ even though a is not a classical contradiction, so that also $min_<(|a|_W) = \emptyset$, vacuously giving $a \,|\!\sim_< f$.

  Nevertheless, for preferential models satisfying the stoppering condition, consistency is preserved with respect to a suitable supraclassical consequence operation. We may take this to be the pivotal-valuation consequence $|\!-_W$ determined by the set W of all valuations in the preferential model. In other words, for stoppered preferential models we *do have*: whenever $A \,|\!\sim_W f$ then $A \,|\!-_W f$.

Preferential consequence relations have been studied quite extensively. We do not wish to lose the reader in details; see for example Makinson (1994) for an overview. The main point of this section is that conceptually, monotonic pivotal-valuation consequences Cn_W serve as a natural bridge between classical Cn and nonmonotonic default-valuation consequences $C_<$. The definitions of Cn_W and $C_<$ both allow restriction of the set of classical valuations; the latter also allows this restriction to vary with the premises.

In contrast to the case of assumption-based consequences, the conceptual order also holds on the level of set-inclusion. We have $Cn \leq Cn_W \leq C_<$ when the preferential model determining $C_<$ is understood to be based on the same set W of valuations. This is for the simple reason that always $min_R(X) \subseteq X$, in particular $min_<(|A|_W) \subseteq |A|_W$.

3.3 Generalizations and variants

We end this account of default-valuation consequence by indicating very briefly some of the better-known generalizations of preferential consequence.

<u>Multiple Copies of Valuations</u>

It is evident that not all situations can be fully described in our meagre Boolean language, and we should therefore allow distinct possible worlds to satisfy exactly the same Boolean formulae. In other words, it is natural to permit as many 'copies' v_i of a valuation v as we like in a model. Technically, this can be done by indexing valuations v by elements i of an index set I. Equivalently, one may take valuations to be functions $v(a,i)$ of two arguments instead of one, where a ranges over formulae of the propositional language L as before, and the new argument i ranges over the index set I.

To be sure, such a step could already have been taken in the monotonic case, for pivotal-valuation semantics. But, as is almost immediate from the definitions, it would have made no difference there, generating exactly the same consequence relations with a more complex apparatus. In the nonmonotonic case, however, it makes a significant difference even for finite languages, because two copies of a single valuation may have quite different positions in the ordering. A simple example of this was given by Kraus, Lehmann and Magidor (1990), using the language with only two elementary letters. We add that, from a technical point of view, the availability of 'copies' of a valuation permits the proof of some important representation theorems that otherwise would not hold, at least not in their current simple forms.

With these considerations in mind, consider a set I of arbitrary items, which we call 'states', and an arbitrary relation over I, written as $<$. With each state $i \in I$ we associate a valuation v_i to the formulae of the language. The function (usually called the *labelling function*) that takes a world i to its associated valuation v_i need not be onto the entire set V of valuations (i.e. there may be some $v \in V$ such that $v \neq v_i$ for all $i \in I$), thus allowing us to work with a proper subset of the classical valuations. Nor need it be injective (i.e. we may have $v_i = v_j$ when $i \neq j$), thus offering the possibility of multiple 'copies' of a valuation.

We remark in passing that terminology is fickle here, as it varies from author to author, and even between the papers of a single author. Sometimes the elements i of I are called 'worlds'; at other times this term is used for the associated valuations v_i (or for the sets of formulae that they make true). This makes no difference when the labelling function is injective, as the elements of I may then be identified with the valuations that are their images under the labelling function; but it can cause confusion in the general, non-injective, case.

When multiple copies are allowed, a preferential model becomes a structure made up of a set I, a relation $<$ over I, and a labelling function associating a classical truth-value assignment v_i to each $i \in I$. The consequence operation is defined naturally by the rule: $A \mid\!\sim_< x$ iff $v_i(x) = 1$ for every state i that is minimal among those making A true. That is: iff $v_i(x) = 1$ for every $i \in min_<(|A|)$, where this time $|A| = \{j \in I: v_j(A) = 1\}$.

When so generalized, preferential consequence retains all those properties, positive and negative, that we have mentioned above for the simpler version without copies. It has the intuitive advantage of better allowing for the complexity of the universe, and the formal advantage of admitting smoother statements and proofs of representation theorems, as shown in detail by Moinard and Rolland (2002). *Indeed, it is this form of preferential consequence that has become the 'industry standard'*, and is often known as KLM consequence after the seminal paper of Kraus, Lehmann and Magidor (1990).

Selection Functions

Another line of generalization abstracts on the preference relation over W (or over I, if we are allowing copies), working instead with a selection function γ over the power set of W, with $\gamma(U) \subseteq U$ for every $U \subseteq W$. This selection function may be constrained by suitable conditions as desired. If sufficiently powerful conditions are imposed, then a preference relation can be recovered from the selection function. The relation so recovered is then sometimes called a 'revealed preference', a term borrowed from the theory of choice between goods in economics.

This approach is investigated in detail, for both nonmonotonic inference and belief revision, by Lindström (1991), Rott (2001) and Lehmann (2001).

<u>Non-Classical Valuations</u>

In all of the constructions of this section, we have been dealing with classical valuations, i.e. those functions on the set of formulae into {0,1} that behave in the usual Boolean fashion. It is possible to relax this condition, either partially or entirely. A partial relaxation was studied by Kraus, Lehmann and Magidor (1990). A more complete relaxation was investigated by Makinson (1989) and more recently Lehmann (2002), taking the valuations v to be arbitrary functions on the set of formulae into {1,0}.

<u>Working with Ideals</u>

When we restrict valuations by means of ideals, giving us the notion of pivotal-exception consequence on the monotonic level, we may pass to the nonmonotonic level by allowing the identity of the ideal Δ to vary with the premise set. More specifically, instead of taking Δ to be an ideal over the entire set V, we may take it to be one over the set $|A|$ of all valuations satisfying the premise set A.

Consider any function δ that associates with each set $|A|$ an ideal over $|A|$. Then we may say that x is a *default-exception consequence of A modulo* δ, and write $A \mid\!\sim_\delta x$ *alias* $x \in Cn_\delta(A)$ iff the set of valuations that satisfy A but not x is an element of the ideal $\delta(|A|)$ over $|A|$, i.e. iff $\{v \in V: v(A \cup \{\neg x\}) = 1\} \in \delta(|A|)$.

The relations $\mid\!\sim_\delta$ so defined are supraclassical but nonmonotonic. However, because the ideals associated with different sets $|A|$ need have no relationship to each other, properties such as cut and disjunction in the premises may also fail. In order to recover some or all of them, it is customary to add constraints on the relationship between the ideals $\delta(|A|)$ for different values of A. Usually this is done directly, but one might also reintroduce a relation $<$ over V and relate the ideals to $<$. We shall not enter into the details, referring the reader to the papers mentioned below.

The default-exception approach was developed, in a number of versions and with various names (e.g. 'systems based on comparative size') by Ben-David & Ben-Eliyahu-Zohary (2000), Schlechta (1997), and Friedman & Halpern (2001). A detailed comparison with the default-valuation approach (in a version using selection functions and multiple copies of valuations) was made by Lehmann (2001).

3.4 Recapitulation

The main point that we want to emphasize in this section is that once again, a well-known kind of nonmonotonic consequence relation (preferential consequence) emerges from a monotonic one (pivotal-valuation consequence) by allowing its key ingredient (the set W of valuations allowed) to vary with the premise sets A, the variation being controlled by minimization under a relation. The monotonic operation thus serves as an antechamber to its nonmonotonic counterpart.

4 Third Bridge - Using Additional Rules

4.1 From Classical Consequence to Pivotal Rules

We come now to a third kind of monotonic system serving as a hinge between classical consequence and a well-known kind of nonmonotonic one. The basic idea is similar to that of adding background assumptions, but instead of adding *propositions* we add *rules*. This apparently small difference brings with it considerable divergence, even in the finite case where (as noted towards the end of section 3.1) the pivotal-assumption and pivotal-valuation approaches are equivalent. Rules for propositions do not behave quite like propositions.

By a rule for propositions (briefly, a *rule*) we mean any ordered pair (a,x) of propositions of the language we are dealing with. A set of rules is thus no more nor less than a binary relation R over the language, i.e. a set $R \subseteq L^2$. Given a set X of propositions and a set R of rules, we define the *image* of X under R, written as $R(X)$, in the standard manner of elementary set theory: $y \in R(X)$ iff there is an $x \in X$ with $(x,y) \in R$. A set X is said to be closed under R iff $R(X) \subseteq X$, i.e. iff whenever $x \in X$ and $(x,y) \in R$ then $y \in X$.

Let $R \subseteq L^2$ be a set of rules, which intuitively will be playing the role of a set of background 'inference tickets' ready for application to any set of premises. Let A be a set of formulae, and let x be an individual formula.

We say that x is a *consequence of A modulo the rule set R*, and write $A \mid\!\!-_R x$ *alias* $x \in Cn_R(A)$ iff x is in every superset of A that is closed under both Cn and the rule set R. In other words, iff x is in every set $X \supseteq A$ such that both $Cn(X) \subseteq X$ and $R(X) \subseteq X$.

At least one such set always exists, namely the entire language L. Moreover, the intersection of any non-empty family of sets, each of which is closed under Cn (resp. R), is itself closed under Cn (resp. R). It follows that there is always a unique least such set. Thus the definition may be put as follows: $Cn_R(A)$ is the least superset X of A such that both $Cn(X) \subseteq X$ and $R(X) \subseteq X$.

The idea of using $Cn_R(A)$ as a stepping stone towards default rule consequence is implicit in many papers, but rarely brought to the surface. An early paper that does make it explicit is Sandewall (1985). Sandewell presents his account in the guise of a four-valued logic, but it can be translated into the language of sets of formulae. However the manner in which he then passes from the bridge system to the default one is quite different from that to be introduced here, and much closer to Reiter's own fixed-point procedure.

We call an operation a *pivotal-rule consequence* iff it is identical with Cn_R for some set R of rules. Note once again that there is not a unique such relation, but many – one for each value of R.

It is immediate from the definition that $Cn \leq Cn_R$ so that pivotal-rule consequence Cn_R is supraclassical. It is not difficult to verify that it is monotonic and also satisfies cumulative transitivity, so that it is a closure operation. We are thus still in the realm of paraclassical inference.

Like pivotal-assumption, and unlike pivotal-valuation consequence, it can also be verified to be compact. But there is an important property possessed by both pivotal-assumption and pivotal-valuation operations that it lacks - disjunction in the premises. For example, if $R = \{(a,x), (b,x)\}$ then $x \in Cn_R(a)$ and also $x \in Cn_R(b)$ but $x \notin Cn_R(a \vee b) = Cn(a \vee b)$. This is because the last-mentioned set is vacuously closed under R, i.e. we have $R(Cn(a \vee b)) = \emptyset \subseteq Cn(a \vee b)$ since $a,b \notin Cn(a)$.

It also fails contraposition. This follows from the failure of disjunction in the premises, since any supraclassical closure operation satisfying the former must satisfy the latter. For a direct counterexample, put $A = \{a\}$ and $R = \{(a,x)\}$. Then $x \in Cn_R(a)$ but $\neg a \notin Cn_R(\neg x) = Cn(\neg x)$, because the last-mentioned set is vacuously closed under R, i.e. we have $R(Cn(\neg x)) = \emptyset \subseteq Cn(\neg x)$ since $a \notin Cn(\neg x)$. This contrasts with the situation for pivotal-assumption and pivotal-valuation consequence. They both satisfy contraposition although, it should be added to avoid confusion, their default versions do not.

Both of these examples drive home the fact that pivotal-rule consequence cannot be reduced to pivotal-assumption consequence by transforming the rules $(a,x) \in R$ to their materializations, i.e. to the truth-functional (*alias* material) propositions $a \rightarrow x$. In particular, we do not have $Cn_R(A) = Cn(A \cup m(R))$ where $m(R)$ is the set of materializations of pairs in R. We always have the inclusion $Cn_R(A) \subseteq Cn(A \cup m(R))$, but the converse inclusion $Cn(A \cup m(R)) \subseteq Cn_R(A)$ will usually fail. In the first example above, while $x \notin Cn_R(a \vee b)$ we do have $x \in Cn(\{a \vee b, a \rightarrow x, b \rightarrow x\})$; and in the second example, while $\neg a \notin Cn_R(\neg x)$ we do have $\neg a \in Cn(\{\neg x, a \rightarrow x\})$

The relationship between the set of pivotal-assumption consequence operations and the set of pivotal-rule ones is analogous to that between the former and the pivotal-valuation relations studied in section 3. *The pivotal-assumption consequence operations are precisely the pivotal-rule ones that satisfy disjunction in the premises*. For on the one hand, every pivotal-assumption consequence operation is itself a pivotal-rule one – given Cn_K, simply put $R = \{(t,k): k \in K\}$ where t is a tautology, and it is easily checked that $Cn_K = Cn_R$. Conversely, when Cn_R is a pivotal-rule consequence then as we have shown it is a compact supraclassical closure operation; so when in addition it also satisfies disjunction in the premises, then by the representation theorem of section 2.1 it is also a pivotal-assumption consequence operation.

As a curiosity, we thus have the following relationship between the three: *the set of pivotal-assumption consequence operations is the intersection of the set of pivotal-valuation and the set of pivotal-rule consequence operations*. Indeed, we have already shown that the left set is included in each of the two right ones, and when an operation is in both of the two right sets we can again apply the representation theorem of section 2.1.

We can also prove a representation theorem for pivotal-rule consequence. It says: *for every compact supraclassical closure operation Cn', there is a relation R such that for all A, $Cn'(A)$ is the least superset of A that is closed under both classical Cn and R.*

- *Construction and proof.* The construction is trivial, the proof easy. For the construction, simply put R to be Cn' itself, restricted to singletons, i.e. put $R = \{(a,x)$: x

$\in Cn'(\{a\})$. The proof is in two parts. First we show that $Cn'(A)$ is a superset of A and is closed under both classical Cn and R. Then we show that it is the least such superset.

For the first part, we have $A \subseteq Cn'(A)$ since Cn' is a closure operation. Also $Cn(Cn'(A)) \subseteq Cn'(A)$ since by supraclassicality and idempotence of Cn' we have $Cn(Cn'(A)) \subseteq Cn'(Cn'(A)) \subseteq Cn'(A)$. Again, $R(Cn'(A)) \subseteq Cn'(A)$ by the definition of R together with the monotony and idempotence of Cn'. This shows that $Cn'(A)$ is a superset of A that is closed under both classical Cn and R.

To show that it is the least such superset, let X be another one; we need to check that $Cn'(A) \subseteq X$. But if $x \in Cn'(A)$ then by compactness of Cn' there is a finite $B \subseteq A$ with $x \in Cn'(B)$. Also $Cn'(B) = Cn'(b)$ where b is the conjunction of all the finitely many elements of B, using supraclassicality and the supposition that Cn' is a closure operation. Since $x \in Cn'(b)$ we have by the definition of R that $(b,x) \in R$, so $x \in R(\{b\}) \subseteq R(Cn(\{b\})) = R(Cn(B)) \subseteq R(Cn(A)) \subseteq R(Cn(X)) = R(X) = X$ since by hypothesis X includes A and is closed under both R and classical Cn.

The above construction works only because there are no constraints placed on the relation R in the definition of a pivotal-rule consequence relation. This allows us to 'cheat' by taking R to be the restriction to singletons of the very consequence relation that we are seeking to represent. Evidently, it would be more interesting to have representation results when R is constrained in some natural way. This is an area that does not seem to have been explored.

To end this section, we note a point that will be vital when we pass to default-rule consequence in the next section. It is possible to reformulate the definition of pivotal-rule consequence in an inductive manner. We simply put $Cn_R(A) = \cup\{A_n : n < \omega\}$ where $A_1 = A$ and $A_{n+1} = Cn(A_n \cup R(A_n))$. The equivalence of this with the original definition is easy to verify using the compactness of classical Cn.

We can indeed go further and reformulate the inductive definition so that it has only singleton increments in the inductive step. To do this, fix an ordering $<R>$ of the set R of rules - that is, fix a sequence $(a_i,x_i)_{i<\alpha}$ of all the rules in R, without repetitions, with α a positive integer or ω according to the cardinality of R (which we assume to be finite or countable). Then put $Cn_{<R>}(A) = \cup\{A_n : n < \omega\}$ where $A_1 = A$ as before but now $A_{n+1} = Cn(A_n \cup \{x\})$, where (a,x) is the *first* rule in $<R>$ such that $a \in A_n$ but $x \notin A_n$. In the case that there is no such rule, then we put $A_{n+1} = Cn(A_n)$.

To avoid a proliferation of symbols, we are using the same notation A_n for the terms of this sequence as for the previous sequence, but evidently it is quite a different one, making much smaller jumps with singleton increments. It should be noted that the construction does not imply that rules are applied in the order in which they occur in $<R>$. When we form A_{n+1}, the rule (a,x) that we apply may occur earlier in the sequence $<R>$ than some rule (b,y) that we have already applied, since the passage from A_{n-1} to A_n may introduce the body a for the first time. Also note that by the terms of the definition, once a rule (a,x) is applied, it is never applied again, since its head x is in all the subsequent sets A_{n+k}.

Although the terms A_n of the second sequence are not the same as those of the first sequence, their union is the same, i.e. $Cn_{<R>}(A) = Cn_R(A)$. This implies also that the choice of ordering makes no difference to the final result.

These inductive definitions of pivotal-rule consequence may seem like long ways of expressing something that the original definition says quite briefly. And so they are. But we will see in the next section that they – and especially the singleton-increment one – provide the key for passing to the nonmonotonic operations of default-rule consequence.

4.2 From Pivotal Rules to Default Rules

The consequence operations Cn_R defined by using pivotal rules are, as we have seen, monotonic. Once again we may obtain nonmonotonicity by allowing the set R of rules to vary with A, or more precisely, by requiring the set of those rules in R that are actually applied in the inductive construction to vary with A. This is done by imposing consistency constraints on their application.

There are a number of different ways of going about this. We will begin by describing one of the best known, due to Reiter (1980), referred to as his system of 'normal defaults'. In the next section we will sketch some generalizations.

We take as our starting point the inductive definition of pivotal-rule consequence, using singleton increments. The additional idea is to allow rules to be applied one by one, under the guidance of the ordering $<R>$, so long as the application does not generate inconsistency.

As before, we fix an ordering $<R>$ of the given set R of rules. As before, we put $C_{<R>}(A) = \cup\{A_n : n < \omega\}$, and set $A_1 = A$, but the subsequent terms are defined differently. We put $A_{n+1} = Cn(A_n \cup \{x\})$, where (a,x) is the *first* rule in $<R>$ such that $a \in A_n$ but $x \notin A_n$ *and x is consistent with* A_n. In the case that there is no such rule, then we put $A_{n+1} = Cn(A_n)$.

The introduction of the consistency constraint in the induction step has many consequences. One is that the identity of $C_{<R>}(A)$ will now vary with the particular ordering $<R>$ of R. For example, if $A = \{a\}$ and $R = \{(a,x), (a,\neg x)\}$, then if we order R in one way we get $C_{<R>}(A) = Cn(\{a,x\})$ while if we take the reverse order we get $C_{<R>}(A) = Cn(\{a,\neg x\})$. We thus have not one but many sets $C_{<R>}(A)$, one for each choice of the ordering $<R>$. Another consequence is that these operations are nonmonotonic. For example, if $A = \{a\}$ and $R = \{(a,x)\}$ then $C_{<R>}(A) = Cn(\{a,x\})$, but when A is increased to $B = \{a,\neg x\}$ then $C_{<R>}(B) = Cn(B) = Cn(\{a,\neg x\})$.

The sets $C_{<R>}(A)$ for varying orderings $<R>$ of R coincide with what Reiter calls *extensions* of A under the (normal) rules R. Each ordering uniquely determines some extension, and every extension is determined by some ordering. The way in which we have defined extensions is not the same as in Reiter's original paper, for we wish to bring out the natural transition through pivotal-rule consequence to default-rule consequence. We also wish to show how they may be defined by standard mathematical induction rather than end-regulated induction or fixpoints. However, the definitions are equivalent. The basic idea for the present one is the same as that for 'prioritized default logic' (*alias* PDL) in the form in which it is set out in Antoniou (1997).

The operations $C_{<R>}$ are themselves of interest, and if one is equipped with a preferred ordering $<R>$ of R, one may wish to go no further. But one can also define a consequence C_R that is independent of any particular order, by intersecting all the $C_{<R>}$. In other words, we may put $C_R(A) = \cap\{C_{<R>}(A)$: $<R>$ an ordering of $R\}$. In the relational notation: $A \mid\!\sim_R x$ iff $A \mid\!\sim_{<R>} x$ for every possible ordering $<R>$ of the rule-set R. We call this relation/operation *default-rule consequence*. It coincides with what is usually known as 'consequence using normal Reiter default rules with the sceptical policy on extensions'.

We summarise the main message. Nonmonotonic default-rule consequence operations $C_{<R>}$ emerge by fixing an ordering $<R>$ of R and requiring the application of the rules in $<R>$ to depend on a consistency constraint, whose effect will vary according to the initial premise set A. If desired, these nonmonotonic operations $C_{<R>}$ may be compressed into a single operation C_R, still nonmonotonic, by simple intersection of values. The monotonic operations $Cn_{<R>}(A)$ alias $Cn_R(A)$ thus serve as stepping-stones to the nonmonotonic ones.

As with the assumption-based consequences (and unlike the valuation-based ones) this conceptual order is reversed on the level of set-inclusion, where we have $Cn \leq C_R \leq C_{<R>} \leq Cn_{<R>} = Cn_R$, as can be verified by easy inductions.

4.3 Generalizations and Variants

We outline briefly some generalizations and a variant of default-rule consequence defined above.

<u>Non-normal Default Rules</u>

In his paper of 1980, Reiter also considered default rules of a much more general kind, usually called 'non-normal' ones. They are no longer pairs (a,x) but triples (a,p,x) of boolean formulae, usually written $(a;p/x)$, or even a little more generally, triples (a,P,x), written $(a;P/x)$, where P is a finite set of formulae. The elements of P are usually called the 'justifications' of the rule. The idea is to pass from a to x whenever each element of P taken separately, as well as each justification in any previously applied rule, is consistent with whatever we have constructed so far, plus x. Normal rules (a,x) may be identified with triples (a,P,x) where P is empty.

This is a notoriously intricate notion, and admits many variations. They are usually defined in terms of 'fixpoints' of a function or in terms of 'end-regulated induction' - see the overview in Delgrande, Schaub and Jackson (1994) or presentations in the books of Lukaszewicz (1990), Marek, V.W. and M. Truszczynski (1993), and Antoniou (1997). However, it is also possible to cover non-normal rules by a standard mathematical induction using orderings $<R>$ of the set R of rules, by suitable refinements of the definition for normal rules from the preceding section.

We give first the simplest such definition, which yields extensions close to those of Lukaszewicz (1990), then show how it may be further modified to give the extensions of Reiter (1980). In both cases, the definition coincides with that of the previous section when the set P of justifications is the singleton of the conclusion, in all rules.

For the simpler system, we again fix an ordering $<R>$ of the given set R of rules, and put $C_{<R>}(A) = \cup\{A_n : n < \omega\}$. The definition of the sequence $A_1, A_2, \ldots$ is like that for normal rules, except that we must accompany each A_n by a 'book-keeping' set Q_n that accumulates the justifications so far used, and check on Q_n as well as on the justifications of the current rule when defining A_{n+1}. These sets Q_n may, if one wishes, be thought of as 'labels' on the corresponding sets A_n, and the definition may be thought of as setting up a 'labelled deductive system' in the sense of Gabbay.

The basis has two components: $A_1 = A$, $Q_1 = \emptyset$. For the principal component of the induction step, we again put $A_{n+1} = Cn(A_n \cup \{x\})$. But this time x is the conclusion of the first rule (a,P,x) in $<R>$ such that $a \in A_n$, $x \notin A_n$ and x is consistent with $A_n \cup \{p\}$ for each $p \in Q_n \cup P$. In words: x is the conclusion of the first rule whose premise is in the current set, whose conclusion is not in the current set, and whose conclusion, together with the current set, is classically consistent with each current or previously used justification taken one by one. In the case that there is no such rule, we again put $A_{n+1} = Cn(A_n)$. For the book-keeping component of the induction step we simply put $Q_{n+1} = Q_n \cup P$.

This does not coincide fully with the Reiter definition. The sets A_n are clearly well-defined no matter how we choose A, R, and its ordering $<R>$. In contrast, as is notorious, there are cases in which a set A has no Reiter extensions under R at all. For example, when $A = \{a\}$ and R consists of the single triple $(a,\neg x,x)$, then we have no Reiter extensions, while $C_{<R>}(A)$ as defined above is $Cn(a)$. Roughly speaking, this definition treats a rule simply as inapplicable in certain situations where, under Reiter's vision, the construction aborts.

Reiter's notion of an extension may be obtained by suitably modifying the above definition at the specification of A_{n+1}. We take the first rule (a,P,x) in $<R>$ such that $a \in A_n$, $x \notin A_n$, and A_n is consistent with each separate $p \in P$, and consider two cases. *Case 1*: Suppose that also x is consistent with $A_n \cup \{p\}$ for each separate $p \in Q_n \cup P$. Then we put $A_{n+1} = Cn(A_n \cup \{x\})$. *Case 2*: Suppose otherwise. Then the construction aborts: A_{n+1} is left undefined and we continue no further, taking also $C_{<R>}(A)$ to be undefined.

We emphasize that this way of defining Reiter extensions under non-normal rules is *still a standard mathematical induction*, notwithstanding the clause that allows the construction to abort. In effect, it gives an inductive definition of a partial (rather than a total) function. It is not a fixed-point definition, nor one by end-regulated induction, as are the definitions of Reiter extensions that are familiar from the literature on default rules. Indeed, as far as the author knows, this is the first purely inductive formulation of the concept to be given.

<u>Maxfamilies of Rules</u>

A rather different approach to default rules, which will surely have occurred to readers who remember how we defined default-assumption consequence in section 2, is to take maximal subsets of the set R of rules.

We go back to the simple notion of a rule as a pair (a,x) of Boolean formulae, i.e. without introducing a 'justification set' P. Given a set R of rules, and a premise set A, we consider the family of all maximal subsets $S \subseteq R$ such that $Cn_S(A)$ is consistent. This is called the *maxfamily* of A under R. We define $C_{\underline{R}}(A)$ (this time with subscript underlined) to be the intersection of the sets $Cn_S(A)$ for all S in the maxfamily of A.

These operations $C_{\underline{R}}$ are not quite the same as the default-rule operations C_R of the preceding section. An example is given by the well-known Möbius strip, where $R = \{(a,x), (x,y), (y,\neg a)\}$. Here the premise set $\{a\}$ gives rise to only one extension, namely $Cn(\{a,x,y\}$, while we have three maxfamilies S, namely the three two-element subsets of R, which give rise to three sets $Cn_S(A)$, namely $Cn(\{a\}$, $Cn(\{a,x\}$, $Cn(\{a,x,y\})$. In general, there are more maxfamilies than extensions, and so the intersection of the maxfamilies will be smaller than the intersection of the extensions.

The difference between the default-rule consequence operations C_R and the maxfamily consequence operations $C_{\underline{R}}$ is subtle. But it is also important. It is in effect a difference between two ways of guarding against contradictions when applying rules to a set of premises: by restricting the *generating process* (giving us the operations $C_{<R>}$ for orderings $<R>$ of R), or by constraining the *generating apparatus* (giving us the operations $C_S(A)$ for sets $S \subseteq R$ in the maxfamily of A).

Maxfamily operations are studied by Makinson and van der Torre (2000) in a broader context where inputs (premises) are not necessarily authorized to reappear as outputs (conclusions).

4.4 Recapitulation

The principal message of this section is that by working with any fixed set of rules we obtain a natural supraclassical consequence, which is monotonic and indeed a closure operation, and which we have called a pivotal-rule consequence. By allowing the set of rules to vary with the premises of the inference – more precisely, by placing a suitable consistency constraint on their step-by-step application – we also obtain quite naturally the well-known normal default consequence operations of Reiter, as well as their generalizations to non-normal rules and a range of variants.

5 Conclusion

We have studied three main ways of obtaining more from a set of premises than is classically authorized: by using background *assumptions* that are added to the premises of each inference, by restricting the set of *valuations* of the premises, and by deploying additional *rules* alongside classical consequence.

These generate three kinds of *monotonic* relation, which we have called *pivotal-assumption*, *pivotal-valuation*, and *pivotal-rule* relations. The first two coincide in the finite case, but can differ in the infinite case; the third is significantly different even in the finite case. Each is of interest in its own right. They are all *paraclassical*, that is, supraclassical closure relations.

Each of these kinds of paraclassical logic serves as an antechamber for the definition of a well-known kind of *nonmonotonic* consequence relation. This is done by allowing the background assumptions, valuations, or rules to sway in a systematic way with the premises under consideration. In this way we obtain *default-assumption*, *default-valuation*, and *default-rule* consequences. They coincide with well-known systems of inference using Poole systems, preferential models, and Reiter default rules. By generalizing and varying the constructions, one also obtains a range of other kinds of nonmonotonic consequence that are known in the literature.

The take-home story is in Table 1 below.

TABLE 1: Classical, Paraclassical and Nonmonotonic Consequence Operations

<table>
<tr><td colspan="3">Classical consequence Cn</td></tr>
<tr><td>Use additional assumptions
Pivotal-assumption consequence operations Cn_K
Paraclassical
- Disjn in premises: yes
- Compact: yes</td><td>Restrict the set of valuations
Pivotal-valuation consequence operations Cn_W
Paraclassical
- Disjn in premises: yes
- Compact: no</td><td>Use additional rules
Pivotal-rule consequence operations Cn_R
Paraclassical
- Disjn in premises: no
- Compact: yes</td></tr>
<tr><td>Vary additional assumptions with the premises
- e.g. using consistency constraint
Default-assumption consequence operations C_K
(Poole systems and variants)</td><td>Vary valuation set with the premises
- e.g. using minimalization under a relation
Default-valuation consequence operations C_W
(preferential systems and variants)</td><td>Vary additional rules with the premises
- e.g. using consistency constraint
Default-rule consequence operations C_R
(Reiter systems and variants)</td></tr>
</table>

References

Alchourrón, Carlos and David Makinson 1982. ‘On the logic of theory change: contraction functions and their associated revision functions’, *Theoria* 48: 14-37.

Alchourrón, Carlos, Peter Gärdenfors and David Makinson 1985. ‘On the logic of theory change: partial meet contraction and revision functions’, *The Journal of Symbolic Logic* 50: 510-530.

Antoniou, Grigoris 1997. *Nonmonotonic Reasoning*. MIT Press: Cambridge Mass.

Ben-David, Shai and Rachel Ben-Eliyahu-Zohary 2000. ‘A modal logic for subjective default reasoning’, *Artificial Intelligence* 116: 217-236.

Delgrande, J.P., T. Schaub and W.K. Jackson 1994. ‘Alternative approaches to default logic’ *Artificial Intelligence* 70: 167-237.

Forget, Lionel, Vincent Risch and Pierre Siegel 2001. 'Preferential logics are *X*-logics' *J. Logic and Computation* 11: 71-83.

Friedman, Nir and Joseph Y. Halpern 2001. ‘Plausibility measures and default reasoning’, *Journal of the ACM* 48: 648-685.

Gärdenfors, Peter and David Makinson 1988. ‘Revisions of knowledge systems and epistemic entrenchment’, in M.Vardi ed., *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*. Los Altos: Morgan Kaufmann, 1988, pp 83-95.

Gärdenfors, Peter and David Makinson 1994. ‘Nonmonotonic inference based on expectations’, *Artificial Intelligence* 65: 197-245.

Kraus, Lehmann and Magidor 1990. ‘Nonmonotonic reasoning, preferential models and cumulative logics’, *Artificial Intelligence* 44: 167-207.

Lehmann, Daniel 2001. ‘Nonmonotonic logics and semantics’, *Journal of Logic and Computation* 11: 229-256.

Lehmann, Daniel 2002. ‘Connectives in quantum and other cumulative logics’, Technical Report TR-2002-28 of the Leibniz Center for Research in Computer Science, Hebrew University of Jerusalem.

Levi, Isaac 1996. *For the Sake of the Argument*. Cambridge University Press.

Lindström, Sten 1991. ‘A semantic approach to nonmonotonic reasoning: inference operations and choice’, Uppsala Prints and Reprints in Philosophy n°6, Department of Philosophy, University of Uppsala.

Lukaszewicz, W. 1990. *Non-Monotonic Reasoning – Formalization of Commonsense Reasoning*. Ellis Horwood.

Makinson, David 1989. ‘General theory of cumulative inference’, in *Nonmonotonic Reasoning* ed. M. Reinfrank et al. Berlin: Springer-Verlag, Lecture Notes on Artificial Intelligence n°346, 1989, 1-17.

Makinson, David 1994. ‘General Patterns in Nonmonotonic Reasoning’, in *Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3*, ed. Gabbay, Hogger and Robinson. Oxford University Press, pages 35-110.

Makinson, David and Leendert van der Torre 2000. ‘Input/output logics’, *Journal of Philosophical Logic* 29: 383-408.

Marek, V.W. and M. Truszczynski 1993. *Nonmonotonic Logic: Context Dependent Reasoning*. Springer, Berlin.

Moinard, Y. and R. Rolland 2002. 'Characterizations of preferential entailments' *Logic Journal of the IGPL* 10: 243-270.

Poole, David 1988. ‘A logical framework for default reasoning’, *Artificial Intelligence* 36: 27-47.

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

Rott, Hans and Maurice Pagnucco 1999. ‘Severe withdrawal - and recovery’, *Journal of Philosophical Logic* 28: 501-547 (reprinted with corrections of corrupted text in *ibid* 2000, 29:121 ff).

Rott, Hans 2001. *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*, Clarendon Press, Oxford UK (Oxford Logic Guides n ° 42).

Sandewall, Erilk 1985. 'A functional approach to non-monotonic logic' *Computational Intelligence* 1: 80-87. Also appeared as Research Report LITH-IDA-R-85-07 of the Department of Computer and Information Science, Linköping University, Sweden.

Schlechta, Karl 1992. ‘Some results on classical preferential models’, *Journal of Logic and Computation* 2: 676-686.

Schlechta, Karl 1997 *Nonmonotonic Logics: Basic Concepts, Results and Techniques*. Springer, Berlin: LNAI Series n° 1187.

Shoham, Yoav 1988. *Reasoning About Change*. MIT Press, Cambridge USA.

Siegel, Perre and Lionel Forget (1996). 'A representation theorem for preferential logics', *Fifth Conference on Principles of Knowledge Representation (KR-96)* ed. L.Aiello et al. Cambridge: USA.

Acknowledgements

Thanks to Björn Bjurling, Michael Freund, Daniel Lehmann, Karl Schlechta for valuable comments on drafts, and the referee of the *LJIGPL* for suggestions that improved the final version.

Correction

This text corrects an oversight in the published version, where it was claimed that the simpler inductive definition of an extension using non-normal rules gives exactly the same extensions as the fixpoint system of Lukaszewicz. In fact, there is a difference. On the other hand, the inductive definition for Reiter's system, does coincide with it as claimed. For a proof of this point and further observations and discussion, see the author's book *Bridges from Classical to Nonmonotonic Logic* (2005)