A REDUCTION OF THE THEORY OF CONFIRMATION TO THE NOTIONS OF DISTANCE AND MEASURE

Karl Schlechta
Laboratoire d'Informatique de Marseille, URA CNRS 1787
CMI, Technopôle de Château-Gombert
F-13453 Marseille Cedex 13, France
ks@gyptis.univ-mrs.fr

December 15, 2008

Abstract

We present an analysis and formalization of confirmation of a theory through observation. The basic ideas are, first, to carry the results of single observations over to neighbouring cases by analogy, using an abstract distance relation as in the Stalnaker/Lewis semantics for counterfactual conditionals. A theory is then, in a second step, considered confirmed iff we have thus concluded positively for a "large" part of the universe - where "large" is interpreted by a weak filter. Formal semantics as well as sound and complete axiomatizations for the (trivial) first order and the propositional case are given.

1 Introduction

1.1 Overview of the paper

This article treats the problem of confirmation of a theory through observation of single cases.

The paper consists of two parts. The first part is more philosophical, and tries to defend the proposed idea - a combination of analogical reasoning and generalization - by arguing that it reflects scientific practice, that this practice is not unsound, and by defending the reductionist approach in general.

The second part is a formal treatment of the idea. A formal semantics is given, using a binary relation as abstract measure of distance, and a (weak) filter as abstract measure of size. Finally, a formal language and logic are given, and their correspondence with the semantics, i.e. soundness and completeness, is shown.

1.2 The basic idea

Our problem is the confirmation of a theory through observation of single cases, a classical problem of philosophy of science.
Our intuitive idea for an answer is very simple. The results of observations actually performed are carried over in a first step to "neighbouring cases" by analogical reasoning. "Neighbourhood" is determined by an abstract notion of distance, which is assumed to be given. In a second step, a resume is made, and if, in a sufficient number of cases, observation itself or the result of the analogical reasoning corresponds to the theory, the theory as a whole is considered to be confirmed. Whether the number of confirmed cases is sufficient, is determined by an abstract measure of size, a weak filter. This measure of size is again assumed to be given. In particular, "some" observations contrary to theory can be tolerated, if they are not "too many". The intuitive (and formal) semantics is thus built on two given primitives: an abstract notion of distance, and an abstract notion of size.
Note that the first step, i.e. the analogical reasoning, can be expressed by the counterfactual conditional "observation carried out $\Rightarrow$ result corresponding to theory". This will be used in the discussions of Section 2, as well as the formal development of Section 3.
Philippe Besnard, Rennes, has pointed out to me that N.Goodman, see [Goo55], had had this idea of expressing confirmation by counterfactuals too, before the advent of the Lewis/Stalnaker semantics (see e.g. [Lew73], [Sta68]). Our argument is, and was found, however, "one layer below" counterfactuals, and based directly on the intuition of distance.

1.3 Justification of the approach by scientific practice

It was Popper's insight that a large number of scientific theories have the logical form of a universally quantified formula, and consequently cannot be confirmed in the strict sense of classical logic: There are simply too many cases to consider over time and space. Popper concluded that many scientific theories can (at best) only be falsified, but never confirmed. Yet, scientists continue to speak about theories being confirmed by observation and experiment.
Our approach tries to capture in abstract terms some of this scientific practice of testing and confirmation of a theory. The choice of experiments in this process is not arbitrary. From a strictly numerical point of view, only the number of confirmed cases would count. This is obviously not the way of practice. For example, scientists will not perform the same experiment over and over again, unless there is reason to assume that the situation may change over time. So, not all cases are considered to be of equal weight. Moreover,

scientists will try to change the frame of an experiment as much as possible, choosing e.g. different substances for an experiment on gravity.

I think it is safe to say that underlying this procedure of testing is an assumption of inertia or homogenousness. The closer two situations are in many aspects, the more the outcome of experiments in both situations concerning a "related" aspect will be expected to coincide. More precisely, scientists have some confidence that "neighbouring cases" behave in the same way, and testing one of them is considered sufficient. One should note, that this distance relation defining "neighbourhood" encodes a large amount of scientific theories and experience. It is an abstraction of many things. It is, so to say, a conceptual interface between experiment and background theories. Our analysis halts at this border, we do not pretend to do more.

Scientists will often admit exceptions to theories. These exceptions might be interpreted by malfunctions of the measuring apparatus, by unknown disturbing effects, or as some true exceptions to the theory, which one cannot as yet explain.

We have therefore choosen a weak filter (or a weak universal quantifier) instead of the universe itself (or the classical universal quantifier) to describe when a theory as a whole is considered confirmed by individual observations or the results of analogical reasoning.

To put things differently, scientists follow reasonable strategies to test and confirm theories. (This might also be an analysis of Hempel's black raven paradox: Testing ravens for blackness seems to be a much better strategy than testing non-black objects for non-ravenness. Strategies are then not automatically prescribed by the logical form of the theory to be confirmed - upon closer reflection, this seems quite obvious.) Scientists choose "strategically placed" cases, and generalize by (probable) analogy to their neighbourhood. This strategy is encoded in our approach into the chosen measures of distance and size.

We now attempt to give a pragmatic justification of this scientific pratice. First, an examination of all cases is impossible. So, we can either resign, or try to do the best, to our knowledge. But we want to act, make predictions, etc. We cannot test all cases, so we have to do with less. We have observed that some cases behave in some aspects the same way. It seems a good strategy to choose one case for each such "equivalence class" of cases, defined by analogous behaviour in some aspects, test this "prototypical" case, and generalize to the rest. Furthermore, we want to organize our knowledge, a "workable" theory is often sufficient from a pragmatic point of view. We are modest, and do not claim "divine insight" into truth, we work our way along. So we admit some exceptions, and hope for future refinements.

1.4 A justification of the reductionist approach

We reduce the problem of confirmation to the problem of choosing the right distance and size, thus reduce one concept to others. Such a reductionist approach does not claim absolute explanations. It rather tries to establish connections. Moreover, it can point out similarities, by reducing different concepts to the same basic ideas. For instance, the Lewis/Stalnaker semantics for counterfactual conditionals, preferential semantics for nonmonotonic and deontic logic reduce to the notion of distance, so does our semantics for one part of confirmation. Consequently, results from one field can be carried over to other cases. (We do this below, e.g. by using the techniques for constructing preferential models to give a completeness result for counterfactual conditionals.)

A reduction of one notion to another can also help clarify intuitive adequacy. Take for example a logic which is sound and complete for a certain semantics. It seems to be much easier to judge the intuitive value of a semantics (i.e. its correspondence to an intuitive abstraction of the situation considered) than to judge directly the value of a logic. The particular notions of distance and size are perhaps easier to judge for adequacy than the more complex notion of confirmation.

Finally, in defense of this approach, it may be said that also natural sciences do not pretend to explain things in an absolute way, but they describe laws governing the relations between notions, e.g. between mass, speed, and forces. So we are in a good and fruitful tradition.

Of course, this defense of the reductionist approach is no argument against explaining some notion by still other means.

1.5 Conclusion

Criticism of the intuition Perhaps the most evident criticism of our basic approach as too simplistic is, that it does not allow any grading between different degrees of confirmation. More generally, the idea of absolute confirmation of a theory is perhaps a philosophical artifact, which tries to find certainty where it does not exist: in the realm of the real world. Scientists speak about one theory being better confirmed than another one. We can enrich our approach at two points to incorporate such grading. First, in the analogical part, we can measure the confidence in the analogical reasoning by the closeness of observed cases. (This extension applies to counterfactual conditionals too: We can measure confidence in counterfactual reasoning in the same way.) Second, in the part of generalization, we can still differentiate between “large” subsets of different sizes. Thus, the approach seems to be apt to further development and research.

See the end Section 2, for a somewhat more detailed discussion.

Conclusion of the formal part From a more abstract point of view, the problem and its formal development seem to show two things.

First, that the modest approach, developing first an intuition about the problem, condensing this intuition into a formal semantics, and then developing a logic which is sound and complete for that semantics, results at least in sound handcraft.

Second, it seems to confirm the author’s view that there are some, perhaps not too many, basic semantical concepts underlying many different nonclassical logics and reasoning mechanisms. The notions of distance and measure are most probably among those concepts: The notion of distance is basic to the Stalnaker/Lewis semantics of counterfactual conditionals, and preferential semantics for nonmonotonic and deontic logics can be seen as expressing an abstract distance from an ideal point of maximal "normality" or "morality". It seems to be a basic notion for analogical reasoning, too. The notion of size seems to be essential for "soft" generalizations, as in the present case, and basic to analysis of the "usual case", giving a semantics to another aspect of default or nonmonotonic reasoning.

1.6 Technical outline of the paper

We first describe the logical type of theories we want to examine - certain cases will be excluded, either presenting unnecessary complications for a first approach, or being not really theories apt to confirmation by single cases. We then prepare the logical framework by describing the kind of reasoning we want to be able to perform, leading to an informal description of the semantics and the language.

We then turn to the formal part, treat the first order case, and the - essentially - propositional case (with some quantifiers added). In the first case, there is nothing to show, as we can formulate conditionals in the first order framework via the underlying order. In the second case, we give a completeness proof, organized in several layers, according to the definition of the language. En passant, we also give a new completeness result for preferential models with a (global) smallest element, as well as a new completeness proof for flat conditionals.

We conclude by discussing the KLM axioms (see e.g. [Gab85], [KLM90], [LM92]) in our framework and show that several crucial axioms cannot hold for very strong reasons.

2 Making the problem precise

2.1 Limitations of the approach

It seems necessary to describe some restrictions and clear up some notions, which we do here in the first order framework, the treatment carries over without difficulties to the propositional case.

We work in a fixed universe U, for a language $\mathcal{L}$ which still needs some specifying, assume for the moment that it is the first order language, with a constant for each $a \in U$ - which we will denote for simplicity a too.

We are interested in theories which can be confirmed and disconfirmed locally.

(1) A theory which is uniformly true or uniformly false in our universe U (containing e.g. $\forall x \forall y (\phi(x) \leftrightarrow \phi(y))$) shall not interest us.

(2) A theory like $\{\exists x \phi(x)\}$ cannot be disconfirmed locally - unless the universe has only one element.

(3) A theory like $\{\forall x \phi(x) \vee \forall x \neg \phi(x)\}$ cannot be falsified locally - we need to examine at least two elements.

Whereas theories described in (1) and (2) can be considered degenerate cases, (3) presents a more complicated situation, where the basic units of observation consist of more than one element, and, in principle, the formal development can also be carried out over such units (a relation of distance between pairs of elements etc.), but this would rather complicate things without real necessity. One should bear in mind, however, that many "real life" theories express interdependencies between elements in the universe.

We turn these negative examples into positive - semantic - conditions:

For simplicity, we consider "real" universal formulas of the type $\forall x \phi(x)$. (In a generalization, we might also consider a formula ϕ for which there is an associated set $\phi_U := \{\phi_a : a \in U\}$ of formulas, which replace the $\phi(a)$.)

(a) (Equivalence with ϕ) For any interpretation I of $\mathcal{L}$ in to U, $\forall x \phi(x)$ holds in I iff all $\phi(a)$ hold in I.

(b) (Independence of the $\phi(a)$) For an arbitrary truth assignment τ to the individual $\phi(a)'s$ there is an interpretation I_τ of $\mathcal{L}$ into U which respects it, i.e. I_τ makes $\phi(a)$ true iff τ assigned true to $\phi(a)$.

The above examples are excluded: (1) Choose I and I' making all $\phi(a)$ true or making at least one $\phi(a)$ false. (2) and (3) Choose for $a \in U$ I_a making all $\phi(b)$ true but for $b = a$.

2.2 The problem of observation

We want to say something like: "Wherever we have made an observation, we found ravens to be black". But what kind of observation? If we have counted their toes, this will tell us nothing about their (the raven's) colour. So we are interested in *relevant* observation. We shall presume an observation relevant to ϕ to be a partition of truth (i.e. a set of mutually exclusive, consistent formulas ψ_i, whose disjunction is equivalent to truth), which permit to judge on ϕ, i.e. at least one of the $\psi_i's$ permits to deduce ϕ, and one $\neg\phi$.

This is like a scale from 0 to 10, one value has to hold, and we can say whether we have measured 5 or not. We further assume our observations to be correct. So we have really examined colour, and we found black, or red, or whatever. Some observations might be inconclusive, i.e. consistent with ϕ and $\neg\phi$, we just do as if we had not observed.

So we want to say: "We have made a ϕ-relevant observation, and we have found indeed ϕ''.

2.3 The intended type of reasoning

We would like to conclude from evidence to hypotheses. Evidence shall be a formula describing a state of affairs at a point in the universe. A hypothesis shall essentially be a universally quantified formula. As we intend to pass via analogy, by a proximity relation, we express conditionals which permit to conclude from the state of affairs at $a_1 \ldots a_n$

to the state of affairs at a. It does not seem necessary to treat nested conditionals, but boolean combinations of conditionals are desirable. We want to express that a property holds "almost everywhere", so we extend first order language by the generalized quantifier ∇ of [Sch89-n1].

We soften the conditions (i.e. ∇ instead of $\forall$) to tolerate exceptions, otherwise, as each element is closest to itself in any reasonable notion of distance, we could never confirm a theory which has one counterexample. We choose to apply ∇ after the conditional $\Rightarrow$, because a small set of counterexamples might have devastating effects on a theory, if they cover "strategic points" in the sense that there are many other points without observation, which are very close to them.

To elaborate further, we want to say "had we made a ϕ-relevant observation at a, we would have found ϕ at a" in counterfactual terms, or, in more detail, and more intuitive too: "At the elements b closest to a, where we have made a ϕ-relevant observation, we have found $\phi(b)''$. There is still a certain impreciseness there, because we am really speaking about $\phi(b)$-relevant observations. So, finally, we want to say: "At the elements b closest to a, where we have made a $\phi(b)$-relevant observation, we have found $\phi(b)''$. We shall abbreviate this by $a \models (o_\phi \Rightarrow \phi)$ - the conditional evaluates to true at a.

Reasoning will then be e.g. of the form: If for all b $\phi(b) \rightarrow \psi(b)$, and, at the elements b closest to a, where we have made a $\phi(b)$-relevant observation, we have found $\phi(b)$, so we have $a \models (o_\phi \Rightarrow \phi)$, then we can conclude by analogy also to $\psi(a)$. (Note that in this argument, we use that any $\psi(b)$-negative observation will also be a $\phi(b)$-negative observation.)

To summarize: We have made a $\phi(a)$-relevant observation at some $a \in U$, we have found in some cases indeed $\phi(a)$, some were inconclusive, and some were counterexamples. We now use this evidence for reasoning by analogy, formally, we examine for which a $a \models (o_\phi \Rightarrow \phi)$ holds, and evidence supports ϕ iff $\nabla x.x \models (o_\phi \Rightarrow \phi)$ holds.

2.4 Possible extensions and further development

Elaboration of the analogical reasoning Some elements x of the universe might be too distant from a given element a to reasonably count in analogical reasoning concerning a. We might then wish to add an unary predicate $d_a(x)$ for each element a, which expresses that x is too far away from a to count. Well-behaviour with $\prec_a$ will be needed: $d_a(x) \wedge (x \prec_a y) \rightarrow d_a(y)$. In the propositional case, d_a will just be a new propositional variable.

We might also be interested in some kind of degree of confirmation. Suppose we have two theories, ϕ and ψ, and look at one point a, where both are confirmed. It seems to be reasonable to say that, at a, ϕ is at least as well confirmed by analogy as ψ is at a, if the positive observations of ϕ are at least as close to a as the positive observations of ψ. This can be expressed by the conditional $a \models (o_\phi \vee o_\psi \Rightarrow o_\phi)$. Note that our formalization - even in the propositional case - will allow us to express that everywhere, or on a large subset, ϕ is at least as well confirmed as is ψ. This seems to be a generalization of Carnap's

idea, as we work with a general partial order, which is not necessarily ranked.
More generally, we can develop a theory of strength of belief into a conditional formula at some x. It seems a reasonable first approach to measure this strength of belief by the set of worlds that enter into the evaluation. The suitable construction will be a tree, because the truth of e.g. $\neg(\sigma \Rightarrow \tau)$ can be demonstrated separately by several worlds, the same holds for disjunctions.

Confirming a "nearby" theory We might consider a theory confirmed in a wider sense, iff "close to it", there is a theory confirmed in our, more restricted sense. We then need a measure of distance between theories as discussed in [Sch92-n1].

3 Definitions and results

The First Order case is disappointing from the formal point of view: There is nothing to do! The propositional case is a little more interesting, but there is not much work to do either: We can largely use the ideas, results and techniques of [Sch89-n1] and [Sch92] to construct a completeness result as wanted.

3.1 The first order case

So assume to be given a first order language with a constant for each element of the universe (by Loewenheim/Skolem/Tarski countably many will do - and the completeness proof for ∇ in [Sch89-n1] works via classical models, so it is accessible to the same arguments), generalized quantifier ∇ and a binary relation symbol $\prec_a$ for each element a.
"$a \models \alpha \Rightarrow \beta$" can now be expressed by the $\Rightarrow$-free formula $(\alpha(a) \wedge \beta(a)) \vee \neg\alpha(a) \wedge \forall x(\alpha(x) \wedge \neg\exists y(y \prec_a x \wedge \alpha(y)) \rightarrow \beta(x))$.
So there is no need to do any further work.
On the other hand, the fact that we can express finite cardinalities in First Order Logic (FOL) (by e.g. $\exists x_1 \ldots x_n \forall x(\alpha(x) \rightarrow (x = x_1 \vee \ldots \vee x_n))$), shows that we are not free to do all kinds of operations we might want to: In particular, we cannot multiply elements by making copies as is often done in completeness proofs.
For the reader's convenience, we repeat the definitions and the main result of [Sch89-n1] - they will also be used (in slightly modified form) in the propositional case. On the semantic side, we formalize "large" subsets by a kind of weak filter, an $\mathcal{N}$-system. On the syntactic side, we introduce a new, generalized, quantifier, ∇, which should be read (approximatively) as "most". Soundness and completeness say that they correspond. For further motivation and discussion, the reader is referred to the original article.

Definition of the semantics

Definition 3.1 Call $\mathcal{N}(M) \subseteq \mathcal{P}(M)$ (= the powerset of M) a $\mathcal{N}$-system over M iff
a. $M \in \mathcal{N}(M)$,
b. $A \in \mathcal{N}(M)$, $A \subseteq B \subseteq M \rightarrow B \in \mathcal{N}(M)$,
c. $A, B \in \mathcal{N}(M) \rightarrow A \cap B \neq \emptyset$ if $M \neq \emptyset$ (thus, $\emptyset \notin \mathcal{N}(M)$, if $M \neq \emptyset$)
(Note that this is weaker than the corresponding axiom for filters.)

Remark: Our semantics covers the two extremes:
- fix one element a of the universe U, then $\{A \subseteq U\colon\ a \in A\}$ will be a $\mathcal{N}$-system,
- let some probability measure be given on U, then $\{A \subseteq U\colon\ p(A) > 0.5\}$ will be a $\mathcal{N}$-system.

(Note, however, that the former can also be expressed by a suitable point measure on U.)
We can thus cover both the "prototypical" and the "average" case.
To facilitate proofs and enable normal forms, we introduce a complementary quantifier, $\clubsuit$, too, with the meaning $\clubsuit x\phi(x) :\leftrightarrow \neg\nabla x\neg\phi(x)$. The intuitive reading of $\clubsuit x\phi(x)$ is thus roughly: "for at least a few x, $\phi(x)$ holds".

Definition 3.2 We augment the language of first order logic by the new quantifiers: If ϕ and ψ are formulas, then so are $\nabla x\phi(x)$, $\clubsuit x\phi(x)$, $\nabla x\phi(x) : \psi(x)$, $\clubsuit x\phi(x) : \psi(x)$ for any variable x. The results and definitions for the relativized versions are straightforward generalizations of the simple case. We call any formula of $\mathcal{L}$, possibly containing ∇ or $\clubsuit$ a $\nabla - \mathcal{L}-$ formula.

Definition 3.3 Let $\mathcal{L}$ be a first order language, and M be a $\mathcal{L}-$ structure. Let $\mathcal{N}(M)$ be a $\mathcal{N}$-system over M. Define $< M, \mathcal{N}(M) > \models \phi$ for any $\nabla - \mathcal{L}-$ formula inductively as usual, with two additional induction steps:
$< M, \mathcal{N}(M) > \models \nabla x\phi(x)$ iff there is $A \in \mathcal{N}(M)$ such that $\forall a \in A$ ($< M, \mathcal{N}(M) > \models \phi[a]$),
$< M, \mathcal{N}(M) > \models \clubsuit x\phi(x)$ iff $\{a \in M\colon < M, \mathcal{N}(M) > \models \neg\phi[a]\} \notin \mathcal{N}(M)$.

Lemma 3.4 $< M, \mathcal{N}(M) > \models \clubsuit x\phi(x)$ iff $\forall A \in \mathcal{N}(M) \exists a \in A(< M, \mathcal{N}(M) > \models \phi[a])$. $\Box$

Definition 3.5 Let any axiomatization of predicate calculus be given. Augment this with the axiom schemata
1. $\nabla x\phi(x) \wedge \forall x(\phi(x) \rightarrow \psi(x)) \rightarrow \nabla x\psi(x)$,
2. $\nabla x\phi(x) \rightarrow \neg\nabla x\neg\phi(x)$,
3. $\forall x\phi(x) \rightarrow \nabla x\phi(x) \rightarrow \exists x\phi(x)$,
4. $\clubsuit x\phi(x) :\leftrightarrow \neg\nabla x\neg\phi(x)$,
5. $\nabla x\phi(x) \leftrightarrow \nabla y\phi(y)$ if x does not occur free in $\phi(y)$ and y does not occur free in $\phi(x)$
(for all ϕ, ψ).

Lemma 3.6 The following formulae are derivable:
a. $\nabla x\phi(x) \wedge \nabla x\psi(x) \rightarrow \exists x(\phi \wedge \psi)(x)$,

b. $\nabla x\phi(x) \wedge \neg\nabla x\psi(x) \rightarrow \exists x(\phi \wedge \neg\psi)(x)$,
c. $\neg\nabla x\neg\phi(x) \rightarrow \exists x\phi(x)$,
d. $\clubsuit x\phi(x) \rightarrow \exists x\phi(x)$,
e. $\nabla x\phi(x) \wedge \clubsuit x\psi(x) \rightarrow \exists x(\phi \wedge \psi)(x)$,
f. $\forall x(\phi(x) \leftrightarrow \psi(x)) \rightarrow (\nabla x\phi(x) \leftrightarrow \nabla x\psi(x)) \wedge (\clubsuit x\phi(x) \leftrightarrow \clubsuit x\psi(x))$,
g. $\forall x\phi(x) \rightarrow \clubsuit x\phi(x)$.
It is usually *not* derivable: $\clubsuit x\phi(x) \wedge \clubsuit x\psi(x) \rightarrow \exists x(\phi \wedge \psi)(x)$. (To see this, use Theorem 3.7 below and argue semantically.) $\Box$

Theorem 3.7 The axioms given in Definition 3.5 are sound and complete for the semantics of Definition 3.3.

3.2 The (semi-) propositional case

We shall modify our completeness proofs of [Sch92] and [Sch89-n1] to adapt to counterfactuals.
The first step uses largely ideas from [Sch92], which, as a matter of fact, need only a slight adaptation to account for the existence of a (global) smallest model. So, en passant, we also find a new completeness proof for another version of preferential reasoning. We adapt the proof to cover flat counterfactual conditionals, the essential difference is that we work with single formulas, not with arbitrary theories. Our central algebraic result in [Sch92] is general enough to permit this adaptation. This gives a new completeness proof for flat conditionals.

Outline of the proof: We first give a completeness proof for theories without the generalized quantifiers, and use for the full version the $\nabla-$ and $\clubsuit$-free consequences to construct a model by the first completeness result, on which we define a $\mathcal{N}$-system, and show that in the full structure the full theory holds.
If you look more closely, you will observe that this is a three layer proof: In the first layer, we have completeness for classical propositional logic. With that result, we construct - via the main algebraic characterization of preferential structures - the completeness for counterfactual conditionals. (Soundness and completeness of classical propositional logic with respect to classical models is ALL we need for the underlying logic - so we could replace it by any other logic! In the single formula version of the counterfactuals, we will also use its compactness.) In the third layer, we transform a consistent ∇-theory into a consistent $\forall/\exists$-theory one layer below, and construct the final result on that basis.

Definition of the language We work in a fragment of FOL with a generalized quantifier - in order to be able to do some reasoning. If you like, impoverish it. (What we really need is to differentiate e.g. between $\forall m(\phi \vee \psi)(m)$ and $(\forall m\phi(m)) \vee (\forall m\psi(m))$.

Essentially, we can replace $\forall m\phi(m)$ by $R \vdash \phi$ and $\exists m\phi(m)$ by $Con(R,\phi)$ - where R is a classical background theory to be specified below.

Definition 3.8 Level 0: a classical propositional language.
Level 1: formulas of the type $m \models \alpha$ or $m \models \alpha \Rightarrow \beta$, where α and β are in Level 0, ($\Rightarrow$ expresses a conditional, and m, $\models$ have the usual meaning) and boolean combinations of such formulas.
Level 2: formulas of the type $\forall m\phi(m)$, $\exists m\phi(m)$, $\nabla m\phi(m)$, $\clubsuit m\phi(m)$, where $\phi(m)$ are of Level 1, and boolean combinations of such formulas.

We might also introduce constants for models, but this does not really change the picture.

Definition of the $\nabla-$ semantics We can take over almost all definitions and results of the first order case, with a slight modification of Definition 3.3 and the subsequent Lemma 3.4.

Definition 3.9 ($\mathcal{N}-$ Model)
Let $\mathcal{L}$ be a propositional language, M be a set of $\mathcal{L}-models$, and $< M, \{\prec_a : a \in M\} >$ be a structure for level 1 of the language, let $\mathcal{N}(M)$ be a $\mathcal{N}$-system over M. Define $< M, \{\prec_a : a \in M\}, \mathcal{N}(M) > \models \nabla m\phi(m)$ ($\clubsuit m\phi(m)$) with two additional induction steps:
$< M, \{\prec_a : a \in M\}, \mathcal{N}(M) > \models \nabla x\phi(x)$ iff there is $A \in \mathcal{N}(M)$ such that $\forall a \in A$ $(< M, \{\prec_a : a \in M\} > \models \phi[a])$,
$< M, \{\prec_a : a \in M\}, \mathcal{N}(M) > \models \clubsuit x\phi(x)$ iff $\{a \in M : < M, \{\prec_a : a \in M\} > \models \neg\phi[a]\} \notin \mathcal{N}(M)$.

Lemma 3.10 $< M, \{\prec_a : a \in M\}, \mathcal{N}(M) > \models \clubsuit x\phi(x)$ iff $\forall A \in \mathcal{N}(M) \exists a \in A (< M, \{\prec_a : a \in M\} > \models \phi[a])$. $\Box$

The logic for ∇ and $\clubsuit$: We can take verbatim the above axioms, where FOL is suitably restricted. The crucial Lemma for the Completeness Theorem can be simplified, as we have only flat $\nabla's$.
We now repeat the basic definitions of [Sch92] on preferential structures.

Definition 3.11 We use $\mathcal{P}$ to denote the power set operator, $\Pi\{X_i : i \in I\} :=$ {g: $g : I \rightarrow \bigcup\{X_i : i \in I\}$, $\forall i \in I . g(i) \in X_i\}$ is the general cartesian product, card(X) shall denote the cardinality of X, and V the set-theoretic universe we work in - the class of all sets. Given a class of pairs $\mathcal{X}$, and a set X, we denote by $\mathcal{X}\lceil X := \{< x, i > \in \mathcal{X} : x \in X\}$, so if $\mathcal{X}$ is a function f, $f\lceil X$ is the usual notation for the restriction of f to a subset of its domain.

Let $\mathcal{L}$ be a propositional language, we denote by $v(\mathcal{L})$ the set of its variables, by $M_{\mathcal{L}}$ the set of its classical models, ϕ etc. shall denote formulas, T etc. theories in $\mathcal{L}$ (i.e. $T \subseteq \mathcal{L}$), and $M_T \subseteq M_{\mathcal{L}}$ the models of T.
For any classical model m, let $Th(m)$ be the set of formulas valid in m, likewise $Th(M) := \{\phi : m \models \phi$ for all $m \in M\}$, if M is a set of classical models. For two theories T and T', let $T \vee T' := \{\phi \vee \psi : \phi \in T, \psi \in T'\}$. $\overline{T} \subseteq \mathcal{L}$ will denote the closure of T under classical logic, and $\vdash$ the classical consequence relation. Given some other logic, $\overline{\overline{T}}$ will denote the set of consequences of T under that logic, i.e. if the more conventional notation for the logic is $\mid\sim$, then $\overline{\overline{T}} := \{\phi : T \mid\sim \phi\}$.
$\boldsymbol{D}_{\mathcal{L}} \subseteq \mathcal{P}(M_{\mathcal{L}})$ shall be the set of definable subsets of $M_{\mathcal{L}}$, i.e. $A \in \boldsymbol{D}_{\mathcal{L}}$ iff there is some $T \subseteq \mathcal{L}$ such that $A = M_T$. If the context is clear, we omit the subscript $\mathcal{L}$ from $\boldsymbol{D}_{\mathcal{L}}$. □

Definition 3.12 $\mathcal{Z} =< \mathcal{X}, \prec>$ will be called a preferential structure iff $\mathcal{X}$ is a set of pairs and $\prec$ is a binary relation on $\mathcal{X}$. We say that $\mathcal{Z}$ is transitive, irreflexive etc., iff $\prec$ is.
$< y, i >$ is called a minimal element of $\mathcal{X} \lceil Y$ in $\mathcal{Z}$ iff:
1. $< y, i > \in \mathcal{X} \lceil Y$ and
2. there is no $< y', i' > \in \mathcal{X} \lceil Y$ such that $< y', i' > \prec < y, i >$.
Thus, $\mathcal{Z}$ defines a function $\mu_{\mathcal{Z}} : V \rightarrow V$ (V the set-theoretic universe) by $\mu_{\mathcal{Z}}(Y) := \{y :$ there is i such that $< y, i >$ is a minimal element of $\mathcal{X} \lceil Y\}$.
Given a set Z, $\mu_{\mathcal{Z},Z}$ shall denote $\mu_{\mathcal{Z}} \lceil \mathcal{P}(Z)$. □

Definition 3.13 A preferential structure $\mathcal{M} =< \mathcal{X}, \prec>$ will be called a classical preferential model (cpm) for $\mathcal{L}$, iff for all $< x, i > \in \mathcal{X}$, $x \in M_{\mathcal{L}}$.
$\mathcal{M}$ will be called definability preserving (dp) iff $\mu := \mu_{\mathcal{M},M_{\mathcal{L}}} : \mathcal{P}(M_{\mathcal{L}}) \rightarrow \mathcal{P}(M_{\mathcal{L}})$ is definability preserving.
By the above, $\mathcal{M}$ defines a logic on $\mathcal{L}$ by $T^{\mathcal{M}} := T^{\mu}$, i.e. $T^{\mathcal{M}} := \{\phi \in \mathcal{L}$: ϕ holds in all $m \in \mu(M_T)\}$.
A logic $^{=}$ for $\mathcal{L}$ is said to be representable by a cpm, iff there is a cpm $\mathcal{M}$ for $\mathcal{L}$, such that for all $T \subseteq \mathcal{L}$ $T^{\mathcal{M}} = \overline{\overline{T}}$.
For $< m, i > \in \mathcal{X}$, we shall abuse notation and say $< m, i > \models \phi$ iff $m \models \phi$, for $\phi \in \mathcal{L}$. □

We first show an algebraic representation result.

Proposition 3.14 Let $x \in Z$ be fixed, $\mathcal{Y} \subseteq \mathcal{P}(Z)$, $f : \mathcal{Y} \rightarrow \mathcal{P}(Z)$, then there is a preferential structure $\mathcal{Z} :=< \mathcal{X}, \prec>$ with one x-copy as smallest element of the whole structure (i.e. minimal and below all other elements) such that for all $X \in \mathcal{Y}$ $f(X) = \mu_{\mathcal{Z}}(X)$ iff
(f1) $f(X) \subseteq X$,
(f2) $X \subseteq Y \rightarrow f(Y) \cap X \subseteq f(X)$,
(f3) $x \in X \rightarrow$ f(X)={x}
for all $X, Y \in \mathcal{Y}$.
$\mathcal{Z}$ can be chosen irreflexive and transitive.

Proof We extend the proof of Proposition 3.4 in [Sch92].
"$\rightarrow$": (f1), (f2) hold in all preferential structures. (f3): If $x \in X$, then $< x, i >$ for some i will be below all other, and itself minimal, so $\mu_{\mathcal{Z}}(X) = \{x\}$.
"$\leftarrow$": By (f1) and (f2) and Proposition 3.4 of [Sch92], we can represent f by a suitable structure $\mathcal{Z}$. We add a new copy $< x, i >$ to $\mathcal{Z}$, making the new element the smallest of all. Obviously, irreflexivity and transitivity are preserved. It remains to show that the new structure $\mathcal{Z}'$ still represents f. If $x \notin X$, then $\mu_{\mathcal{Z}}(X) = \mu_{\mathcal{Z}'}(X)$. If $x \in X$, then $\mu_{\mathcal{Z}'}(X) = \{x\}$, but by prerequisite f(X)={x} too. $\Box$

Proposition 3.15 Let $R \subseteq S$ be $\mathcal{L} - \mathit{theories}$. Consider the following conditions for a function f on a subset of $\mathcal{P}(M_{\mathcal{L}})$ and a logic $^{=}$:
(f1) $f(X) \subseteq X$,
(f2) $X \subseteq Y \rightarrow f(Y) \cap X \subseteq f(X)$,
(f3) $x \in X \rightarrow$ f(X)={x},
(f4) f is dp
for all $X, Y \in \mathcal{Y}$,
and
($^{=}$1) $\overline{T \cup R} = \overline{T' \cup R} \rightarrow \overline{\overline{T}} = \overline{\overline{T'}}$,
($^{=}$2) $\overline{\overline{T}}$ is closed under the underlying logic $^{-}$,
($^{=}$3) $T \cup R \subseteq \overline{\overline{T}}$,
($^{=}$4) $\overline{T \cup R} \subseteq \overline{T' \cup R} \rightarrow \overline{\overline{T'}} \subseteq \overline{\overline{\overline{T}} \cup T' \cup R}$,
($^{=}$5) There is a consistent complete theory U with $S \subseteq \overline{U}$ such that $U \vdash T \rightarrow \overline{\overline{T}} = \overline{U}$
for $T, T' \in \mathcal{T}$.
(a) Given $\mathcal{Y} \subseteq \mathcal{P}(M_R)$, $m \in M_R$, $m \models S$, $f : \mathcal{Y} \rightarrow \mathcal{P}(M_R)$ such that (f1)-(f4) hold for $m = x$, then the logic defined on $\mathcal{T} := \{T : M_{T \cup R} \in \mathcal{Y}\}$ by $\overline{\overline{T}} := Th(f(M_{T \cup R}))$ satisfies ($^{=}$1) – ($^{=}$5).
(b) Given $\mathcal{T}$ and $^{=}$ satisfying ($^{=}$1) – ($^{=}$5), then the function defined on $\mathcal{Y} := \{M_{T \cup R} : T \in \mathcal{T}\}$ by $f(M_{T \cup R}) := M_{\overline{\overline{T}}}$ satisfies for $m := M_U$ (f1)-(f4).

Proof: We extend the proof of Proposition 3.9 in [Sch92]. As a matter of fact, we need not really show ($^{=}$1) – ($^{=}$4) and their corresponding properties (f1), (f2), (f4), we could just re-interpret the operator $\overline{T}$, which denotes classical closure there, by closure of $T \cup R$ by classical logic. As we only need soundness and completeness of the underlying logic with respect to the points in the structure, we are done. But this might be a little unfair to readers who want to see the details, so we present them.
Note that by dp for (a) and by the definition of f for (b), we always have $f(M_{T \cup R}) = M_{\overline{\overline{T}}}$ for all $T \in \mathcal{T}$.
(a) ($^{=}$1) : $\overline{T \cup R} = \overline{T' \cup R} \rightarrow M_{T \cup R} = M_{T' \cup R} \rightarrow f(M_{T \cup R}) = f(M_{T' \cup R}) \rightarrow \overline{\overline{T}} = \overline{\overline{T'}}$. ($^{=}$2) is trivial. ($^{=}$3) : $f(M_{T \cup R}) \subseteq M_{T \cup R} \rightarrow T \cup R \subseteq \overline{\overline{T}}$. ($^{=}$4) : Let $\overline{T' \cup R} \subseteq \overline{T \cup R} \rightarrow M_{T \cup R} \subseteq M_{T' \cup R} \rightarrow f(M_{T' \cup R}) \cap M_{T \cup R} = M_{\overline{\overline{T'}}} \cap M_{T \cup R} = M_{\overline{\overline{T'}} \cup T \cup R} \subseteq f(M_{T \cup R}) = M_{\overline{\overline{T}}} \rightarrow$

$\overline{\overline{T}} \subseteq \overline{\overline{\overline{T'}} \cup T \cup R}$. $(^=5)$: Let $U := Th(m)$. $S \subseteq U = \overline{U}$. $U \vdash T \rightarrow m \models T \rightarrow m \in M_{T \cup R} \rightarrow f(M_{T \cup R}) = \{m\} \rightarrow \overline{\overline{T}} = U = \overline{U}$.
(b) If T, T' are such that $\overline{T \cup R} = \overline{T' \cup R} \rightarrow \overline{\overline{T}} = \overline{\overline{T'}}$, so f is well-defined. (f1): By $T \cup R \subseteq \overline{\overline{T}}$, $f(M_{T \cup R}) \subseteq M_{T \cup R}$. (f2): Let $M_{T' \cup R} \subseteq M_{T \cup R}$, so $\overline{T \cup R} \subseteq \overline{T' \cup R}$. Then $M_{\overline{\overline{T}}} \cap M_{T' \cup R} = f(M_{T \cup R}) \cap M_{T' \cup R} = M_{\overline{\overline{T}} \cup T' \cup R} \subseteq M_{\overline{\overline{T'}}} = f(M_{T' \cup R})$. (f3) Take $m := M(U)$, by $S \subseteq \overline{U}$ $m \in M_R$ and $m \models S$. If $m \in M_{T \cup R}$, $m \models T$, so $U \vdash T$, thus $f(M_{T \cup R}) = M_{\overline{\overline{T}}} = M_U = \{m\}$. (f4) is trivial by definition.
Finally, f represents $^=$: $\overline{\overline{T}} = T^f$: $\phi \in T^f :\leftrightarrow \forall m \in f(M_{T \cup R}).m \models \phi \leftrightarrow \forall m \in M_{\overline{\overline{T}}}.m \models \phi$ $\leftrightarrow \overline{\overline{T}} \vdash \phi \leftrightarrow \phi \in \overline{\overline{T}}$, as $\overline{\overline{T}}$ is classically closed. □
We have as Corollary of the preceding two propositions the following soundness and completeness result for preferential models with a smallest element:

Proposition 3.16 Let $^=$ be a logic for $\mathcal{L}$ defined on some $\mathcal{T} \subseteq \mathcal{P}(\mathcal{L})$, and $R \subseteq S$ $\mathcal{L}$-theories. Then there is a definability preserving cpm $\mathcal{M}$ with smallest element $< m, i >$ with $m \models S$ and all points in M R-models, such that $\overline{\overline{T}} = T^{\mathcal{M}}$
iff
$(^=1)$ $\overline{T \cup R} = \overline{T' \cup R} \rightarrow \overline{\overline{T}} = \overline{\overline{T'}}$,
$(^=2)$ $\overline{\overline{T}}$ is closed under the underlying logic,
$(^=3)$ $T \cup R \subseteq \overline{\overline{T}}$,
$(^=4)$ $\overline{T \cup R} \subseteq \overline{T' \cup R} \rightarrow \overline{\overline{T'}} \subseteq \overline{\overline{\overline{T}} \cup T' \cup R}$, $(^=5)$ There is a consistent complete theory U with $S \subseteq \overline{U}$ such that $U \vdash T \rightarrow \overline{\overline{T}} = \overline{U}$.
for all $T, T' \in \mathcal{T}$.

Proposition 3.17 The following axioms are sound and complete for flat conditionals (where $\Rightarrow$ expresses the counterfactual conditional):
(m1) $\forall m(m \models \phi \rightarrow ((\phi \Rightarrow \psi) \leftrightarrow \psi))$,
(m2) $\forall m(m \models \psi \rightarrow \psi') \rightarrow \forall m(m \models (\phi \Rightarrow \psi) \rightarrow (\phi \Rightarrow \psi'))$,
(m3) $\forall m(m \models (\phi \Rightarrow \psi) \wedge (\phi \Rightarrow \psi') \rightarrow (\phi \Rightarrow \psi \wedge \psi'))$,
(m4) $\forall m(m \models \phi \leftrightarrow \phi') \rightarrow \forall m(m \models (\phi \Rightarrow \psi) \rightarrow (\phi' \Rightarrow \psi))$,
(m5) $\forall m(m \models \phi \Rightarrow \phi)$,
(m6) $\forall m(m \models (\phi' \wedge \phi \Rightarrow \psi) \rightarrow (\phi' \Rightarrow (\phi \rightarrow \psi)))$
where m ranges over all points in a fixed counterfactual conditional structure.
(with, of course, the usual rules for $\models$: $m \models \phi \wedge \psi$ iff $m \models \phi$ and $m \models \psi$, likewise for $\neg$ and $\vee$.)

Remark: (1) (m1) corresponds to the smallest element, (m2) and (m3) are classical closure (by compactness of classical logic), (m4) is Left Logical Equivalence, (m5) (with (m2)) is strengthening, (m6) is the Deduction Theorem.
(2) Note that for ϕ, ψ classical, we can deduce: $\forall m(m \models \phi) \leftrightarrow \forall \psi \forall m(m \models \psi \Rightarrow \phi) \leftrightarrow \forall m(m \models \neg\phi \Rightarrow \bot)$ (Proof: $\forall m(m \models \phi) \rightarrow$ for any ψ $\forall m(m \models \psi \rightarrow \phi)$, so by $\forall m(m \models$

$\psi \Rightarrow \psi$) and (m2) $\forall m(m \models \psi \Rightarrow \phi)$. Let $m \models \neg\phi$, then $m \not\models \mathit{true} \Rightarrow \phi$. And: $\forall m(m \models \phi)$ $\rightarrow \forall m(m \models \neg\phi \rightarrow \bot)$, so by $\forall m(m \models \neg\phi \Rightarrow \neg\phi)$ and (m2) $\forall m(m \models \neg\phi \Rightarrow \bot)$. Let $m \models \neg\phi$, then $m \not\models \neg\phi \Rightarrow \bot$.)

Proof We use compactness of classical logic repeatedly. Soundness is trivial.
For completeness: Set $R := \{\phi : \forall m(m \models \neg\phi \Rightarrow \bot)\}$, $\mathcal{T} := \{\phi : \phi \in \mathcal{L}\}$, $\overset{=}{\phi} := \overset{=}{\phi}^{m} := \{\psi : m \models \phi \Rightarrow \psi\}$. We show that this is a special case of Proposition 3.15.
$(^{=}1)$: $\overline{\phi \cup R} = \overline{\phi' \cup R}$, $m \models \phi \Rightarrow \psi$. We have to show $m \models \phi' \Rightarrow \psi$. By $\phi' \in \overline{\phi \cup R}$, there is $\rho \in R$ with $\phi \wedge \rho \vdash \phi'$, likewise there is $\rho' \in R$ with $\phi' \wedge \rho' \vdash \phi$, and as for all m' $m' \models R$, so for all m' $m' \models \phi \leftrightarrow \phi'$, so we are done by (m4).
$(^{=}2)$ is trivial by (m2) and (m3).
$(^{=}3)$ $m \models \phi \Rightarrow \phi$ by (m5). As $m' \models \rho$ for any $\rho \in R$, and m', $m' \models \phi \rightarrow \rho$, so $m \models \phi \Rightarrow \rho$ by (m2).
$(^{=}4)$ Let $\overline{\phi \cup R} \subseteq \overline{\phi' \cup R}$, and $m \models \phi' \Rightarrow \psi'$, we have to show that there is ψ such that $m \models \phi \Rightarrow \psi$ and $\rho' \in R$ with $\vdash \rho' \wedge \phi' \wedge \psi \rightarrow \psi'$. Let by prerequisite $\phi' \wedge \rho' \vdash \phi$. As $\forall m'.m' \models R, \forall m'(m' \models \phi' \leftrightarrow \phi' \wedge \rho' \wedge \phi)$, so $m \models \phi' \wedge \rho' \wedge \phi \Rightarrow \psi'$, so $m \models \phi \Rightarrow (\rho' \wedge \phi' \rightarrow \psi')$ by (m6), and $\rho' \wedge \phi' \rightarrow \psi'$ is our ψ.
$(^{=}5)$: Let $S := U := Th(m)$. Let $U \vdash \phi$, i.e. $m \models \phi$, so $m \models \phi \Rightarrow \psi$ iff $m \models \psi$ iff $U \vdash \psi$.
□
Thus, each $^{=m}$ can be represented by a function f_m on the *same* set of models M_R. Our first representation result tells us now that we can represent each f_m by a different relation $\prec_m$ with smallest element m. If one f_m necessitates more copies than another f_m, we can always make them "invisible" by a straightforward technique - too many copies do not bother. Thus, a consistent $\Rightarrow$-theory has a model $< M, \{\prec_m : m \in M\} >$ with $M \subseteq M_R$ - more precisely, M consists of copies of R-models - and we can turn to the last part of the proof, which is easy: We define an $\mathcal{N}$-system on the structure $< M, \{\prec_m : m \in M\} >$ constructed so far by consistency of the ∇-free consequences.

The rest of the completeness proof: We now show that every consistent theory of the full language has a model.
Any formula ϕ of Level 2 can be written as a disjunction of conjunctions of formulas of the type $\forall m\psi(m)$, $\exists m\psi(m)$, $\nabla m\psi(m)$, $\clubsuit m\psi(m)$, ψ of Level 1, and if a theory T is consistent, we can choose for each $\phi \in T$ one of such conjunctions, preserving consistency. So assume without loss of generality that T is a consistent set of such conjunctions, or, equivalently, we have a set A of formulas of type $\forall m\psi(m)$ or $\exists m\psi(m)$, a set B of formulas of type $\nabla m\psi(m)$, a set C of formulas of type $\clubsuit m\psi(m)$. Take now the set of consequences of T which do not contain the new quantifier. This set is consistent, and, by prerequisite, it has a model. We take this structure and define a $\mathcal{N}$-system on its underlying set M.
If $B = \emptyset$, set $\mathcal{N} := \{M\}$. Otherwise:
By the above Lemma 3.6, for each $\nabla m\psi(m) \in B$, we can conclude that there is some m with $\psi(m)$, as $\exists m\psi(m)$ is a consequence thereof, likewise for pairs in B, pairs in $B + C$,

singletons in C. Choose then for each $\nabla m\psi(m) \in B$ the set $X_{\nabla m\psi(m)}$ by putting into it all the choices for $\nabla m\psi(m)$, $\nabla m\psi(m)$ and $\nabla m\psi'(m)$ etc. Trivially, all such $X_{\nabla m\psi(m)}$ will have pairwise non-empty intersection. Thus, $\{V \subseteq M : V$ contains some such $X_{\nabla m\psi(m)}\}$ is a $\mathcal{N}$-system over M. It remains to show that T holds in the full structure. For each $\nabla m\psi(m)$, this is trivial, as $X_{\nabla m\psi(m)} \in \mathcal{N}$, and the points where $\psi(m)$ holds is a superset thereof. For each $\clubsuit m\psi(m)$ likewise, as we need just one m in each $V \in \mathcal{N}$. $\Box$

Theorem 3.18 The axioms given for ∇ and $\Rightarrow$, together with enough of FOL are sound and complete for our semantics.

Proof: Let $T \not\models \phi$. Then there is a model M, such that $M \models T \wedge \neg\phi$. Thus, $Con(T \wedge \neg\phi)$, so $T \not\vdash \phi$. The other direction is analogous. $\Box$

4 Discussion of the KLM axioms in our framework

(For a discussion and motivation of these axioms, see [KLM90], [LM92], and [Mak94].)
We read $\alpha \mid\sim \beta$ as follows: α is evidence and perhaps background theory, β is a hypothesis, $\alpha \mid\sim \beta$ says that evidence and background theory α confirm hypothesis β.
The KLM axioms Right Weakening, Reflexivity, and Left Logical Equivalence will always hold in an approach that is based on a semantics whose base units consist of models of the underlying logic.
But already AND fails: Evidently, the same evidence can support many hypotheses. This is why it is sometimes necessary to carry out an "experimentum crucis" deciding between hypotheses. But evidence will not support a contradiction, the conjunction of the two. This seems to be a deep reason. We should not hope that some logic generates for us from evidence - how feeble it might be, just one element examined - the ideal hypothesis. Logic should have something to do with truths that hold in all possible worlds, and a logic of confirmation should speak about hypotheses that can be reasonably upheld in the light of evidence, so it speaks about speculations about the nature of the universe at hand, which can be erroneous. The logic of confirmation should permit diverging hypotheses to be supported by the same evidence - but it should not support nonsense (a contradiction)!
Cautious Monotony fails for the same reasons: Evidence might support two contradictory hypotheses, but adding one to the evidence should not permit to deduce the other!
More formally, work in the extended FOL. Fix a structure with its universe U, $\mathcal{N}$-system, and the orders $\prec_a$, $a \in U$, take some observations and look at the hypotheses which are confirmed by evidence and perhaps a background theory. Consider Example 1: $U := \{a, b\}$ - the orders are trivial here - $\mathcal{N}(U) := \{U\}$, and consider a unary predicate p with $p(a)$ as evidence. Both $h := \forall x p(x)$ and $h' := \forall x(p(x) \leftrightarrow x = a)$ are confirmed by the evidence, but of course not their conjunction. Likewise, adding h to the evidence will disconfirm h', showing failure of Cautious Monotony.

OR fails in the presence of tolerated counterexamples. Let $e_i = e^+ \cup e_i^-$ be evidence supporting hypothesis h, with fixed positive evidence e^+, and varying negative evidence. Taking a (non-exclusive) or on the left hand side will add up all negative evidence - and in any finite reading one will arrive at the conclusion that there might be too many counterexamples.

More formally again, consider Example 2: $U := \{a, b, c\}$, $\mathcal{N}(U) := \{X \subseteq U : card(X) \geq 2\}$ with a closest neighbour of both b and c. Assume a positive evidence for ϕ, b negative. Then $c \models o_\phi \Rightarrow \phi$, so ϕ is confirmed. If we take c instead of b as negative evidence, then $b \models o_\phi \Rightarrow \phi$, so ϕ is confirmed again. But taking a as positive, and both b and c as negative evidence will not confirm ϕ any more. So OR does not hold.

Rational Monotony: As, in general, we can derive contradictory hypotheses, the axiom of rational monotony does not make much sense.

4.1 Acknowledgements

Peter Flach, Tilburg, gave me through several articles and a talk given at a DRUMS meeting in fall 1994 at Malaga, Spain, the impression that some intuitive notions about induction should be clarified before trying a formatization. Both are attempted here.

An (anonymous) referee incited me to make this paper more readable.

References

[Gab85] D.M.Gabbay, "Theoretical foundations for non-monotonic reasoning in expert systems". In: K.R.Apt (ed.), "Logics and Models of Concurrent Systems", Springer, Berlin, 1985, p.439-457

[Goo55] N.Goodman, "Fact, Fiction, and Forecast", Harvard Univ. Press, 1955

[KLM90] S.Kraus, D.Lehmann, M.Magidor, "Nonmonotonic reasoning, preferential models and cumulative logics", Artificial Intelligence, 44 (1-2), p.167-207, July 1990

[Lew73] D.Lewis, "Counterfactuals", Harvard, Cambridge, Mass., USA, 1973

[LM92] D.Lehmann, M.Magidor, "What does a conditional knowledge base entail?", Artificial Intelligence, 55(1), p. 1-60, May 1992

[Mak94] D.Makinson, "General patterns in nonmonotonic reasoning", in D.Gabbay, C.Hogger, Robinson (eds.), "Handbook of Logic in Artificial Intelligence and Logic Programming", vol. III: "Nonmonotonic and Uncertain Reasoning", Oxford University Press, 1994, p. 35-110

[Sch89-n1] K.Schlechta: "Defaults as Generalized Quantifiers", Journal of Logic and Computation, Oxford, Vol.5, No.92-27:4, pp.1-22, 1995

[Sch92] K.Schlechta: “Some Results on Classical Preferential Models”, Journal of Logic and Computation, Oxford, Vol.2, No.6 (1992), p. 675-686

[Sch92-n1] K.Schlechta: “Logic, Topology, and Integration”, Journal of Automated Reasoning, Kluwer, 14:353-381, 1995

[Sta68] Stalnaker, “A theory of conditionals”, in N.Rescher (ed.), “Studies in Logical Theory”, Blackwell, Oxford, p. 98-112