THE JOURNAL OF SYMBOLIC LOGIC
Volume 00, Number 0, XXX 0000

IMPLICATION WITH POSSIBLE EXCEPTIONS

HERMAN JURJUS AND HARRIE DE SWART

Abstract. We introduce an implication-with-possible-exceptions and define validity of rules-with-possible-exceptions by means of the topological notion of a full subset. Our implication-with-possible-exceptions characterises the preferential consequence relation as axiomatized by Kraus, Lehmann and Magidor [Kraus, Lehmann, and Magidor, 1990]. The resulting inference relation is non-monotonic. On the other hand, modus ponens and the rule of monotony, as well as all other laws of classical propositional logic, are valid-up-to-possible exceptions. As a consequence, the rules of classical propositional logic do not determine the meaning of deducibility and inference as implication-without-exceptions.

§0. Introduction. Although non-monotonic logic arose as a subdiscipline of artificial intelligence, the potential of the subject for the field of formal logic should not be underestimated, as non-monotony is connected with the most important difficulties in formalizing practical reasoning. Practical reasoning seems to be governed by general principles that function as rules-with-possible-exceptions. Hence, mathematical axioms, when interpreted in the usual way, will not suffice to describe the underlying calculus.

In this paper, the expressive power of the language used to formulate principles of reasoning has been extended to enable the interpretation of formal axioms as valid-up-to-possible-exceptions. The core of the system is the notion of a *full subset*, a topological generalization of intuitions stemming from Euclidean geometry (a subject that preceded Aristotelean logic). This notion gives rise to a topological semantics of *defeasible implication* in a natural way. A sound and complete axiomatization is established, relying for the proof on a result of [Kraus, Lehmann, and Magidor, 1990] (in fact, we find the same axiom system).

The notion of a full subset is used to define semantics of nested implicational statements. This semantics is associated with two issues. In the first place, we think of an imaginary person that handles (some) rules of inference as rules-with-possible-exceptions. With the semantics defined in this paper, the person turns out to obey-up-to-possible-exceptions every law of classical propositional logic, including modus ponens and the rule of monotony. Thus, it is shown that it is possible for a nonmonotonic formalism to obey-up-to-possible-exceptions the rule of monotony. A much more interesting conclusion is that the rules of classical propositional logic do apparently not determine the usual interpretation of implication (and inference) as implication-without-exceptions.

Received July 6, 1999.

In the second place, the formalism can be seen as extending classical propositional logic with a binary connective denoting "defeasible implication". Seen in this way, the non-monotony of the inference relation has nothing to do with "non-deductive reasoning" or "jumping to conclusions", but is a necessary consequence of the presence of a defeasible conditional in the language; correct, deductive reasoning from defeasible premisses *is* non-monotonic. (More provocatively: correct, deductive reasoning in general is non-monotonic.)

In the literature on non-monotonic logic, non-monotonic formalisms are usually thought of as models of certain types of human reasoning that deviate rather drastically from correct reasoning, but that are remarkably efficient and useful for certain practical purposes. Non-monotonic formalisms, then, are acknowledged to be models of incorrect reasoning. One of the primary motivations for this paper is to take non-monotony more seriously. This can be done in several ways. One possibility is to point out some blind spots of classical logic, formalize the neglected phenomena, and find out that the resulting formalism is non-monotonic. After all, classical logic (say, first order predicate logic) does in fact neglect some phenomena typical for human reasoning that could hardly be called irrational or incorrect and some of these phenomena would lead to non-monotonic formalisms.

Another possibility, not necessarily incompatible with the former, is to realize that a formal system, such as predicate logic, is best seen as a mathematical model (of something that we have called a type of reasoning). But a model is like a caricature, emphasizing certain aspects, neglecting others. And it is not uncommon for an object to have two different models. This leads us to the first main point of departure of this paper: *we want to view both classical logic and non-monotonic logic as caricatures of the same type of reasoning*. Consequently, we will restrict our attention to those forms of non-monotony that are compatible with this point of view.

The type of non-monotony that we want to study is best illustrated by the following example. *John is always home at 6 o'clock, it's 6 o'clock now, and John is not home (now).*

What conclusions can be drawn from this combination of assumptions?

Let us rephrase "John is always home at 6 o'clock" as "if it is 6 o'clock, then John is home." Then the assumptions may be written as $p \to q$, p, $\neg q$, where p denotes "it is 6 o'clock now" and q denotes "John is home now."

If we interpret implication ("$\to$") as implication-without-possible-exceptions, as is usual in mathematics, this set of assumptions is contradictory, and we may conclude anything from it, including "I am the emperor of China," as well as "John is home now." However, in everyday life, we would usually take into account the possible failure of "John is always home at 6 o'clock" in exceptional cases. Hence, we would not consider "John is home now" to be a sensible conclusion to those three assumptions. In particular, the combination of assumptions is not inconsistent. Symbolically represented, classical logic accepts the following argument:

$$\dfrac{\dfrac{\dfrac{p \to q, p, \neg q}{p \to q, p}}{q} \qquad \neg q}{\perp}$$

However, with a less pedantic interpretation of implication *as well as inference*, we could accept

$$\frac{p \to q, p}{q} \text{ as valid, but not } \frac{p \to q, p, \neg q}{q}$$

(which shows that the resulting inference relation would be non-monotonic), nor $\dfrac{p \to q, p, \neg q}{\perp}$ (since "$\to$" amounts to implication-with-possible-exceptions).

In Sections 2 and 3, we will present mathematical definitions of implication-with-possible-exceptions and of an inference relation that behave as indicated.

If we say that implication as used in everyday reasoning rarely amounts to implication-without-possible-exceptions, it is important to understand that this is not a linguistic issue. Although it is also true that conditional sentences as used in natural language may allow exceptions, we consider implication to be an instrument used by humans to organize or represent knowledge *in their mind*, without the help of natural language. It just happens not to be an essential aspect of implication that it cannot have exceptions. On the contrary, implication signifies the existence of a rule. But rules may have exceptions and nevertheless be valid. Hence, in this paper, we will take the position that it is natural to interpret implication as implication-as-a-rule-with-possible-exceptions. We will defend this viewpoint, not by argumentation leading to a conclusion, but by investigating its mathematical consequences and possibilities.

We want to study non-monotonic logic as a subdiscipline of logic, using mathematical methods. There is reason to doubt, however, whether the mathematical methods typically used in logic will suffice. The activity, at present, in the field of non-monotonic reasoning can roughly be divided into two sectors. The first sector consists of proposals for concrete prescriptions, usually stemming from an attempt to simulate human reasoning processes. Most of them either use probabilistic considerations (e.g., [Adams, 1975], [Bacchus, 1990]; see also [Geffner, 1992]), try to model reasoning-by-lack-of-information (e.g., default logic and auto-epistemic logic, see [Brewka, 1991]), or involve some minimalization of possible models, minimizing, for example, the collection of unexplained exceptions to rules or unexpected changes in a changing domain (e.g., circumscription, see [Brewka, 1991]). Since most of these "concrete" approaches were designed with a number of concrete examples in mind, they typically lack generality. The general tendency in this sector is towards more and more complicated machinery (for example, [Abdallah, 1995], 715 pages).

The second sector consists of more axiomatic approaches, in which it is asked what properties some reasonable system should or may have (e.g., [Kraus, Lehmann, and Magidor, 1990], [Flach, 1995], but we can also consider [Alchourrón, Gärdenfors, and Makinson, June 1985] to be in this sector). These approaches mimic classical logic in that completeness results are used to establish connections between semantical constructions and axiom systems. Since we aim at a logical, non-adhoc approach, it will be clear that it is this sector that will be of interest to us. But the approaches in this sector have problems analogous to those in the first sector. Most proposed axiom systems seem either too weak or too strong. Which brings us to another main issue with which we began our inquiry: *why is it so difficult* to find a convincing formal treatment of these types of reasoning, while, at the same time, humans seem to handle them so easily?

Our answer will essentially be to indicate a lack of expressive power of the mathematical modelling methods used in most approaches, in particular, the way in which rules of inference are interpreted. It is in this context that one may ask, *what would happen if we interpret rules of inference, being implicational statements, as rules-with-possible-exceptions*? This paper is, by and large, an attempt to provide mathematical methods by which this (latter) question may be investigated.

In Section 2, we will present a topological interpretation of implication-with-possible-exceptions, inspired by geometrical intuitions. This semantic construction is not intended to bear direct resemblance to any intuitive interpretation of natural implication, it serves as basic material for the rest of the paper. The simplest way to use this definition in a caricature of practical reasoning, however, turns out to be equivalent to the well-known approaches of Shoham or Kraus et al. (see [Shoham, 1988] and [Kraus, Lehmann, and Magidor, 1990]), involving partially ordered sets of possible worlds.

In Section 3, the same topological notion will be used to interpret inference. This allows us to talk about rules of inference (that is, laws of logic) as rules-with-possible-exceptions. It is then established that, in at least one natural extension of the simple system of Section 2, each of the laws of classical logic is valid-as-a-rule, including the rule of monotony. The system, however, is non-monotonic from a mathematical point of view (since the rule of monotony has exceptions). In [Jurjus, 1997, Chapter 4], the topological notion of Section 2 is also used to interpret universal quantification.

§1. Preliminaries. Below we treat some well-known definitions from the theory of Boolean algebras and from topology. More information can be found in any textbook on these subjects, for example in [Abbott, 1969] and [Gaal, 1964], respectively.

DEFINITION 1.1. Let X be a set. A *partial ordering* (on X) is a binary relation $\leq$ satisfying:

i) for all $p \in X : p \leq p$,
ii) if $p, q, r \in X$, $p \leq q$ and $q \leq r$, then $p \leq r$,
iii) if $p, q \in X$, $p \leq q$ and $q \leq p$, then $p = q$.

DEFINITION 1.2. A *Boolean algebra* is a structure $(B, \leq, \wedge, \vee, \neg, \bot, \top)$, where B is a set, $\leq$ is a partial ordering on B, $\wedge$ and $\vee$ are binary operations on B, $\neg$ is a unary operation on B, and $\bot$ and $\top$ are elements of B, such that

i) for all $a, b, c \in B : a \leq b \wedge c$ (only) if $a \leq b$ and $a \leq c$,
ii) for all $a, b, c (\in B) : a \vee b \leq c$ (only) if $a \leq c$ and $b \leq c$,
iii) for all a, b, c : if $b \wedge a \leq c$ and $b \wedge (\neg a) \leq c$, then $b \leq c$,
iv) for all a, b, c : if $a \leq b$ and $a \leq \neg b$, then $a \leq c$,
v) for all $a : \bot \leq a$ and $a \leq T$.

EXAMPLE 1.3. If X is a set, then $(\mathscr{P}(X), \subseteq, \cap, \cup, .^c, \emptyset, X)$ is a Boolean algebra.

Instead of "the Boolean algebra $(B, \leq, \wedge, \vee, \neg, \bot, \top)$", we will typically write "the Boolean algebra $(B, \leq)$" or "the Boolean algebra B", if there is no danger of confusion.

DEFINITION 1.4. Let B_1 and B_2 be Boolean algebras (we will write $\leq_1$, $\leq_2$, $\wedge_1$, $\wedge_2$, etc.). A *Boolean translation* from B_1 to B_2 is a map $\phi : B_1 \to B_2$ satisfying:

i) for all $a, b \in B_1$ such that $a \leq_1 b : \phi(a) \leq_2 \phi(b)$,
ii) for all $a, b \in B_1 : \phi(a \wedge_1 b) = \phi(a) \wedge_2 \phi(b)$,
iii) for all $a, b \in B_1 : \phi(a \vee_1 b) = \phi(a) \vee_2 \phi(b)$,
iv) for all $a \in B_1 : \phi(\neg_1(a)) = \neg_2(\phi(a))$.

Typical associations and notations. If B is a Boolean algebra, X is a set and $\phi : B \to \mathscr{P}(X)$ is a Boolean translation, then the elements of B are sometimes thought of as sentences or *propositions*. The elements of X are typically called *possible worlds*. If $a \in B$ and $w \in X$, we say that

$$w \models_{(X,\phi)} a \qquad (\text{“}a \text{ is true in world } w\text{”})$$

whenever $w \in \phi(a)$. This is sometimes abbreviated to $w \models_X a$, $w \models_\phi a$ or even to $w \models a$, if there is no danger of confusion. Hence, for all $a \in B$,

$$\phi(a) = \{w \in X \quad | \quad w \models a\},$$

"the set of all possible worlds in which a is true". For this reason, $\phi(a)$ is sometimes called the *extension* of the proposition a. Furthermore, if a and b are elements of B, we say

$$a \models_{(X,\phi)} b$$

whenever $\phi(a) \subseteq \phi(b)$. This is sometimes abbreviated to $a \models_X b$, $a \models_\phi b$ or $a \models b$.

Note that with these notations, the properties i), ii), iii) and iv) in Definition 1.4 amount to:

i) for all $a, b \in B$: if $a \leq b$, then $a \models_X b$,
ii) for all $a, b \in B$, $w \in X$: $w \models a \wedge b$ (only) if $w \models a$ and $w \models b$,
iii) for all $a, b \in B$, $w \in X$: $w \models a \vee b$ (only) if $w \models a$ or $w \models b$,
iv) for all $a \in B$, $w \in X$: $w \models \neg a$ (only) if $w \not\models a$.

DEFINITION 1.5. Let X be a set. A *topology on X* is a collection τ of subsets of X satisfying:

i) $X \in \tau$, $\emptyset \in \tau$,
ii) If a and $b \in \tau$, then $a \cap b \in \tau$,
iii) If $a_i \in \tau$ for all $i \in I$, then $\cup_{i \in I} a_i \in \tau$.

"Let (X, τ) be a topological space" means: let τ be a topology on X. Sometimes we just say "let X be a topological space". $a \subseteq X$ is called an *open* subset of X whenever $a \in \tau$.

In this paper, variable names like O, O', O_1, etc. are used exclusively for open sets. In addition, we will skip the word "open" as much as possible.

DEFINITION 1.6. A space X is called a *Hausdorff space* whenever for every $p, q \in X$ such that $q \neq p$ there exist $O_p, O_q \subseteq X$ such that $p \in O_p$, $q \in O_q$ and $O_p \cap O_q = \emptyset$.

EXAMPLE 1.7. If X is a set, then $\mathscr{P}(X)$, the collection of all subsets of X, is a topology on X, called the *discrete* topology. The collection $\{\emptyset, X\}$ is also a topology on X, called the *trivial* topology. Any set equipped with the discrete topology is a Hausdorff space. A set equipped with the trivial topology is not a Hausdorff space, unless the underlying set is empty or contains only one element.

EXAMPLE 1.8. The Euclidean plane, $\mathbb{R}^2$, is standardly assumed to be equipped with the *Euclidean topology*, defined as follows: $O \subseteq \mathbb{R}^2$ is open whenever for all $p \in O$ there is a circular-disc-without-boundary contained in O and containing p. Hence, circular-discs-without-boundary are open, all unions of such sets are open

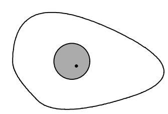

and every open set is a union of such sets. That is, the open sets of the Euclidean plane are precisely the unions of circular-discs-without-boundary.

EXAMPLE 1.9. Likewise, $\mathbb{R}$ is canonically equipped with the topology consisting of all unions of *open intervals*.

Intuitively, a topology on a set X is a device that describes how the points of X are geometrically "glued together". Typically there exists more than one topology on every set, corresponding to different ways of glueing the points. Not all topological spaces are equally convincing as carriers of some geometrical intuition. For example, topologies on finite spaces are typically difficult to interpret geometrically. Such spaces are sometimes called "pathological spaces" (a notion that does not have a precise definition).

DEFINITION 1.10. If (X, τ) is a topological space and $a \subseteq X$, then a inherits a topology from X, called the *induced topology (on a)*:

$$\tau_a := \{O \cap a \quad | \quad O \in \tau\}$$

This amounts to $b \subseteq a$ being open in a (only) if there is some O, open in X, such that $b = O \cap a$. (Again, it is easy to check that this defines a topology on a.)

For example, the interval $(1/2, 1]$ is not open in $\mathbb{R}$, but it is *open in* $[0, 1]$.

It is custom to generalize this definition, by calling any $b \subseteq X$ open in a whenever $b \cap a$ is open in a.

§2. Full subsets.

DEFINITION 2.1. Let X be a topological space, and let a be a subset of X. We say that a is *full* (in X) whenever every nonempty open O contains a nonempty open O' (open in X) with $O' \subseteq a$.

EXAMPLE 2.2. If E denotes the Euclidean plane, and ℓ a line in it, then $E \backslash \ell \; (= \ell^c)$ is full in E ("almost all elements of E are elements of $E \backslash \ell$.")

The following picture might suffice to see this:

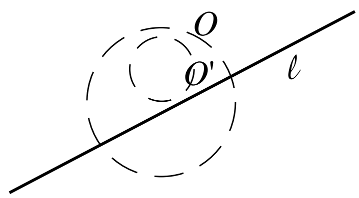

Definition 2.3. Given a topological space X and $a, b \subseteq X$, we say that b *is full in* a (written as "$a \rightarrow b$") whenever $a \cap b$ is a full subset of a-with-induced-topology.

"$a \rightarrow b$" is considered to be a possible interpretation of the phrase "almost all elements of a are also in b".

Proposition 2.4. *If X is a topological space and $a, b \subseteq X$, then $a \rightarrow b$ (only) if every O (open in X) such that $O \cap a \neq \emptyset$ contains an O' such that $O' \cap a \neq \emptyset$ and $O' \cap a \subseteq b$. (Provable by elementary check.)*

Example 2.5. If E denotes the Euclidean plane, and ℓ some line in it, then ℓ^c is full in $E (E \rightarrow \ell^c)$ but not in ℓ $(\ell \nrightarrow \ell^c)$. After all, almost all elements of E are in $E \backslash \ell$. But of the elements of ℓ, not even a single one is in $E \backslash \ell$. If $a \subseteq b$, then $a \rightarrow b$; but not conversely: $E \rightarrow \ell^c$, while not $E \subseteq \ell^c$.

The notion of "b is full in a", written as "$a \rightarrow b$" could be seen as a topological notion capturing the intuition of inclusion-up-to-possible-exceptions. It has very nice properties. The example above shows that it is non-monotonic: $E \rightarrow \ell^c$ in E, but $E \cap \ell \nrightarrow \ell^c$ in E; hence, $a \rightarrow b$ does, in general, not imply $a \cap c \rightarrow b$.

Definition 2.6. Let $(B, \leq)$ be a Boolean algebra. A binary relation, $\mid\!\sim$, on B is called a *preferential consequence relation* (on B) (see [Kraus, Lehmann, and Magidor, 1990]) whenever

(P1) for all $a \in B$, $a \mid\!\sim a$,
(P2) for all $a, b, c \in B$, if $a \mid\!\sim b$ and $b \leq c$, then $a \mid\!\sim c$,
(P3) for all $a, b, c \in B$, if $a \mid\!\sim b$ and $a \wedge b \mid\!\sim c$, then $a \mid\!\sim c$,
(P4) for all $a, b, c \in B$, if $a \mid\!\sim b$ and $a \mid\!\sim c$, then $a \wedge b \mid\!\sim c$,
(P5) for all $a, b, c \in B$, if $a \mid\!\sim c$ and $b \mid\!\sim c$, then $a \vee b \mid\!\sim c$.

Theorem 2.7. *Let X be a topological space. The relation $\rightarrow$ on $\mathscr{P}(X)$ is a* preferential consequence *relation on* $(\mathscr{P}(X), \subseteq)$.

Proof. (P1) is trivial. (P2) is a direct consequence of the fact that for all a, if a is full in X and $a \subseteq a' \subseteq X$, then a' is full in X.

(P3) Suppose that $a \rightarrow b$ and $a \cap b \rightarrow c$. For every O such that $O \cap a \neq \emptyset$, let $O' \subseteq O$ be such that $O' \cap a \neq \emptyset$ and $O' \cap a \subseteq b$. Then $O' \cap (a \cap b) \neq \emptyset$. Let $O'' \subseteq O'$ be such that $O'' \cap (a \cap b) \neq \emptyset$ and $O'' \cap (a \cap b) \subseteq c$. Then $O'' \cap a \neq \emptyset$ and $O'' \cap a \subseteq O'' \cap (a \cap b) \subseteq c$. Hence, $a \rightarrow c$.

(P4) Suppose that $a \rightarrow b$ and $a \rightarrow c$. For every O such that $O \cap a \cap b \neq \emptyset$, let $O' \subseteq O$ be such that $O' \cap a \neq \emptyset$ and $O' \cap a \subseteq b$. Let $O'' \subseteq O'$ be such that $O'' \cap a \neq \emptyset$ and $O'' \cap a \subseteq c$. Then $O'' \cap (a \cap b) = O'' \cap a \neq \emptyset$ and $O'' \cap (a \cap b) \subseteq c$. Hence, $a \cap b \rightarrow c$.

(P5) Suppose that $a \rightarrow c$ and $b \rightarrow c$. Suppose that $O \cap (a \cup b) \neq \emptyset$. Say $O \cap a \neq \emptyset$. Let $O' \subseteq O$ be such that $O' \cap a \neq \emptyset$ and $O' \cap a \subseteq c$. If $O' \cap b = \emptyset$, then

$O' \cap (a \cup b) \neq \emptyset$ and $O' \cap (a \cup b) \subseteq c$. If $O' \cap b \neq \emptyset$, then let $O'' \subseteq O'$ be such that $O'' \cap b \neq \emptyset$ and $O'' \cap b \subseteq c$. Then $O'' \cap (a \cup b) \neq \emptyset$, and $O'' \cap a \subseteq O' \cap a \subseteq c$, hence $O'' \cap (a \cup b) \subseteq c$. Hence, $a \cup b \to c$. ⊣

Definition 2.8. Let B be a finite Boolean algebra. A *topological model* (of B) is a pair (X, ϕ), where

i) X is a topological space,
ii) $\phi : B \to \mathscr{P}(X)$ is a Boolean translation.

(The elements of X are nicknamed "possible worlds", and for all $a \in B$, $\phi(a)$ is the set of all possible worlds in which a is said to be true.)

For $a, b \in B$, we will say "$a \to b$ is *true in* (X, ϕ)" whenever $\phi(a) \to \phi(b)$ in the topological space X, that is: almost all possible worlds in X that satisfy a do also satisfy b.

Example 2.9. Let B be the free Boolean algebra generated by two (distinct) basic formulas, a and b. Then there exist topological models of B in which $a \to b$ is true, but $(a \wedge \neg b) \to b$ is not.

Proof. If E is the Euclidean plane and ℓ a line in E, then there is one and only one Boolean translation $\phi : B \to \mathscr{P}(E)$ such that $\phi(a) = E$ and $\phi(b) = E \backslash \ell$. In the topological model (E, ϕ), $a \to b$ is true (since $E \to \ell^c$), while $(a \wedge \neg b) \to b$ is not: $E \cap \ell \nrightarrow \ell^c$, since ℓ is not empty. ⊣

We could think of the formulas a and b as signifying the sentences "x is a bird" and "x can fly", respectively. Then the elements of X should rather be thought of as possible instances that x could refer to. It is not claimed, of course, that the set of all birds bears any resemblance with the Euclidean plane, but that the relation between birds and birds that cannot fly could be thought to resemble the relation between E and ℓ. Note that, for this reading of a and b, the topological model above shows that the statement "if x is a bird, then it can fly" does not imply "if x is a bird and x cannot fly, then x can fly". According to classical logic, however (that is, interpreting implication as implication-without-exceptions), this latter statement is a necessary consequence of the former. In this way, certain combinations of sentences considered inconsistent when using the more usual interpretation of implication are nevertheless representable using topological models. We will call them "topologically representable" or even "topologically consistent".

Now it is time for a complete characterization of all topologically representable situations, or equivalently, for a complete axiomatization of topological models. Below we will prove the completeness theorem, saying that (P1) – (P5) is a complete set of axioms for the relation "$\to$" in topological models.

Definition 2.10. Let $(X, \leq)$ be a partial ordering and $a, b \in X$. An element $x \in a$ is minimal in a whenever there is no $y \in a$ such that $y \leq x$ except x itself. We will use $a \mid\!\sim_{\leq} b$ to denote : every element of a that is minimal in a is an element of b.

If we think of the elements of X as possible worlds, the partial ordering is supposed to describe a "normality" ordering on the possible worlds, or a preference of some possible worlds over others. For $x_1, x_2 \in X$, $x_1 \leq x_2$ is to be read as "x_1 is a more normal world than x_2" or "x_1 is preferred over x_2". Then $a \mid\!\sim_{\leq} b$ amounts to:

"of all the worlds in a, at least the most normal (preferred) ones are in b". It is easy to see, that for every *finite* partial ordering $(X, \leq)$, the relation $\mid\!\sim_{\leq}$ is a preferential consequence relation on the Boolean algebra $(\mathscr{P}(X), \subseteq)$.

THEOREM 2.11 (Kraus, Lehmann, Magidor 1990). *Let B be a finite Boolean algebra, and $\mid\!\sim$ a preferential consequence relation on B. Then there is a finite partially ordered set $(X, \leq)$ and a Boolean translation $\phi : B \to P(X)$ such that for all $a, b \in B$, $a \mid\!\sim b$ (only) if $\phi(a) \mid\!\sim_{\leq} \phi(b)$.*

For the proof, the reader is referred to [Kraus, Lehmann, and Magidor, 1990].

We are now ready to prove a similar theorem concerning topological models. The completeness result will be an easy corollary of the following theorem.

THEOREM 2.12. *For every partial ordering, $\leq$, on a finite set X there is at least one topology on X such that, for all $a, b \in X$, $a \mid\!\sim_{\leq} b$ (only) if $a \to b$.*

PROOF. Let $(X, \leq)$ be a partial ordering. Define $\tau := \{v \subseteq X \mid \text{if } p \in v \text{ and } q \leq p \text{ then } q \in v\}$. This is a topology on X. For every $p \in X$, let O_p denote $\{q \in X \mid q \leq p\}$. Suppose that $a \to b$. If p is minimal in a, then $O_p \cap a = \{p\}$. Since $a \to b$ is true, there is an $O' \subseteq O_p$ such that $O' \cap a \neq \emptyset$ and $O' \cap a \subseteq b$. Hence, $O' \cap a = \{p\}$ and $p \in b$. Hence $a \mid\!\sim_{\leq} b$. On the other hand, suppose that $a \mid\!\sim_{\leq} b$. Let O be such that $O \cap a \neq \emptyset$. Say $p \in a$ and $p \in O$. Since X is finite, there is a q that is minimal in a such that $q \leq p$. For such q : $O_q \subseteq O_p \subseteq O$ and $q \in b$ (since $a \sim_{\leq} b$). Hence, $O_q \subseteq O$, $O_q \cap a = \{q\} \neq \emptyset$, and $O_q \cap a \subseteq b$. Hence $a \to b$. Hence, for all $a, b \in X$, $a \to b$ (only) if $a \mid\!\sim_{\leq} b$. ⊣

COROLLARY 2.13 (Completeness Theorem). *Let B be a finite Boolean algebra, and $\mid\!\sim$ a preferential consequence relation on B. Then there is a (finite) topological space X and a Boolean translation $\phi : B \to \mathscr{P}(X)$ such that for all $a, b \in B$, $a \mid\!\sim b$ (only) if $\phi(a) \to \phi(b)$. (This is an immediate consequence of the two theorems above.)*

This theorem characterizes the relation "$\to$" in topological models of finite Boolean algebras, thus settling the task of characterizing the "topologically representable situations" as meant above.

DEFINITION 2.14. i) Let $(X, \leq)$ be a partial ordering. A subset $a \subseteq X$ is called *smooth* whenever for all $x \in a$, there is a minimal element x' of a such that $x' \leq x$.

ii) Let B be a Boolean algebra. A *preferential model* (of B) is a triple $(X, \leq, \phi)$ such that:

$(X, \leq)$ is a partial ordering,
ϕ is a Boolean translation $B \to \mathscr{P}(X)$,
for every $a \in B$, $\phi(a)$ is smooth in $(X, \leq)$.

In comparison to preferential semantics, our topological presentation enables us to state and prove insights that would otherwise be hard to achieve (for example, the three Lemmas below, or Theorems 8.1, 8.3 and 8.7 in [Jurjus, 1997]). Moreover, as shown in the same [Jurjus, 1997], the class of topological models is strictly larger than the class of preferential models. Note further that although Theorem 2.13 is restricted to finite Boolean algebras, it is not difficult to extend the result to infinite Boolean algebras, using Definition 2.14, and an analogous completeness theorem in [Kraus, Lehmann, and Magidor, 1990]. See also [Jurjus, 1997].

The three lemmas below are, in the first place, technically useful results. They do not seem to have a direct equivalent in preferential semantics.

LEMMA 2.15 ("the open-lemma"). *For all* $a, b, c \subseteq X$, $a \to b$ *implies* $a \cap c \to b$ *if* c *is open* (*either in* X *or in* a).

PROOF. If $a \to b$ and c is open (either in X or in a), then for every O such that $(O \cap c) \cap a \neq \emptyset$, there is an $O' \subseteq (O \cap c)$ such that $O' \cap a \neq \emptyset$ and $O' \cap a \subseteq b$. Then $O' \cap (a \cap c) = O' \cap a \neq \emptyset$ and $O' \cap (a \cap c) \subseteq b$. Hence, $a \cap c \to b$. ⊣

LEMMA 2.16 ("the dense-lemma"). *For all* $a, b, c \subseteq X$, $a \to b$ *implies* $a \cap c \to b$ *if* c *is dense in* a.

PROOF. If $a \to b$ and c is dense in a, then for every O such that $O \cap c \cap a \neq \emptyset$ (hence, $O \cap a \neq \emptyset$), there is an $O' \subseteq O$ such that $O' \cap a \neq \emptyset$ and $O' \cap a \subseteq b$. Since c is dense in a, this implies $O' \cap a \cap c \neq \emptyset$. But $O' \cap a \cap c \subseteq b$. Hence, $a \cap c \to b$. ⊣

LEMMA 2.17 ("the closed-lemma"). *For all* $a, b, c \subseteq X$, $a \to b$ *implies* $a \cap c \to b$ *if* b *is closed* (*either in* X *or in* a).

PROOF. If b is closed (either in X or in a), and $a \to b$, then b is both closed in a and dense in a. Hence, $a \subseteq b$. Hence, for all c, $a \cap c \to b$. ⊣

Full subsets in general, and the closed lemma in particular, are connected to a subject that could be called the "calculus of degenerate cases" in classical geometry.

For example, suppose we have two statements,

$$A(P_1, \dots P_n, l_1, \dots l_m) \qquad \text{and} \qquad B(P_1, \dots P_n, l_1, \dots l_m)$$

about a number of points (P_i) and lines (l_i) in the Euclidean plane, and suppose that we want to prove a theorem stating that A implies B (for all points $P_1, \dots P_n$, and all lines $l_1, \dots l_m$). Let Π be the product space $E^n \times G^m$, where E denotes the Euclidean plane, and G the space of all lines in E. Then we are dealing with two subsets $a, b \subseteq \Pi$. One of the consequences of the closed lemma is that if there is a set $f \subseteq \Pi$ such that

$$a \to f \qquad \text{and} \qquad f \subseteq b \qquad (\text{which implies } a \to b)$$

then we can conclude that $a \subseteq b$, provided b is closed in Π. Note that many theorems in classical geometry indeed have a "closed statement" as their conclusion:

"$P_1 = P_2$", "P_1 is on l_1", "l_1, l_2 and l_3 have one point in common", etc.

Typical candidate-statements for f are "exclusions of degenerate cases":

"$P_1 \neq P_2$", "P_1, P_2 and P_3 are not on one line", etc.

Many 'classical' proofs of theorems silently use one or more 'innocent' extra assumptions like this. The closed lemma may serve to explain why addition of this kind of assumptions is often innocent, indeed.

A more detailed treatment of this "calculus of degenerate cases" is beyond the scope of this article.

§3. Non-monotonic inference. L will denote a language generated via $\wedge, \vee, \to$ and $\neg$, from a finite number of basic formulas, "$\to$" being an additional binary connective. We will use $\mathtt{a}, \mathtt{b}, \mathtt{c}, \dots$ as variables ranging over L. $a, b, c, \dots$ denote subsets of some topological space. Hence, $\mathtt{a} \to \mathtt{b}$ denotes a formula in L, while

$a \to b$ is not a formula, but a sentence about two sets, a and b. This multiple use of the symbol "$\to$" should not cause much confusion.

DEFINITION 3.1. An extended model (of L) is a tuple (X, I, ϕ), where X is a topological space, $I(\ ,\ )$ is a binary operation on $\mathscr{P}(X)$ such that for all $a, b \subseteq X$, $I(a, b) \subseteq X$, and ϕ is a map $L \to \mathscr{P}(X)$ such that for all $a, b \in L$,

$$\begin{aligned} \phi(\mathtt{a} \wedge \mathtt{b}) &= \phi(\mathtt{a}) \cap \phi(\mathtt{b}), \\ \phi(\mathtt{a} \vee \mathtt{b}) &= \phi(\mathtt{a}) \cup \phi(\mathtt{b}), \\ \phi(\neg \mathtt{a}) &= \phi(\mathtt{a})^c, \\ \phi(\mathtt{a} \to \mathtt{b}) &= I(\phi(\mathtt{a}), \phi(\mathtt{b})). \end{aligned}$$

When there is no danger of confusion, we will freely speak about "model X", "model (X, ϕ)", "model X based on I" etc., instead of "model (X, I, ϕ)".

DEFINITION 3.2. A formula $\mathtt{a} \in L$ is called *true in* (X, I, ϕ), whenever $\phi(\mathtt{a}) = X$.

It will be clear that (X, I, ϕ) is not for all $I(\ ,\ )$ acceptable as something that provides sensible interpretations for formulas containing implication. For example, for every $a, b \in L$, we would like $a \to b$ to be true in (X, I, ϕ) (only) if $\phi(a) \to \phi(b)$ in the topological space X.

DEFINITION 3.3. $I(\ ,\ )$ is called an *implication operator* on X whenever for all $a, b \subseteq X$, $I(a, b) = X$ (only) if $a \to b$.

PROPOSITION 3.4. *If (X, I, ϕ) is an extended model, and $I(\ ,\ )$ is an implication operator on X, then, for every $\mathtt{a}, \mathtt{b} \in L$, $\mathtt{a} \to \mathtt{b}$ is true in (X, I, ϕ) (only) if $\phi(\mathtt{a}) \to \phi(\mathtt{b})$ in the topological space X. (Provable by elementary check.)*

EXAMPLE 3.5. For every topological space, the "trivial" operator, defined by $T(a, b) = X$ if $a \to b$, and $T(a, b) = \emptyset$ if $a \nrightarrow b$, is an implication operator. The operator $Q(\ ,\ )$ ("material implication"), defined by $Q(a, b) = a^c \cup b$ is not an implication operator: $Q(a, b) = X$ means $a \subseteq b$ which is not equivalent to $a \to b$.

In connection with the geometrical motivation of the notion of a full subset, there is a very natural implication operator on every topological space:

DEFINITION 3.6. Let X be a topological space, and $a, b \subseteq X$. Then $O(a, b)$ is the union of all (open) $O \subseteq X$ such that $O \cap a \to b$.

Hence, for all $w \in X$, $w \in O(a, b)$ (only) if there is some O such that $w \in O$ and $O \cap a \to b$. This amounts to the following: a sentence $\mathtt{a} \to \mathtt{b}$ is defined to be true in world w whenever "in some neighbourhood of w, practically all worlds that satisfy $\mathtt{a}$ do also satisfy $\mathtt{b}$". Note that $O(a, b)$ is open for all $a, b \subseteq X$.

EXAMPLE 3.7. If E denotes the Euclidean plane, and ℓ a line in E, then $O(E, \ell^c) = E$, $O(\ell, \ell^c) = \ell^c$. Note that $O(E, \ell^c) \cap E \nsubseteq \ell^c$, but $O(E, \ell^c) \cap E \to \ell^c$ in E.

PROPOSITION 3.8. *Let X be a topological space. Then:*

i) *for all $a, b \subseteq X$, $O(a, b) \cap a \to b$,*
ii) *$O(\ ,\ )$ is an implication operator.*

PROOF. Suppose that $O \cap (O(a,b) \cap a) \neq \emptyset$. Say $p \in O \cap (O(a,b) \cap a)$. Then $p \in O(a,b)$, hence, there is an O' such that $O' \cap a \rightarrow b$ and $p \in O'$, hence $O \cap (O' \cap a) \neq \emptyset$. Since $O' \cap a \rightarrow b$, there is an $O'' \subseteq O$ such that $O'' \cap (O' \cap a) \neq \emptyset$ and $O'' \cap (O' \cap a) \subseteq b$. But $O''' := O'' \cap O' \subseteq O(a,b)$, hence $O''' \cap (O(a,b) \cap a) \neq \emptyset$ and $O''' \cap (O(a,b) \cap a) \subseteq b$. Hence, $O(a,b) \cap a \rightarrow b$. To see that $O(\ ,\ )$ is an implication operator, suppose that $a \rightarrow b$. Then $X \cap a \rightarrow b$, and X is an open set. Hence $O(a,b) = X$. On the other hand, suppose that $O(a,b) = X$. Then i) implies $a \rightarrow b$. ⊣

COROLLARY 3.9. *$O(a,b)$ is the largest (open) O such that $O \cap a \rightarrow b$. Note that, for all $a, b \subseteq X$, $a^c \cup b$ is the largest $V \subseteq X$ such that $V \cap a \subseteq b$.*

Using the implication operator $O(\ ,\ )$, we can now interpret rules of inference as follows.

DEFINITION 3.10. If $\mathtt{a}_1, \dots, \mathtt{a}_m, \mathtt{b}_1, \dots, b_n$ are formulas of L, then the *rule*

$$\frac{\mathtt{a}_1, \dots, \mathtt{a}_m}{\mathtt{b}_1, \dots, \mathtt{b}_n}$$

is called topologically valid or O-valid whenever the formula $\mathtt{a}_1 \wedge \dots \wedge \mathtt{a}_m \rightarrow \mathtt{b}_1 \wedge \dots \wedge \mathtt{b}_n$ is true in every extended model based on $O(\ ,\ )$ (that is, in every extended model $(X, O(\ ,\ ), \phi)$).

Note that this definition involves nested implicational statements whenever any of the formulas $\mathtt{a}_1, \dots, \mathtt{a}_m, \mathtt{b}_1, \dots, \mathtt{b}_n$ contains the symbol "$\rightarrow$".

EXAMPLE 3.11. The rule $\dfrac{\mathtt{a} \rightarrow \mathtt{b}, \mathtt{a}}{\mathtt{b}}$ is topologically valid (this is an easy corollary of the Proposition 3.8 i). However, the rule $\dfrac{\mathtt{a} \rightarrow \mathtt{b}, \mathtt{a}, \neg\mathtt{b}}{\mathtt{b}}$ is not topologically valid: if E is the Euclidean plane, and ℓ is a line in E, then $O(E, \ell^c) = E$, and $O(E, \ell^c) \cap E \cap \ell \rightarrow \ell^c$ is not true.

THEOREM 3.12. $\dfrac{\mathtt{a} \rightarrow \mathtt{b}}{\mathtt{a} \wedge \mathtt{c} \rightarrow \mathtt{b}}$ *is topologically valid.*

PROOF. Let X be a topological space, and $a, b, c \subseteq X$. Suppose that $O \cap O(a,b) \neq \emptyset$. We may assume that $O \subseteq O(a,b)$. Then $O \cap a \rightarrow b$. If c is dense in $O \cap a$, then $O \cap a \cap c \rightarrow b$, hence $O \subseteq O(a \cap c, b)$. If c is not dense in $O \cap a$, then there is an $O' \subseteq O$ such that $O' \cap a \neq \emptyset$, and $O' \cap a \cap c = \emptyset$. Then $O' \cap a \cap c \rightarrow b$, hence $O' \subseteq O(a \cap c, b)$. Moreover, $O' \cap O(a,b) \neq \emptyset$, since $O' \neq \emptyset$. In both cases, there exists an $O' \subseteq O$ such that $O' \cap O(a,b) \neq \emptyset$ and $O' \subseteq O(a \cap c, b)$. ⊣

Theorem 3.12 allows us to think of the rule of monotony as *valid*, be it only up to possible exceptions. Our system is non-monotonic, nevertheless, in that "$\mathtt{a} \rightarrow \mathtt{b}$ is true in the extended model (X, O, ϕ)" does not imply "$\mathtt{a} \wedge \mathtt{c} \rightarrow \mathtt{b}$ is true in the extended model (X, O, ϕ)". Nor does "$\mathtt{a} \rightarrow \mathtt{b}$ is true in every extended model (X, O, ϕ)" imply "$\mathtt{a} \wedge \mathtt{c} \rightarrow \mathtt{b}$ is true in every extended model (X, O, ϕ)". In other words, the topological validity of $\dfrac{\mathtt{a}}{\mathtt{b}}$ does not imply the topologically validity of $\dfrac{\mathtt{a} \wedge \mathtt{c}}{\mathtt{b}}$ (see the example above). In general, *the topological validity of* $\dfrac{\mathtt{a}}{\mathtt{b}}$ and $\dfrac{\mathtt{b}}{\mathtt{c}}$ does not imply the topologically validity of $\dfrac{\mathtt{a}}{\mathtt{c}}$.

THEOREM 3.13. *Each of the following rules is topologically valid*:

$$\frac{\mathtt{c} \to \mathtt{a}, \mathtt{c} \to \mathtt{b}}{\mathtt{c} \to \mathtt{a} \wedge \mathtt{b}} \qquad \frac{\mathtt{c} \to \mathtt{a} \wedge \mathtt{b}}{\mathtt{c} \to \mathtt{a}} \frac{\mathtt{c} \to \mathtt{a} \wedge \mathtt{b}}{\mathtt{c} \to \mathtt{b}}$$

$$\frac{\mathtt{c} \to \mathtt{a}}{\mathtt{c} \to \mathtt{a} \vee \mathtt{b}} \qquad \frac{\mathtt{c} \to \mathtt{b}}{\mathtt{c} \to \mathtt{a} \vee \mathtt{b}} \qquad \frac{\mathtt{c} \to \mathtt{a} \vee \mathtt{b}, \mathtt{c} \wedge \mathtt{a} \to \mathtt{d}, \mathtt{c} \wedge \mathtt{b} \to \mathtt{d}}{\mathtt{c} \to \mathtt{d}}$$

$$\frac{\mathtt{c} \wedge \mathtt{a} \to \mathtt{b}}{\mathtt{c} \to \neg\mathtt{a} \vee \mathtt{b}} \qquad \frac{\mathtt{c} \to \mathtt{a}, \mathtt{c} \to \neg\mathtt{a} \vee \mathtt{b}}{\mathtt{c} \to \mathtt{b}}$$

$$\frac{\mathtt{c} \wedge \mathtt{a} \to \mathtt{b}, \mathtt{c} \wedge \mathtt{a} \to \neg\mathtt{b}}{\mathtt{c} \to \neg\mathtt{a}} \qquad \frac{\mathtt{c} \to \mathtt{a}, \mathtt{c} \to \neg\mathtt{a}}{\mathtt{c} \to \mathtt{b}}$$

As well as:

$$\frac{}{\mathtt{a} \to \mathtt{a}} \qquad \frac{\mathtt{c} \wedge \mathtt{a} \wedge \mathtt{b} \to \mathtt{d}}{\mathtt{c} \wedge \mathtt{b} \wedge \mathtt{a} \to \mathtt{d}} \qquad \frac{\mathtt{c} \wedge \mathtt{a} \wedge \mathtt{a} \to \mathtt{b}}{\mathtt{c} \wedge \mathtt{a} \to \mathtt{b}}$$

$$\frac{\mathtt{a} \to \mathtt{b}, \mathtt{a}}{\mathtt{b}} \ (\textit{modus ponens}) \qquad \frac{\mathtt{c} \to \mathtt{a}}{\mathtt{c} \wedge \mathtt{b} \to \mathtt{a}} \ (\textit{monotony})$$

PROOF. Let X be a topological space, and $a, b, c, d \subseteq X$. For most of the rules above, it is possible to prove an even stronger statement, as follows.

If $O_1 = O(c, a)$ and $O_2 = O(c, b)$ then $O_1 \cap O_2 \cap c \to a$ and $O_1 \cap O_2 \cap c \to b$ (by the open-lemma). Hence $O_1 \cap O_2 \cap c \to a \cap b$, hence $O_1 \cap O_2 \subseteq O(c, a \cap b)$. Hence, $O(c, a) \cap O(c, b) \subseteq O(c, a \cap b)$, which immediately implies $O(c, a) \cap O(c, b) \to O(c, a \cap b)$. Likewise: $O(c, a \cap b) \subseteq O(c, a), O(c, a \cap b) \subseteq O(c, b), O(c, a) \subseteq O(c, a \cup b), O(c, b) \subseteq O(c, a \cup b), O(c, a \cup b) \cap O(c \cap a, d) \cap O(c \cap b, d) \subseteq O(c, d)$ (since $(O \cap c) \cap a \to d, (O \cap c) \cap b \to d$ implies $(O \cap c) \cap (a \cap b) \to d$ and $(O \cap c) \cap (a \cup b) \to d$, $(O \cap c) \to a \cup b$ implies $(O \cap c) \to d)$, $O(c \cap a, b) \subseteq O(c, a^c \cup b)$ (since $v \cap a \to b$ implies $v \to a^c \cup b$), $O(c, a^c \cup b) \cap O(c, a) \subseteq O(c, b)$ (since $(a^c \cup b) \cap a \subseteq b$), $O(c \cap a, b) \cap O(c \cap a, b^c) \subseteq O(c, a^c)$, $O(c, a) \cap O(c, a^c) \subseteq O(c, b)$. The topological validity of modus ponens and the rule of monotony was established in Example 3.11 and Theorem 3.12, respectively. The rest is trivial. ⊣

Let us think of an imaginary person who uses the rules of propositional logic as rules-with-possible-exceptions. He knows about the distinction between object-language and meta-language. On the object level, he continuously uses $\neg \ldots \vee \ldots$ to interpret implicational statements. On the metalevel, he draws conclusions using the rules of inference of Theorem 3.13. The symbol "→" is the symbol that we use to denote the person's deducibility relation (it is not part of the person's object-language). For example, what we write as

$$\frac{\mathtt{a} \to \mathtt{b}, \mathtt{a} \to \mathtt{c}}{a \to \mathtt{b} \wedge \mathtt{c}}$$

is known by him as "if b is deducible from a, and c is deducible from a, then b-and-c is deducible from a" or, simpler, as the process of writing

$$\frac{\mathtt{b} \qquad \mathtt{c}}{\mathtt{b} \wedge \mathtt{c}}$$

somewhere within an argument or proof tree. (Note that modus ponens, in the form by which it was included in Theorem 3.13, also functions as such a rule in practical reasoning: if $\mathtt{a}$ and $\mathtt{b}$ are formulas such that $\mathtt{b}$ is deducible from $\mathtt{a}$, and, in some situation, with some concrete interpretation of the basic formulas occurring in $\mathtt{a}$ or $\mathtt{b}$, $\mathtt{a}$ is true, then modus ponens allows us to conclude that $\mathtt{b}$, with the same interpretation of the basic formulas, is also true in that situation.)

However, our imaginary person handles some of these rules (namely modus ponens and the rule of monotony) as if they were rules-with-possible-exceptions. The notion of topological validity, now, can be seen as a mathematical image of the reasoning behaviour of such a person. Thus, our imaginary person appreciates each of the laws of classical propositional logic, distinguishes between implication $(\neg \ldots \vee \ldots)$ and deducibility $(\rightarrow)$, uses an unambiguously defined and fixed language (having $\wedge$, $\vee$, and $\neg$), distinguishes between an objectlanguage and a metalanguage etc, etc. Note that the terms objectlanguage and metalanguage are meant, here, to refer to the person's objectlanguage and metalanguage. They are not to be confused with our objectlanguage and metalanguage, when reading, for example, the definition of topological validity or a theorem like 3.12. For example, "$\rightarrow$" is not part of the person's object language.

The rules of inference of Theorem 3.13, when interpreted as valid-without-exceptions, are known to be a complete characterization of classical propositional logic. Thus, the person could be said to reason non-monotonically "on the meta-level" while accepting each of the laws of classical propositional (monotonic) logic as valid-with-possible-exceptions.

The laws that the person appreciates as valid (be it only up-to-possible-exceptions) include the principle of monotony: the person appreciates "if $\mathtt{c}$ is deducible from the assumptions $\mathtt{a}_1, \ldots, \mathtt{a}_n$, then $\mathtt{c}$ is deducible from $\mathtt{a}_1, \ldots, \mathtt{a}_n, \mathtt{a}_{n+1}$" as valid. Nevertheless, "if $\mathtt{c}$ is deducible from the assumptions $\mathtt{a}_1, \ldots, \mathtt{a}_n$, and $\mathtt{c}$ is not deducible from $\mathtt{a}_1, \ldots, \mathtt{a}_n, \mathtt{a}_{n+1}$" is regarded as a consistent situation:

$$\frac{\mathtt{a}_1 \wedge \ldots \wedge \mathtt{a}_n \rightarrow \mathtt{c}}{\mathtt{a}_1 \wedge \ldots \wedge \mathtt{a}_n \wedge \mathtt{a}_{n+1} \rightarrow \mathtt{c}} \quad \text{is topologically valid, but}$$

$$\frac{\mathtt{a}_1 \wedge \ldots \wedge \mathtt{a}_n \rightarrow \mathtt{c}, \neg(\mathtt{a}_1 \wedge \ldots \wedge \mathtt{a}_n \wedge \mathtt{a}_{n+1} \rightarrow \mathtt{c})}{\mathtt{a}_1 \wedge \ldots \wedge \mathtt{a}_n \wedge \mathtt{a}_{n+1} \rightarrow \mathtt{c}} \quad \text{is not.}$$

Likewise, it follows by Example 3.11 that modus ponens, as handled by our person, is accepted as valid, but may have exceptions. It is important to note that our person handles the other rules of Theorem 3.13 as rules without exceptions. This is easily seen by reconsidering the proof of the Theorem.

For example, $\dfrac{\mathtt{a} \rightarrow \mathtt{b}, \mathtt{a} \rightarrow \mathtt{c}, \mathtt{d}}{\mathtt{a} \rightarrow \mathtt{b} \wedge \mathtt{c}}$ is topologically valid. (For all X and all $a, b, c \subseteq X$, $O(c, a) \cap O(c, b) \subseteq O(c, a \cap b)$.) Hence, if our person has concluded that $\mathtt{b}$-and-$\mathtt{c}$ is deducible from $\mathtt{a}$ on the grounds that both $\mathtt{b}$ and $\mathtt{c}$ are deducible from $\mathtt{a}$, this conclusion will not be withdrawn on the arrival of whatever new information $(\mathtt{d})$.

Hence, although the mathematical elaboration may leave room for improvement, Theorem 3.13 can be said to support the following tentative conclusion. The typical monotonic character of classical propositional logic is not a consequence of any of the following:

1) accepting the rule of monotony,
2) accepting that rule and all “classical laws”,
3) using an unambiguously defined formal language,
4) making a distinction between objectlanguage and meta-language, or between implication and inference,
5) using $\neg \ldots \vee \ldots$ to interpret implication on the objectlevel,
6) any combination of 1) - 5).

It is a consequence of *adopting monotonic reasoning habits on the metalevel*. The collection of axioms that is usually assumed to characterize propositional logic does not suffice to characterize these reasoning habits. Most probably, there is no other set of rules that manages to capture, completely, these habits, since any such set of rules will involve implicational statements, in one way or another, that will be susceptible of a non-strict interpretation. This latter statement, however, is no more than a conjecture, and is certainly not supported by Theorem 3.13. Let us state it, more provocatively, as follows: it is thinkable that monotonic reasoning defies complete axiomatization in the eyes of a person that persistently reasons non-monotonically.

§4. Comparison with other approaches. The properties of inference-as-a-rule-with-possible-exceptions, as found in the preceding section, rely on the choice of the implication operator. The operator $O(\ ,\ )$ that we used turned out to be quite suitable. In section 12 of [Jurjus, 1997] some alternatives for the implication operator are considered, resulting in the notions of U-validity, O_Π-validity, S-validity and T_Π-validity. It turns out that neither of these notions could have been used to obtain Theorem 3.13.

Definition 4.1. Let $(X, \leq)$ be a partial ordering and $a, b \subseteq X$. An element $x \in a$ is *minimal in a* whenever there is no $y \in a$ such that $y \leq x$ except x itself. We will use $a \mathrel{|\!\!\sim}_{\leq} b$ to denote: every element of a that is minimal in a is an element of b.

If $(X, \leq)$ is a partial ordering (of possible worlds), we may define, for $a, b \subseteq X$, $Dc(a,b) := \{w \in X \mid \underset{\wedge}{w} \cap a \mathrel{|\!\!\sim}_{\leq} b\}$ and $Uc(a,b) := \{w \in X \mid \hat{w} \cap a \mathrel{|\!\!\sim}_{\leq} b\}$ where, for every $w \in X$, $\underset{\wedge}{w} := \{v \in X \mid v \leq w\}$ and $\hat{w} := \{v \in X \mid w \leq v\}$, called the *downcone* and the *upcone* of w, respectively.

Each of these operators can be used to evaluate (sentences containing) nested conditionals in preferential models. In [Makinson, 1993], this is referred to as the *downcone construction* and the *upcone construction*, respectively. It will be clear that it is also possible to define downcone-validity and upcone-validity of *rules*, similar to our definition of O-validity.

The operator $Uc(\ ,\ )$ is, in general, not an implication operator. For, if $(X, \leq)$ is the following partial ordering:

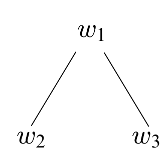

$a = \{w_1, w_3\}$ and $b = \{w_2, w_3\}$, then $a \mid\!\sim_{\leq} b$, but $w_2 \notin Uc(a, b)$, hence $Uc(a, b) \neq X$. Moreover, $X \not\mid\!\sim_{\leq} Uc(a, b)$. Hence, the upcone construction seems to be not very appropriate.

The downcone construction, on the other hand, is much more interesting. The operator $O(\ ,\ )$, as defined above, generalizes the downcone construction, in the following sense: if $(X, \leq)$ is a partial ordering, we may define $a \subseteq X$ is open whenever $\underset{\wedge}{w} \subseteq a$ for all $w \in a$. It is now easy to see that, with respect to this topology on $X, O(a, b) = Dc(a, b)$ for all $a, b \subseteq X$. Hence, every downcone operator is the $O(\ ,\ )$ operator of some topology. As a corollary, the downcone construction shares all phenomena found for $O(\ ,\ )$, in particular those indicated in Theorems 3.12 and 3.13. Although the downcone-validity of modus ponens and the rule of monotony was mentioned in [Boutilier, 1990], Theorem 3.13 was never noticed before, as far as we know.

Note that not every $O(\ ,\ )$ operator is the downcone operator of some partial ordering, not even on spaces for which there does exist a partial ordering providing the right consequence relation. For example, if $(X, \leq)$ is the following partial ordering:

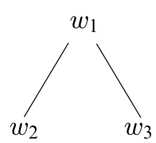

Then the definition above yields the topology $\tau_1 = \{\emptyset, \{w_2\}, \{w_3\}, \{w_2, w_3\}, X\}$, which has the property that (for all $a, b \subseteq X$) $a \to b$ (only) if $a \mid\!\sim_{\leq} b$. But the topology $\tau_2 := \{\emptyset, \{w_2, w_3\}, X\}$ also has this property. If we let $O_i(\ ,\ )$ denote the O-operator that is based on the topology τ_i, for $i = 1, 2$, then $O_1(\ ,\ )$ equals $Dc(\ ,\ )$, but $O_1(\ ,\ )$ does not equal $O_2(\ ,\ )$. (For example, $O_1(\{w_2\}, \emptyset) = \{w_3\}$, but $O_2(\{w_2\}, \emptyset) = \emptyset$.) Since $\leq$ is the only partial ordering on X yielding this consequence relation, there is no (other) partial ordering on X for which the downcone operator equals $O_2(\ ,\ )$. Thus, $O_2(\ ,\ )$ is not a downcone operator on the set X.

In general: on a single preferential model, there typically exist several different topologies providing the right consequence relation, each one leading to an $O(\ ,\ )$ operator, one of them being the downcone operator. Summarizing, we could say that, above, we found some new results about the downcone construction, and a number of other operators displaying the same behaviour. It is important, though, to note the following as well.

Makinson criticizes the downcone construction on the grounds that "like the upcone one, (it) does not correspond closely to the basic idea underlying normality semantics". In other words, the downcone construction is not well motivated for preferential semantics. From our topological perspective, however, the operator $O(\ ,\ )$ is very natural; $O(a, b)$ is the largest region in which $a \to b$ holds, where it is taken for granted that a region is an open set. In a preferential model, however, this identification is not so natural. Essentially, because the topological space associated with a preferential model will typically be a "pathological" space: a structure that happens to satisfy the requirements of a topological space without being a "carrier of geometrical intuition". In short, the downcone construction is better motivated by our topological presentation than by preferential semantics.

One more (important) reason to prefer the topological intuition over the preferential idea, is the following. To appreciate the interpretative remarks concerning our imaginary person, it is necessary to appreciate the non-standard conditional as a plausible interpretation of inference. While preferential semantics depicts the defeasible conditional as something that takes only the most normal possibilities into account (which may be a negligible minority among all possibilities), the topological semantics pretends to take every possibility into account, be it in a not too pedantic way. It is not claimed that our construction is entirely convincing as a plausible interpretation of inference. But the downcone construction (that is, the preferential interpretation of the operator $O(\ ,\ )$), does not support the remarks about our imaginary person at all.

In [Gabbay, 1995], a "fibring" construction was used to evaluate nested conditional sentences. This construction, unlike the downcone construction, does correspond to the idea underlying normality semantics.

In [Jurjus, 1997], some topological variants of this fibring construction were investigated, called T_Π-validity, O_Π-validity and O'-validity. The idea behind these constructions was to work with a 'topological space of topological spaces'. Two of these variants turned out to be unsuitable to obtain results like Theorem 3.13. The other one (O'-validity) turned out to be equivalent to topological validity.

Some existing approaches towards nested defeasible conditionals are strongly connected with modal logic. For example, the main advantage of the downcone construction is its connection with the modal system S4. This connection has a topological equivalent. We will not give a detailed exposition, since it is a straightforward adaptation of the connection as described in [Boutilier, 1990]. It is based on the fact that for all topological spaces, X, and all $a, b \subseteq X$,

$$O(a,b) = \left((a^c)^o \cup \overline{(a \cap (a^c \cup b)^o)}\right)^o$$

(where a^o and $\bar{a}$ denote the interior and the closure of a, respectively). As is well-known, the calculus of $\cap$, $\cup$, ^{c} and ^{o} in topological spaces is equivalent to the modal system S4, ^{o} playing the role of the $\Box$-modality.

Hence, the formula above gives rise to a translation from L into the modal language of S4 such that every formula of L is O-valid (only) if its translation is a tautology of S4. (As a corollary of Boutilier's work and the remarks above, rules are downcone-valid (only) if they are topologically valid.)

This and similar connections have lead some people ([Lamarre, 1991], [Boutilier, 1994]) to tackle the task of finding suitable strengthenings of P1–P5 by studying extensions of S4. For example, [Boutilier, 1990] pointed out that joining R3 to P1–P5 corresponds to using the modal system known as S4.3. To incorporate the most important adjustments, however, this approach will not be general enough. For example, if $\dfrac{\mathtt{a} \to \mathtt{b}}{\mathtt{a} \wedge \mathtt{c} \to \mathtt{b}}$ is to be valid, but $\dfrac{\mathtt{a} \to \mathtt{b}}{\mathtt{a} \wedge \neg\mathtt{b} \to \mathtt{b}}$ is not, or if $\dfrac{\mathtt{a} \to \mathtt{b}, \mathtt{a}, \neg\mathtt{b}}{\mathtt{b}}$ is to be invalid, but $\dfrac{\mathtt{a} \to \mathtt{b}, \mathtt{a}, \mathtt{c}}{\mathtt{b}}$ is to be valid ($\mathtt{c}$ being a basic formula not occurring in $\mathtt{a}$ or $\mathtt{b}$), we will have to consider modal systems that do not satisfy the substitution theorem. Since very little is known about such modal systems, it seems useless to proceed in this direction. Likewise, it will be of little use to search for more convincing implication-operators, since this would necessarily lead to systems that do satisfy the substitution theorem.

Conclusions.

1. We introduced an implication-with-possible-exceptions and defined validity of rules-with-possible-exceptions by means of the topological notion of a full subset.
2. Our implication-with-possible-exceptions characterizes the preferential consequence relation as axiomatized by Kraus, Lehmann and Magidor [Kraus, Lehmann, and Magidor, 1990]. However, the class of topological models is strictly larger than the class of preferential models.
3. The resulting inference relation is non-monotonic. On the other hand, modus ponens and the rule of monotony, as well as all other laws of classical propositional logic, are valid-up-to-possible exceptions.
4. As a consequence, the rules of classical propositional logic do not determine the meaning of deducibility and inference as implication-without-exceptions.
5. The typical monotonic character of classical propositional logic is not a consequence of accepting the rule of monotony or accepting that rule and all classical laws; it is a consequence of adopting monotonic reasoning habits on the meta-level.

REFERENCES

[Abbot 69] J.C. Abbott, ***Sets, Lattices and Boolean Algebras***, Allyn and Bacon, Boston, 1969.

[Nait Abdallah 95] A. Nait Abdallah, ***The Logic of Partial Information***, Springer, Berlin / Heidelberg / New York, 1995.

[Adams 75] E.W. Adams, ***The Logic of Conditionals***, Reidel, The Netherlands, 1975.

[AGM 85] C. E. Alchourrón, P. Gärdenfors, and D. Makinson, *On the logic of theory change: partial meet contraction and revision functions*, this Journal, vol. 50 (June 1985), no. 2, pp. 510–530.

[Bacchus 90] F. Bacchus, ***Representing and Reasoning with Probabilistic Knowledge: a logical approach to probabilities***, MIT Press, Cambridge, 1990.

[Boutilier 90] C. Boutilier, *Conditional Logics of Normality as Modal Systems*, ***Proceedings AAAI-90***, vol. 2, AAAI Press / MIT Press, 1990, pp. 594–599.

[Boutilier 94] ———, *Conditional Logics of Normality: A Modal Approach*, ***Artificial Intelligence***, vol. 68 (1994), no. 1, pp. 87–154.

[Brewka 91] G. Brewka, ***Nonmonotonic Reasoning: Logical Foundations of Commonsense***, Cambridge University Press, 1991.

[Flach 95] P.A. Flach, *Conjectures: an inquiry concerning the logic of induction*, ***Ph.D. thesis***, Institute for Language Technology & Artificial Intelligence, Tilburg University, 1995.

[Gaal 64] S.A. Gaal, ***Point Set Topology***, Academic Press, New York - London, 1964.

[Gabbay 95] D.M. Gabbay, *Conditional Implications and Non-monotonic Consequence*, ***Conditionals: from philosophy to computer science***, Oxford University Press, New York, 1995, pp. 337–359.

[Geffner 92] H. Geffner, ***Default Reasoning: causal and conditional theories***, MIT Press, 1992.

[Jurjus 97] H.H. Jurjus, *The Exception Proves the Rule; Non-monotonic Logic via Topology*, ***Ph.D. thesis***, Tilburg University, Faculty of Philosophy, 1997.

[KLM 90] S. Kraus, D. Lehmann, and M. Magidor, *Nonmonotonic Reasoning, Preferential Models and Cumulative Logics*, ***Artificial Intelligence***, vol. 44 (1990), pp. 167–207.

[Lamarre 91] Ph. Lamarre, *S4 as the Conditional Logic of Nonmonotonicity*, ***Proceedings of the 2nd international conference (kr91)*** (J. Allen, R. Fikes, and E. Sandewall, editors), Kaufmann Publishers, 1991, pp. 357–367.

[Makinson 93] D. Makinson, *Five Faces of Minimality*, ***Studia Logica***, vol. 52 (1993), pp. 339–379.

[Shoham 88] Y. SHOHAM, *Efficient Reasoning about Rich Temporary Do-mains*, ***Journal of Philosophical Logic***, vol. 17 (1988), pp. 443–474.

DEPARTMENT OF PHILOSOPHY
TILBURG UNIVERSITY
P.O. BOX 90153, NL 5000 LE TILBURG, NETHERLANDS
E-mail: hhjurjus@hetnet.nl

DEPARTMENT OF PHILOSOPHY
KATHOLIEKE UNIVERSITEIT BRABANT
POSTBUS 90153, 5000 LE TILBURG, NETHERLANDS
E-mail: H.C.M.deSwart@kub.nl