Supraclassical Consequence: Abduction, Induction, and Probability for Commonsense Reasoning

Luis M. Augusto*

Editor-in-Chief

Journal of Knowledge Structures & Systems

April 2023

Vol.: 4 Issue: 1 Pages: 1-46

Abstract

Reasoning over our knowledge bases and theories often requires non-deductive inferences, especially – but by no means only – when commonsense reasoning is the case, i.e. when practical agency is called for. This kind of reasoning can be adequately formalized via the notion of supraclassical consequence, a non-deductive consequence tightly associated with default and non-monotonic reasoning and featuring centrally in abductive, inductive, and probabilistic logical systems. In this paper, we analyze core concepts and problems of these systems in the light of supraclassical consequence.

Key words: Supraclassical Consequence; Commonsense Reasoning; Abduction; Induction; Probability

1 Introduction

Although logic is primarily seen as the science of thinking and reasoning it has always been tightly associated with knowledge, providing it among other aspects with adequate media for expression and analysis; to be sure, thinking and reasoning do require knowledge bases – which can actually be theories if closed under logical consequence – over which to operate (e.g., Augusto, 2020c). In the last decades of the 18th century, the celebrated philosopher I. Kant saw no need for any changes in the field of logic, seeing it as "a closed and completed body of doctrine" (Kant, 1787/1929). G. Boole, G. Frege, and a few others would soon prove him wrong. In particular the development of computer science in the second half of the 20th century and related

*✉ luis.ml.augusto@gmail.com

artificial intelligence (AI) and cognitive modeling have brought and continue to bring ever newer challenges that have to do with reasoning in a world characterized by information that is often vague, uncertain, conflicting, noisy, etc. In effect, as early as in 1974, M. Minsky saw logical reasoning as "not flexible enough to serve as a basis for thinking." According to him, the inadequacy of logic to modeling human reasoning was due to the overly strict requirement of consistency, the reliance on small sets of axioms for representing ordinary knowledge, the abstract character of the rules of inference, and the belief that all knowledge is accessible to deduction. In one way or another, all these shortcomings of the "logistic approach" were connected to the classical notion of logical consequence:

> I do not mean to suggest that "thinking" can proceed very far without something like "reasoning." We certainly need (and use) something like syllogistic deduction, but I expect mechanisms for doing such things to emerge in any case from processes for "matching" and "instantiation" required for other functions. Traditional formal logic is a technical tool for discussing either *everything that can be deduced from some data* or *whether a certain consequence can be so deduced*; it cannot discuss at all what ought to be deduced under ordinary circumstances. (Minsky, 1974)

Under "ordinary circumstances," in 1991 Minsky meant the facts that our world is a commonsense (rather than a rational) one and we are satisficers (rather than optimizers). Interestingly enough, Minsky (1991) reiterated the complaint with respect to the strict requirements of consistency and provability, which made logic inappropriate for the modeling of *commonsense reasoning*. This notion of logical consequence criticized by him is deductive, monotonic consequence, in particular the classical one rooting in Aristotelian logic; but this has been greatly "relaxed" or even *ad-hocly* removed for the modeling of both non-deductive and non-classical – i.e. commonsense and heuristic – reasoning, making it even possible to program computers to reason in these ways.

This has been so greatly due to the realization, from within the broadly conceived field of cognitive science, that logic is not merely normative, but also (or perhaps above all) *descriptive*. If we envisage human cognizers as practical reasoners, then we need logics that are capable of describing (instances of) practical agency, and this to a great extent means that we must be prepared to enlarge the properties of interest from the viewpoint of logic to include relevance, analogy, and plausibility, inter alia (see, e.g., Gabbay & Woods, 2003). This *new logic* (Gabbay & Woods, 2001) thus requires modified notions of logical consequence, with the modifications reflecting the properties that one wishes be preserved from a set of premises to a conclusion that is believed to follow from them.

This kind of reasoning replaces the generality of classical deduction by specific (empirical) facts and, most importantly, it typically rejects the rule of substitution, which prescribes that if a formula ϕ is a theorem, then any of its substitution instances $\sigma\phi$ is also a theorem. These aspects make it so that the operation of logical consequence over a knowledge base or a theory often produces more consequences compared to the classical consequence operation; for this reason, this is called *supraclassical consequence*. Default and non-monotonic reasoning are supraclassical in essence; abductive, inductive, and probabilistic inference are the core non-deductive supraclassical inferences.

2 From Classical Deduction to Supraclassicality

2.1 Logical Consequence

We begin with generalities with respect to an unspecified consequence operation $\bigstar$.[1] We do this for a propositional language L, and remark that the notions below generalize to L^* in a straightforward way provided that they are not defined with reference to propositional variables.[2] Structurality (cf. Def. 2.2 below), however, poses some problems that are not without solution (see Wójcicki, 1988, p. 407).

Definition 1. We here take the view of a *logical system* L as a pair $\left(\mathsf{L}^{(*)}, \bigstar\right)$ where $\mathsf{L}^{(*)}$ is an object language and $\bigstar \in \mathsf{L}^{\mu^\mu}$ is a consequence operation.

Definition 2. Let $\bigstar$ be a consequence operation on a set of formulae F, $F \subseteq \mathsf{L}$, $\Gamma \subseteq F$.

1. $\bigstar(\Gamma)$ denotes the set of all the consequences of Γ, i.e. the set of all the formulae ϕ such that $\phi \in \bigstar(\Gamma)$.

2. If furthermore for all substitutions σ it is the case that $\sigma\bigstar(\Gamma) \subseteq \bigstar(\sigma\Gamma)$, then $\bigstar$ is said to be *structural*.

3. $\bigstar$ is *finitary* if it is the case that

$$\bigstar(\Gamma) = \bigcup\left\{\bigstar\left(\Gamma'\right) | \Gamma' \text{ is a finite subset of } \Gamma\right\}.$$

   Otherwise, it is *infinitary.*

4. $\bigstar$ is said to be *standard* if it is both finitary and structural.

5. The strongest consequence operation on L is the operation $\bigstar$ such that $\bigstar(\Gamma) = F = \mathsf{L}$. It is called the *inconsistent* or *trivial consequence operation* on F.

6. The weakest consequence operation on L is the operation $\bigstar$ defined by $\bigstar(\Gamma) = \Gamma$. This is called the *idle consequence operation* on F.

[1]See Augusto (2020a) for a development; see Wójcicki (1988) for a comprehensive elaboration on logical consequence from the viewpoint of mathematical logic. We assume familiarity with the jargon of formal logic, in particular of formal semantics, and with its folklore.

[2]We shall consider a logical language whose formulae are built in the usual way (see Augusto, 2019, for details). This is essentially standard propositional logic with the usual operators collected in set $O = \{\neg, \wedge, \vee, \rightarrow, \leftrightarrow\}$ and a set $V_p = \{p, q, r, s, t, ...\}$ of propositional variables. We denote this language by L and (objects of) its metalanguage will be easily identifiable by means of the superscript μ. Sets of formulae of this language are denoted by F_{L}, arbitrary sets of formulae in the metalanguage L^μ are denoted by the Greek upper-case letters Γ and Δ, and arbitrary formulae in the metalanguage are lower-case letters with or without subscripts at the end of the Greek alphabet, i.e. $F_{\mathsf{L}}^\mu = \{\phi, \varphi, \chi, \psi, \omega, \phi_1, ...\}$. We denote the standard extension to the first-order predicate language (FOL) by L^* and we shall often write $\mathsf{L}^{(*)}$ to denote indifferently the propositional language or its first-order extension. We shall actually make a negligible use of $\mathsf{L}^{(*)}$, remaining mostly at the metalanguage level, with the odd incursion into its meta-metalanguage L^{μ^μ}.

2.2 Inference and Deduction

Presenting it as a rule-based system is a fairly common way of specifying a logical consequence operation $\bigstar$. In turn, this provides us with the definition of an inference system.

Definition 3. Given a language $\mathsf{L}^{(*)}$, a consequence operation $\bigstar$ on $F_{\mathsf{L}^{(*)}}$ is called an *inference operation* on $\mathsf{L}^{(*)}$ if it satisfies $\bigstar 1$ and $\bigstar 2$ for every $\Gamma \subseteq F \subseteq \mathsf{L}^{(*)}$:

$$(\bigstar 1) \quad \Gamma \subseteq \bigstar(\Gamma) \qquad \textit{Inclusion}$$
$$(\bigstar 2) \quad \bigstar(\bigstar(\Gamma)) = \bigstar(\Gamma) \qquad \textit{Idempotency}$$

An *inference system* on $\mathsf{L}^{(*)}$ is a pair $\left(\mathsf{L}^{(*)}, \bigstar\right)$ where $\bigstar$ is an inference operation on $\mathsf{L}^{(*)}$.

Definition 4. Over a language $\mathsf{L}^{(*)}$:

1. an *inference* is a couple (Γ, ϕ) such that $\Gamma \subseteq F \subseteq \mathsf{L}^{(*)}$ and $\phi \in F$. $\phi \in \bigstar(\Gamma)$ just is another notation for representing an inference. Inferences of the form $(\emptyset, \phi)$ are said to be *axiomatic* and we have it that $\phi \in AX$, where AX denotes the set of *axioms*.

2. an *inference rule* $\mathbf{r}_i \in RI$ on a set of formulae F is a mapping assigning to some sequences $\phi_1, ..., \phi_n \in \Gamma$ of formulae (the premises) a formula ψ (the conclusion), i.e. $\mathbf{r} : \Gamma \longrightarrow F$ where $\Gamma \subseteq F^n$ for some $n = 1, 2, ...$ We write $\mathbf{r}(\phi_1, ..., \phi_n) = \psi$ or, more frequently, in the *argument* form $\phi_1, ..., \phi_n/\psi$.

3. a rule $\mathbf{r}$ is said to *preserve* a set of formulae Γ, and Γ is said to *be closed under* $\mathbf{r}$ iff for all Γ' and for all formulae ϕ, if $\Gamma' \subseteq \Gamma$ and $\mathbf{r}(\Gamma', \phi)$, then $\phi \in \Gamma$.

4. a subset $F_0 \subseteq F$ is said to be *closed under a rule of inference* $\mathbf{r}$ provided that $(\phi_1, ..., \phi_n) \in F_0^n \cap \Gamma$ implies that $\mathbf{r}(\phi_1, ..., \phi_n) \in F_0$.

5. a formula ϕ is *provable from* a set Γ by means of axioms in the set AX and/or rules in the set RI iff there is a finite sequence of formulae $\phi_1, ..., \phi_n$ that is a *proof* or a *derivation of* ϕ *from* Γ, i.e. there is a finite sequence $\phi_1, ..., \phi_n$ such that

   (a) $\phi_1 \in (\Gamma \cup AX \cup RI)$;

   (b) for every $1 < i \leq n$, either $\phi_i \in (\Gamma \cup AX \cup RI)$ or ϕ_i is the conclusion of one of the rules of inference $\mathbf{r}_j$, $j = 1, ..., k$ of which the premises are some of the $\phi_1, ..., \phi_{i-1}$;

   (c) $\phi_n = \phi$.

An arbitrary $\mathbf{r}_i$ is a *rule of* $\bigstar$ iff for all Γ, ϕ we have it that $\mathbf{r}(\Gamma, \phi)$ entails $\phi \in \bigstar(\Gamma)$. It is thus obvious that each operation $\bigstar$ is uniquely determined by its rules of inference. There are, however, two rules of inference that are very general and ubiquitous, to wit, *modus ponens*

$$(\mathrm{MP}) \qquad \frac{\phi, \phi \rightarrow \psi}{\psi}$$

and the *rule of substitution*

$$\text{(SUB)} \qquad \frac{\phi}{\sigma\phi}.$$

SUB asserts that if a formula ϕ is a theorem, then any of its substitution instances $\sigma\phi$ is also a theorem.

Definition 5. Two consequence operations can be ordered according to their *strength* (denoted by $\preccurlyeq$): we say that $\bigstar_1 \preccurlyeq \bigstar_2$ ($\bigstar_2$ is stronger than $\bigstar_1$) iff, for all Γ, $\bigstar_1(\Gamma) \subseteq \bigstar_2(\Gamma)$. In fact, the following conditions are equivalent:

1. $\bigstar_1 \preccurlyeq \bigstar_2$
2. $\Theta_{\bigstar_2} \subseteq \Theta_{\bigstar_1}$
3. $\Vdash_{\bigstar_1} \subseteq \Vdash_{\bigstar_2}$

Definition 6. If conditions 1-3 of Def. 5 hold, then we say that the logical system with the consequence operation $\bigstar_2$ is an *extension* of the logical system with the consequence operation $\bigstar_1$.

Definition 7. We say that $\bigstar(\Gamma)$ is a *system of* $\bigstar$. $\bigstar(\emptyset)$ is the system of all logically provable or valid sentences of $\bigstar$. We call this system *a logic* and denote it generally by **L**.

Theorem 8. *A logical system is* structurally complete *iff each structural rule that preserves* $\bigstar(\emptyset)$ *is a rule of* $\bigstar$.

A deductive system is a system of $\bigstar$. In other words, a deductive system $\bigstar(\Gamma)$ is the least theory of $\bigstar$ containing Γ. More formally:

Definition 9. A *deductive system* S over a language $\mathsf{L}^{(*)}$ is the pair

$$\mathrm{S} = \left(\mathsf{L}^{(*)}, (AX \cup RI)\right).$$

Additionally, there can be a distinction between *rules of derivation* (RId) and *rules of theoremhood* ($RI\vartheta$), that is, inference rules that can be applied to hypotheses and those that can be applied only to generate theorems, respectively. We can then define a deductive system over $\mathsf{L}^{(*)}$ as the triple $\left(\mathsf{L}^{(*)}, RId, RI\vartheta\right)$, where it is often (but by no means necessarily) the case that $RId \cap RI\vartheta = \emptyset$. This distinction might come in handy when considering the axiomatization of a logical system.

Definition 10. Two sets of formulae Γ, Δ are *(logically* or *inferentially) equivalent (under* $\bigstar$*)*, denoted by $\Gamma \equiv_{(\bigstar)} \Delta$, iff $\bigstar(\Gamma) = \bigstar(\Delta)$. Two formulae ϕ, ψ are equivalent iff their unit sets are equivalent.

2.3 Classical Consequence

2.3.1 The Classical Consequence Operation

Classical consequence was firstly elaborated on in Tarski (1930). We provide its core properties.

Definition 11. Let L be a propositional language and let $\Gamma, \Delta \subseteq F$. An operation $\bigstar \in O^{\mu^{\mu}}$ defined on Γ, Δ is said to be a *Tarskian, or classical consequence operation* if it is a mapping $Cn : 2^F \longrightarrow 2^F$ satisfying the following conditions:

(C1)	$\Gamma \subseteq Cn(\Gamma)$	*Inclusion*
(C2)	$Cn(Cn(\Gamma)) = Cn(\Gamma)$	*Idempotency*
(C3)	If $\Gamma \subseteq \Delta$, then $Cn(\Gamma) \subseteq Cn(\Delta)$	*Monotonicity*

C2 is often replaced by C2*:

(C2*) $\quad Cn(Cn(\Gamma)) \subseteq Cn(\Gamma) \quad$ *Closure*

It can be easily verified that conditions C1-C3 amount to, for every $\Gamma, \Delta \in F$:

$$\text{(C0)} \qquad \Gamma \subseteq Cn(Cn(\Gamma)) \subseteq Cn(\Gamma) \subseteq Cn(\Gamma \cup \Delta)$$

In mathematical terms, the consequence operation $Cn \in O^{\mu}$ is a closure operation defined on the power set of the set of formulae of a propositional language. Compare conditions C1-3 with conditions c1-3 in the following definition:

Definition 12. Let A be a set. A mapping $\mathsf{c} : 2^A \longrightarrow 2^A$ is called a *closure operator* or *operation on A* if the following properties are satisfied for all subsets $X, Y \subseteq A$:

(c1)	$X \subseteq \mathsf{c}(X)$	*Extensivity*
(c2)	$\mathsf{c}(X) = \mathsf{c}(\mathsf{c}(X))$	*Idempotency*
(c3)	If $X \subseteq Y$, then $\mathsf{c}(X) \subseteq \mathsf{c}(Y)$	*Monotonicity*

A subset $\mathsf{c}(X) \subseteq A$ is said to be *closed* (with respect to the operator c) and it is called the *closed set generated by X*.

Definition 13. Let F be a set of formulae and *Cn* a unary operation defined on the set 2^F:

1. If C0 is satisfied by all $\Gamma, \Delta \in F$, then *Cn* is a *closure operation on F*.

2. $\Gamma \in \Theta_{Cn}$ iff $\Gamma = Cn(\Gamma)$, i.e. Γ is a theory of *Cn* iff it is closed under *Cn*.

3. For each set Cq of consequences on L, the consequence

$$Cn_{Cq}(\Gamma) = \bigcap \{Cn(\Gamma) \mid Cn \in Cq\}$$

   is the greatest lower bound of *Cq*.

4. Let $\mathscr{F}$ be a family of sets. For each consequence operation *Cn*, Θ_{Cn} is a closure system, i.e. for each $\mathscr{F} \subseteq \Theta_{Cn}$, $\cap\mathscr{F} \in \Theta_{Cn}$, with $Cn(\emptyset)$ and F as the least and greatest elements thereof. The former is called the *base theory* and the latter the *trivial theory* of *Cn*.

5. Let $\mathscr{F} = 2^F$. Then, $\mathscr{F}$ is a *closure base* for a consequence operation *Cn* iff for each Γ,

$$Cn(\Gamma) = \bigcap \{\Delta \in \mathscr{F} \mid \Gamma \subseteq \Delta\}.$$

   Note that Θ_{Cn} is the greatest closure base for *Cn*.

6. If Θ_{Cn} is a closure system, then the pair $(\mathscr{F}, \subseteq)$ is a *complete lattice* s.t. for every $\mathscr{F} \subseteq \Theta_{Cn}$,
   (a) $inf(\mathscr{F}) = \bigcap \mathscr{F}$,
   (b) $sup(\mathscr{F}) = inf\{\Gamma \in \Theta_{Cn} \mid \bigcup \mathscr{F} \subseteq \Gamma\}$.

With respect to Def. 13 above, we remark that any set of formulae closed under Cn is a *logical theory*, henceforth abbreviated as *theory* and denoted here by Θ. In contrast with this, we shall see a *knowledge base* as a set Δ of propositions about the world assumed to be *true*.

We next present some of the key results of Tarski (1930). These have to do with consistency, completeness, and axiomatizability of logical systems.

Definition 14. *Consistency:* The set $\Gamma \subseteq F_{\mathsf{L}(*)}$ is consistent relative to Cn iff $Cn(\Gamma) \neq F_{\mathsf{L}(*)}$; otherwise, Γ is said to be inconsistent.

This is equivalent to the classical definition of consistency according to which a set Γ of formulae is consistent if no formula together with its negation both belong to $Cn(\Gamma)$.

Definition 15. *Completeness:* A set of formulae Γ is complete if every consistent set $\Pi \supseteq \Gamma$ satisfies the condition $Cn(\Gamma) = Cn(\Pi)$.

Again, this is equivalent to the classical definition of completeness: a theory Θ of classical logic is complete if, for every formula ϕ, we have it that either $\phi \in Cn(\Theta)$ or $\neg\phi \in Cn(\Theta)$.

Definition 16. *Axiomatizability:* A set of sentences Γ is said to be finitely axiomatizable under Cn iff there is a finite set Γ' that is equivalent to Γ with respect to Cn. Tarski's criterion runs as follows: A system Γ of a finitary consequence operation Cn is *not* finitely axiomatizable iff there exists a strictly increasing sequence of systems of Cn

$$\Gamma_0 \subset \Gamma_1 \subset ... \subset \Gamma_n \subset ...$$

such that

$$\Gamma = \bigcup_{n \in \omega} \Gamma_n$$

where ω denotes the set of all finite ordinals, which has cardinality $\aleph_0$.

2.3.2 The Classical Consequence Relation

Given the natural association between operations and relations, with $\bigstar$ there is associated a consequence relation that can be said to be Tarskian or classical.[3] In effect, $\bigstar$ induces a *consequence relation*

$$(R_{\bigstar}) \qquad (\Gamma, \phi) \in R \text{ iff } \phi \in \bigstar(\Gamma)$$

[3]Given an operation O of degree n on a set A there is a naturally associated relation R of degree $n+1$ on A. For instance, to the operation of addition on the set $\mathbb{R}$ of the real numbers for which $x + y = z$ consisting of the pair $((x, y), z)$ there is associated the relation consisting of the triple (x, y, z), for $x, y, z \in \mathbb{R}$.

where $(\Gamma, \phi) \in R$ denotes that ϕ is a consequence of Γ, and R induces the c*onsequence operator*

$$(\bigstar_R) \qquad \bigstar(\Gamma) = \{\phi \in F_{\mathsf{L}^{(*)}} \mid (\Gamma, \phi) \in R\}.$$

This association allows us to reformulate any of the results for the consequence operation in terms of the consequence relation, and vice-versa.

Definition 17. A *Tarskian, or classical consequence relation* is a relation $\mathtt{R} \subseteq 2^F \times F$ satisfying the following conditions:

(R1) If $\phi \in \Gamma$, then$(\Gamma, \phi) \in \mathtt{R}$.
(R2) If $(\Gamma, \phi) \in \mathtt{R}$ and $\Gamma \subseteq \Delta$, then $(\Delta, \phi) \in \mathtt{R}$.
(R3) If $(\Gamma, \phi) \in \mathtt{R}$ and $(\Delta, \psi) \in \mathtt{R}$ for every $\psi \in \Gamma$, then $(\Delta, \phi) \in \mathtt{R}$.

In truth, both the above $\bigstar$ and $\mathtt{R}$ are elements of a meta-metalanguage; in the metalanguage we shall represent the former as Cn, as is usual, and the latter as $\vdash$ or $\models$, depending on whether we mean $\mathtt{R}$ in a proof-theoretical or in a model-theoretical way, respectively. When we mean no emphasis on such distinction, we shall denote a general (though not *generalized*!; see below) consequence relation by the symbol $\Vdash$. This symbol may actually also stand for adequateness, i.e. the coincidence of $\vdash$ and $\models$, or for some hybridization.

Example 18. Although *semantic tableaux* (cf. Example 29) is basically a *proof* system for propositional and predicate logics, its character is largely *hybrid* in the sense that a tableaux construction is in fact a countermodel with respect to some semantics (e.g., Gabbay, 2014). By *countermodel* it is understood that the proof procedure is so with respect to the negation of the formula we wish to prove. In informal terms, a countermodel corresponds here to a closed tree $\mathfrak{T}$ whose branches are partial descriptions of the model; a branch is said to close if both literals L and $\neg L$ are in it (i.e. if they are nodes of the same branch), and the tree itself is said to close if all its branches close. This is thus a *refutation* proof procedure: if the tableaux $\mathfrak{T}$ for $\neg\phi$ closes, the falsity of $\neg\phi$ is considered refuted and the truth of ϕ is proven. This said, this calculus can also be used to test directly for satisfiability, especially in the case of finite sets of formulae: an open tableau is proof of satisfiability. This proof system was invented by Beth (1955) and it was greatly simplified by Smullyan (1968) into the variant known as *analytic tableaux*.

Moving from the meta-metalanguage to the metalanguage, we have the following definition of the consequence relation:

Definition 19. From $R_\bigstar$ and $\bigstar_R$, we have:

$$(\Vdash_{def}) \qquad \Gamma \Vdash \phi \quad \text{iff} \quad \phi \in Cn(\Gamma)$$

This allows for the redefinition of a logical system as follows:

Definition 20. Given the natural association between the consequence operation Cn and the consequence relation $\Vdash$, a *logical system* L can now be seen as a pair $\left(\mathsf{L}^{(*)}, \Vdash\right)$.

We next present the properties that are shared by both the syntactical and the semantical consequence relations and that additionally hold for most logical systems. As is usual, we abbreviate $\emptyset \Vdash \phi$ as $\Vdash \phi$, $\Gamma \Vdash \{\phi\}$ as $\Gamma \Vdash \phi$, and $\Delta \cup \{\phi\} \Vdash \{\psi\}$ as $\Delta, \phi \Vdash \psi$.

Definition 21. The following properties define the classical consequence relation $\Vdash$:

($\Vdash$ 1) If $\phi \in \Gamma$, then $\Gamma \Vdash \phi$
($\Vdash$ 2) If $\Gamma \Vdash \phi$ and $\Gamma \subseteq \Delta$, then $\Delta \Vdash \phi$
($\Vdash$ 3) If $\Gamma \Vdash \phi$ and $\Delta \Vdash \psi$ for every $\psi \in \Gamma$, then $\Delta \Vdash \phi$
($\Vdash$ 4) $\phi \Vdash \phi$
($\Vdash$ 5) If $\Vdash \phi$, then $\Gamma \Vdash \phi$
($\Vdash$ 6) If $\Gamma \Vdash \phi$ and $\phi \Vdash \psi$, then $\Gamma \Vdash \psi$
($\Vdash$ 7) If $\Gamma \Vdash \phi$ and $\Delta, \phi \Vdash \psi$, then $\Gamma \cup \Delta \Vdash \psi$
($\Vdash$ 8) If $\Gamma \cup \{\phi_1, ..., \phi_n\} \Vdash \psi$ and $\Gamma \Vdash \phi_i$ for $i = 1, ..., n$, then $\Gamma \Vdash \psi$

Property $\Vdash$ 1 just is the definition of the classical consequence relation and we now rewrite R1-R3 and $\Vdash$ 1-$\Vdash$ 4 as follows:

Definition 22. The consequence relation $\Vdash$ is classical if it has the following properties of *reflexivity* (R), *monotonicity* (M), and *transitivity* (T):

(R) $\phi \Vdash \phi$
(M) If $\Gamma \Vdash \phi$, then $\Gamma, \Delta \Vdash \phi$
(T) If $\Gamma \Vdash \phi$ and $\Gamma, \phi \Vdash \psi$, then $\Gamma \Vdash \psi$

Definition 23. A *generalized consequence relation* (GCR) on a set of formulae $F \subseteq \mathsf{L}^{(*)}$ is a consequence relation $\Vdash \subseteq 2^F \times 2^F$ satisfying the following three conditions:

(R) $\phi \Vdash \phi$
(M*) If $\Gamma \Vdash \Delta$, then $\Gamma, \Gamma' \Vdash \Delta, \Delta'$
(T*) If $\Gamma, \phi \Vdash \Delta$ and $\Gamma \Vdash \phi, \Delta$, then $\Gamma \Vdash \Delta$

2.4 Classical Deduction, Compactness, and Structurality

Deduction is a fundamental aspect of a logical system if we are interested in preserving truth (vs. falsity). Thus, we define it semantically, even if we use the symbol $\Vdash$; of course, we assume adequateness of the logical systems in consideration.

Definition 24. A consequence relation $\Vdash$ is *deductive* if it is truth-preserving, i.e. if in the valid argument Γ/ϕ the premises Γ are true, then the conclusion ϕ is necessarily true.

This defines essentially *deductive reasoning*, namely in contrast to abductive and inductive reasoning forms. This definition is conveniently made in semantical terms – *truth* and, to a lesser extent, *truth-preservation* are somehow more intuitive notions than *proof* –, but both the deduction and the deduction-detachment theorems (see below) make the bridge between this and the syntactical component of deduction.

Classical logic is the standard logical system realizing deduction, and we speak of deduction as *classical deduction* when in the context of deduction as realized in classical logic. As it is, classical deduction is adequately expressed via the deduction theorem and, even more strongly, the deduction-detachment theorem. We next discuss these theorems, namely from the viewpoint of the so-called *deductive systems*.

A consequence relation $\Vdash$ is said to satisfy the *deduction theorem* when the following condition is verified:[4]

[4] We leave the proofs of theorems and propositions as exercises for the reader.

Theorem 25. *For Γ a (possibly empty) set of formulae and formulae ϕ, ψ we have it that*

$$\text{(DT)} \qquad \text{If } \Gamma, \phi \Vdash \psi, \text{ then } \Gamma \Vdash \phi \to \psi.$$

DT just means that the formula on the right can be obtained, by successive applications of rules in a proof system, from axioms, Γ, and/or derived rules (syntactical version), or that $\phi \to \psi$ is a tautology (semantical version).

Importantly, DT should be distinguished from the *deduction-detachment theorem.*

Theorem 26. *For Γ a (possibly empty) set of formulae and formulae ϕ, ψ we have it that*

$$\text{(DDT)} \qquad \Gamma, \phi \Vdash \psi \text{ iff } \Gamma \Vdash \phi \to \psi.$$

Obviously, DDT is satisfied iff both DT and MP are satisfied.

We now make more precise the earlier definition of deductive system in Def. 9:

Definition 27. A *deductive system* is a logical system in which the consequence operation satisfies C1-C3 and, in addition, satisfies *compactness*, i.e. for a set Γ and a finite subset Γ' thereof:

$$Cn(\Gamma) \subseteq \bigcup \{Cn(\Gamma') \mid \Gamma' \subseteq \Gamma\}$$

That is, a deductive system is a logical system in which the consequences of a set Γ can be obtained from a finite subset of Γ.

In terms of the corresponding consequence relation, then $\Vdash$ defines a deductive system if, in addition to satisfying conditions R1-R3, it also satisfies, for some finite $\Gamma' \subseteq \Gamma$,

$$\text{If } \Gamma \Vdash \phi, \text{ then } \Gamma' \Vdash \phi.$$

Armed with this definition of a deductive system, we can now define the important notion of a *deductive set.*

Definition 28. A set of formulae $F \subseteq \mathsf{L}$ is deductive iff the following conditions are satisfied:

(D1) F is closed under substitution.
(D2) For all $\phi \in F$, $(\phi \to \phi) \in F$.
(D3) F is closed under the following rules of inference:

1. *Rearrangement of the antecedent* (RA): $\phi \to p / \phi' \to p$, where ϕ, ϕ' are any formulae such that $Vp(\phi) = Vp(\phi')$ and the only connective appearing in both ϕ and ϕ' is $\wedge$.[5]
2. *Enlargement of the antecedent*: $p \to q / (p \wedge r) \to q$.
3. *Composition*: $p_1 \to q_1, p_2 \to q_2 / (p_1 \wedge p_2) \to (q_1 \wedge q_2)$.
4. *Transitivity*: $p \to q, q \to r / p \to r$.

[5] RA is a class of sequential rules rather than a single rule.

5. *MP.*

6. *Cancellation of a valid conjunct*: $p, (p \wedge q) \rightarrow r / q \rightarrow r$.

Structurality, or *closure under substitution* (cf. D1 and Def. 2.2 above), is a particularly important property in that it means that the inference patterns in a given logical system L are invariant with respect to substitutions in the language L. This property actually distinguishes between two major classes of logical systems, to wit, *classical* vs. *supraclassical* systems, and in particular *deductive* vs. *non-deductive* systems. More generically, though, it distinguishes between purely formal logic, which considers only *logical form* regardless of any facts, and logics that take individual (empirical) *facts* into consideration. These latter logics allow for a formal modeling of commonsense reasoning, as they are largely free from the strict constraints imposed by (classical) deduction. Besides structurality, the reader will see that also D3 is essentially not satisfied in the logics we shall focus on in Section 3 below.[6]

2.5 Two Classical Deduction Systems

For the sake of self-containment, we elaborate briefly on two classical deductive systems, also spoken of as *proof systems* or *calculi*, that will be required below in Section 3, to wit, Gentzen-style systems and the semantic/analytic tableaux system. These systems are structurally complete, i.e. they preserve all the classical tautologies. We introduce them in the order they will be required. With respect to the latter mentioned system, we remain at the propositional level; the reader is referred to Augusto (2019) for the first-order predicate extension.

Example 29. In a semantic/analytic tableau, one has proven that formulae $\phi_1, ..., \phi_n$ logically entail ψ – and thus ψ is a logical consequence of $\phi_1, ..., \phi_n$ – when it is the case that $\{\phi_1, ..., \phi_n, \neg\psi\}$ is unsatisfiable. The rules to be applied are the *expansion* and *closure rules* and they can best be expressed in the language of set theory as follows:

$$(\wedge_{\mathrm{RE}}) \quad \frac{\Gamma \cup \{\phi \wedge \psi\}}{\Gamma \cup \{\phi, \psi\}}$$

$$(\vee_{\mathrm{RE}}) \quad \frac{\Gamma \cup \{\phi \vee \psi\}}{\Gamma \cup \{\phi\} \mid \Gamma \cup \{\psi\}}$$

$$(\mathsf{X}) \quad \frac{\Gamma \cup \{L, \neg L\}}{\mathsf{X}}$$

The expansion rules can be conveniently reduced to two; this is known as the $\alpha\beta$ classification, summarized in the following tables (where evidently χ_1, χ_2 are subformulae of χ):

α	α_1	α_2
$\phi \wedge \psi$	ϕ	ψ
$\neg(\phi \vee \psi)$	$\neg\phi$	$\neg\psi$
$\neg(\phi \rightarrow \psi)$	ϕ	$\neg\psi$
$\neg\neg\phi$	ϕ	ϕ

β	β_1	β_2
$\neg(\phi \wedge \psi)$	$\neg\phi$	$\neg\psi$
$\phi \vee \psi$	ϕ	ψ
$\phi \rightarrow \psi$	$\neg\phi$	ψ

[6] D2 poses problems to relevance logic(s).

There are thus in effect only two rules, A (for α) and B (for β):

$$\text{(A)} \qquad \frac{\alpha}{\begin{array}{c}\alpha_1\\ \alpha_2\end{array}}$$

$$\text{(B)} \qquad \frac{\beta}{\beta_1 \mid \beta_2}$$

Intuitively, these two rules mean that if α is a conjunction of α_1 and α_2, then both these two subformulae are logical consequences of α (rule A) and thus are nodes in the one and same branch of the tree $\mathfrak{T}$ (i.e., analytic tableau); if β is a disjunction of β_1 and β_2, then the latter subformulae originate two different branches as a logical consequence of β (rule B). Figure 1 shows an analytic tableau proof. Note the numbering of steps on the left and the application of the rules on the right of the steps. Closure of a branch is indicated by X.

1. $\neg((p \to q) \to (\neg(q \wedge r) \to \neg(r \wedge p)))$
2. $p \to q$ $\quad \alpha(1)$
3. $\neg(\neg(q \wedge r) \to \neg(r \wedge p))$ $\quad \alpha(1)$
4. $\neg(q \wedge r)$ $\quad \alpha(3)$
5. $\neg(\neg(r \wedge p))$ $\quad \alpha(3)$
6. $r \wedge p$ $\quad \alpha(5)$
7. r $\quad \alpha(6)$
8. p $\quad \alpha(6)$

9. $\neg q$ $\quad \beta(4)$ 10. $\neg r$ $\quad \beta(4)$

X (under 10)

11. $\neg p$ $\quad \beta(2)$ 12. q $\quad \beta(2)$

X X

Figure 1: Theoremhood of $(p \to q) \to (\neg(q \wedge r) \to \neg(r \wedge p))$: An analytic tableau proof.

Example 30. Frege systems and extended Frege systems are subsumed under the notion of *Hilbert(-style) systems*. These systems are *axiom systems*, i.e. their set of rules of inference is often a singleton. Frege and Hilbert systems are distinct from *Gentzen(-style) systems* in that the latter rely on no, or very few, axioms. These proof systems are characterized by the fact that a proof is a sequence of *sequents* (subsets of formulae rather than formulae), and each of the sequents is derivable from earlier sequents in the sequence by means of rules of inference. A sequent $\mathbf{s}$ for subsets $\Gamma = \{\phi_1, ..., \phi_m\}$ and $\Delta = \{\psi_1, ..., \psi_n\}$, m and n are non-negative integers, is of the form

$$(\mathbf{s}) \qquad \Gamma \Rightarrow \Delta$$

where the symbol $\Rightarrow$ denotes "Δ follows from Γ", and Γ (the *antecedent*) is equivalent to $\phi_1 \wedge ... \wedge \phi_m$ and Δ (the *succedent* or *consequent*) is equivalent to $\psi_1 \vee ... \vee \psi_n$. We focus here on Gentzen's (1934-5) *logical calculus* $\mathcal{LK}$.[7] We have the important fact of $\mathcal{LK}$: A sequent $(\phi_1, ..., \phi_m) \Rightarrow (\psi_1, ..., \psi_n)$ is provable in $\mathcal{LK}$ iff the sequent $(\phi_1 \wedge ... \wedge \phi_m) \Rightarrow (\psi_1 \vee ... \vee \psi_n)$ is provable in it. This means that s is equivalent to:

$$\text{(DER)} \qquad \Rightarrow (\phi_1 \wedge ... \wedge \phi_m) \rightarrow (\psi_1 \vee ... \vee \psi_n)$$

This is in fact the definition of *derivability*, or *proof-theoretical validity* for conclusions that are actually sets of formulae. An expression of the form "$\Rightarrow \phi$" denotes that ϕ is a theorem. $\mathcal{LK}$ has a single axiom known as *axiom of identity*:

$$\text{(Ax)} \qquad \frac{}{\phi \Rightarrow \phi}$$

In what follows, symbols occur as in $F^{\mu}_{\mathsf{L}^{(*)}}$ with the following (novel) denotations: Γ, $\triangle$, Σ, and Π denote *contexts*, i.e. finite, possibly empty, sequences of formulae. As usually, $\phi(t)$ denotes a formula ϕ where the occurrence of t is of interest, and $\phi^x_{t,a}$ denotes the formula obtained by replacing specified occurrences of x in $\phi(x)$ by a term t or a parameter a. The rules of inference are either *logical* or *structural*, and each comes in left and right versions. The former are basically rules for the use of the connectives (but also of the quantifiers), reason why they are also known as *operational rules*, whereas the latter rules impact on a logical system in ways that will become clear below. We begin with the logical rules.

Logical rules of $\mathcal{LK}$:

Left rules		Right rules	
$(\wedge L_1)$	$\frac{\Gamma,\phi \Rightarrow \Delta}{\Gamma,\phi\wedge\psi \Rightarrow \Delta}$	$(\wedge R)$	$\frac{\Gamma \Rightarrow \phi,\Delta \quad \Sigma \Rightarrow \psi,\Pi}{\Gamma,\Sigma \Rightarrow \phi\wedge\psi,\Delta,\Pi}$
$(\wedge L_2)$	$\frac{\Gamma,\psi \Rightarrow \Delta}{\Gamma,\phi\wedge\psi \Rightarrow \Delta}$		
$(\vee L)$	$\frac{\Gamma,\phi \Rightarrow \Delta \quad \Sigma,\psi \Rightarrow \Pi}{\Gamma,\Sigma,\phi\vee\psi \Rightarrow \Delta,\Pi}$	$(\vee R_1)$	$\frac{\Gamma \Rightarrow \phi,\Delta}{\Gamma \Rightarrow \phi\vee\psi,\Delta}$
		$(\vee R_2)$	$\frac{\Gamma \Rightarrow \psi,\Delta}{\Gamma \Rightarrow \phi\vee\psi,\Delta}$
$(\rightarrow L)$	$\frac{\Gamma \Rightarrow \phi,\Delta \quad \Sigma,\psi \Rightarrow \Pi}{\Gamma,\Sigma,\phi\rightarrow\psi \Rightarrow \Delta,\Pi}$	$(\rightarrow R)$	$\frac{\Gamma,\phi \Rightarrow \psi,\Delta}{\Gamma \Rightarrow \phi\rightarrow\psi,\Delta}$
$(\neg L)$	$\frac{\Gamma \Rightarrow \phi,\Delta}{\Gamma,\neg\phi \Rightarrow \Delta}$	$(\neg R)$	$\frac{\Gamma,\phi \Rightarrow \Delta}{\Gamma \Rightarrow \neg\phi,\Delta}$
$(\forall L)$	$\frac{\Gamma,\phi^x_t \Rightarrow \Delta}{\Gamma,\forall x\phi \Rightarrow \Delta}$	$(\forall R)$	$\frac{\Gamma \Rightarrow \phi^x_a,\Delta}{\Gamma \Rightarrow \forall x\phi,\Delta}$
$(\exists L)$	$\frac{\Gamma,\phi^x_a \Rightarrow \Delta}{\Gamma,\exists x\phi \Rightarrow \Delta}$	$(\exists R)$	$\frac{\Gamma \Rightarrow \phi^x_t,\Delta}{\Gamma \Rightarrow \exists x\phi,\Delta}$

In the above rules $\forall R$ and $\exists L$, the *restrictions* are that the parameter a must not be a free variable in Γ and Δ, or it must not appear anywhere in the respective lower sequents. We now present the structural rules; these are rules for *weakening* (W), *contraction* (C), and *permutation* (P).

[7] This is often presented with formulae and sets of formulae of the object language. We chose to remain in the metalanguage for $\mathsf{L}^{(*)}$.

Structural rules of $\mathcal{LK}$:

Left rules		Right rules	
(WL)	$\dfrac{\Gamma\Rightarrow\Delta}{\Gamma,\phi\Rightarrow\Delta}$	(WR)	$\dfrac{\Gamma\Rightarrow\Delta}{\Gamma\Rightarrow\phi,\Delta}$
(CL)	$\dfrac{\Gamma,\phi,\phi\Rightarrow\Delta}{\Gamma,\phi\Rightarrow\Delta}$	(CR)	$\dfrac{\Gamma\Rightarrow\phi,\phi,\Delta}{\Gamma\Rightarrow\phi,\Delta}$
(PL)	$\dfrac{\Gamma_1,\phi,\psi,\Gamma_2\Rightarrow\Delta}{\Gamma_1,\psi,\phi,\Gamma_2\Rightarrow\Delta}$	(PR)	$\dfrac{\Gamma\Rightarrow\Delta_1,\phi,\psi,\Delta_2}{\Gamma\Rightarrow\Delta_1,\psi,\phi,\Delta_2}$

There is a supplementary rule that is of fundamental interest to $\mathcal{LK}$– and to logic in general–, to wit, the *cut rule*:

$$\text{(CUT)} \qquad \frac{\Gamma \Rightarrow \Delta, \phi \quad \phi, \Sigma \Rightarrow \Pi}{\Gamma, \Sigma \Rightarrow \Delta, \Pi}$$

Fig. 2 shows a $\mathcal{LK}$ proof of $\forall x\,(A\,(x) \rightarrow B) \rightarrow \exists x\,(A\,(x) \rightarrow B)$. We make $t = a$. The proof does not use CUT.

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{A\,(a) \Rightarrow A\,(a)\,, B}\,\text{Ax}}{\Rightarrow A\,(a)\,, A\,(a) \rightarrow B}\,(\rightarrow\text{R})}{\Rightarrow A\,(a)\,, \exists x A\,(x) \rightarrow B}\,(\exists\text{R}) \qquad \dfrac{\dfrac{\dfrac{}{A\,(a)\,, B \Rightarrow B}\,\text{Ax}}{B \Rightarrow A\,(a) \rightarrow B}\,(\rightarrow\text{R})}{B \Rightarrow \exists x A\,(x) \rightarrow B}\,(\exists\text{R})}{\dfrac{A\,(a) \rightarrow B \Rightarrow \exists x A\,(x) \rightarrow B}{\forall x A\,(x) \rightarrow B \Rightarrow \exists x A\,(x) \rightarrow B}\,(\forall\text{L})}\,(\rightarrow\text{L})$$

$$\dfrac{\forall x A\,(x) \rightarrow B \Rightarrow \exists x A\,(x) \rightarrow B}{\Rightarrow \forall x\,(A\,(x) \rightarrow B) \rightarrow \exists x\,(A\,(x) \rightarrow B)}\,(\rightarrow\text{R})$$

Figure 2: A $\mathcal{LK}$ proof of $\Rightarrow \forall x\,(A\,(x) \rightarrow B) \rightarrow \exists x\,(A\,(x) \rightarrow B)$.

2.6 Non-Deductive Reasoning: Non-Monotonic and Default Reasoning

A deductive system is a logical system, but not all logical systems are deductive systems. The kind of inference that we are about to approach, roughly known as *defeasible reasoning*, can be seen as either deductive or non-deductive; it can be deductive because at play is often truth-preservation, and defeasible reasoning has a mostly corrective character, but it can be non-truth-preserving in that the conclusion may hold for most typical cases when the premises hold, but not for the case at hand. In the latter case, it has an ampliative character. In any case, truth-preservation and deductive closure are central preoccupations of defeasible reasoning. Also, non-monotonicity, on which this kind of reasoning is fundamentally based, can actually have a monotonic, deductive basis, and we thus decide for the essentially deductive character of this kind of reasoning.

Crucially, non-monotonic consequence is classical consequence, as non-monotonic consequence relations do not typically reject, but actually include, elements from the latter. Indeed:

Definition 31. *Non-monotonic consequence* is an extension of the classical consequence, a property known as *supraclassicality*, expressed formally as

$$\text{(SCL)} \qquad \text{If } \Gamma \Vdash \phi, \text{ then } \Gamma \mid\!\!\mid\!\!\sim \phi$$

where $\mid\!\!\mid\!\!\sim$ denotes an unspecified non-monotonic consequence relation.[8]

Recall condition $\Vdash 2$ for the classical consequence relation:

$$(\Vdash 2) \qquad \text{If } \Gamma \Vdash \phi \text{ and } \Gamma \subseteq \Delta, \text{ then } \Delta \Vdash \phi$$

Recall also that this is related to condition M:

$$\text{(M)} \qquad \text{If } \Gamma \Vdash \phi, \text{ then } \Gamma, \Delta \Vdash \phi$$

Condition M basically states that adding new information to a knowledge base does not change it. However, early on in the history of AI and cognitive modeling (e.g., McCarthy, 1959; Minsky, 1961, 1974), it became evident that the monotonicity condition for the classical consequence relation made classical logic inappropriate to modeling human reasoning when dealing with everyday problems. In effect, we more often than not reason with partial information about a situation, and in order to deduce a class of immediate consequences wide enough to meet this partiality we have to make use of *commonsense.* So as to carry out commonsense reasoning a wide set of assumptions about "normalcy" in the world is required. But this set of assumptions is not static: it is actually not only being made on the going, but it is also being updated or revised in a principled way. Thus, the addition of new data may invalidate some or all the assumptions that were made at an earlier stage in the reasoning process.

Example 32. The "classical" example is that of Tweety the penguin: when first hearing that Tweety is a bird, we infer that Tweety flies, as this is in agreement with our knowledge base about birds (i.e. we know that most birds fly); however, we retract this conclusion once we learn that Tweety is a penguin, because we know that these birds do not fly.

This retraction is considered the main feature of *defeasible inference*, and *non-monotonicity* is at the core of the consequence relation stipulated for a correct formalization of this kind of inference.

We obtain a *non-monotonic logic* by restricting condition $\Vdash 2$ in the following way:

Definition 33. *Restricted monotonicity (RM)* is defined via the condition:

$$\text{If } \Gamma \mid\!\!\mid\!\!\sim \phi \text{ and } \Gamma \mid\!\!\mid\!\!\sim \psi \text{ , then } \Gamma, \phi \mid\!\!\mid\!\!\sim \psi$$

Note that this is also known as *cautious monotonicity.* The meaning of this condition is as follows: if ϕ and ψ are expected to be true by Γ (the "if ... and" in the condition above), then ψ is expected to be true if ϕ is assumed to be true.

[8]Specifications are made below in this Subsection and in Section 3.

Example 34. A classical example of failure of RM is given in the following knowledge base:

$$\texttt{(1) Quaker}(x)\ \Vert\!\sim \texttt{Pacifist}(x)$$

$$\texttt{(2) Republican}(x)\ \not\Vert\!\sim \texttt{Pacifist}(x)$$

Now, suppose x is simultaneously a Quaker and a Republican; we may choose to revise (1) given (2), and we conclude that:

$$\texttt{(3) Quaker}(x)\texttt{, Republican}(x)\ \not\Vert\!\sim \texttt{Pacifist}(x)$$

The applications of defeasible logic thus suggest themselves, with medical diagnosis and legal reasoning high on the list. In particular, defeasible reasoning is central to the field known as *belief revision*, in which it is of fundamental importance to determine what remains the same in a constantly changing world. This, known as the *frame problem* (McCarthy & Hayes, 1969), together with the fact that we often make inferences based solely on absence of information to the contrary, motivated the elaboration of formalisms with *new* rules of inference capable of expressing non-monotonic reasoning. Moreover, these rules were expected to be capable of handling *beliefs*, rather than merely formulae, which somehow alters the meaning of truth-preservation: rather than this, what is sought is a formalism for *consistency* between all the facts about the world (expressible in first-order formulae) and what we believe to be true about the world that is in fact sanctioned by such rules. Default rules, at the core of *default logic* (Reiter, 1980), are such rules.

Definition 35. A *default (rule)* **d** is a rule of the form

$$(\mathbf{d}) \qquad \phi : \Gamma / \psi$$

expressing that if ϕ (the *pre-requisite*) is believed and each $\chi \in \Gamma$ can be consistently believed (the *justification*), then ψ (the *consequent*) should be believed.

Example 36. An obvious example of the application of a default rule is

$$\text{Bird}(x) : \texttt{m\_}\text{Fly}(x)\ /\ \text{Fly}(x)$$

where "m_" is to be read as "it is consistent to assume." If we do not have information that Tweety is a pinguin, then it is consistent to assume that Tweety can fly, and we therefore can conclude that Tweety does fly:

$$\text{Bird(tweety)} : \texttt{m\_}\text{Fly(tweety)}\ /\ \text{Fly(tweety)}$$

Note also that, given **d**, we only have to represent explicitly positive information about the world, which considerably reduces an explicit knowledge base.

Definition 37. This is known as *closed world assumption.*

In effect, given any n-ary relation R and $x_1, ..., x_n$ individuals, we can assume that $\neg R(x_1, ..., x_n)$ whenever it is consistent to do so. This is sanctioned by a *closed world default*:

$$: \mathtt{m\_}\neg R(x_1, ..., x_n) / \neg R(x_1, ..., x_n) .$$

As is obvious, the *consistency requirement* is at the heart of a formalism for default reasoning.

Definition 38. A *default theory* is a pair (RD, WLD) where RD is a set of default rules and WLD is a set of closed first-order formulae representing what is known to be true of the world.

The crucial point in a default theory is that it is assumed that WLD is incomplete, and the default rules *extend* the theory by acting as mappings from the incomplete theory to a more complete theory, that is, an *extension*.

Definition 39. (Reiter, 1980) For any set of sentences $F \subseteq \mathsf{L}^*$ of a closed default theory $\Lambda = (RD, WLD)$, $\triangledown(F)$ is defined as the smallest set satisfying

1. $WLD \subseteq \triangledown(F)$.
2. $\triangledown(F)$ is closed under $\Vdash_{\mathrm{K}}$, i.e. $Cn(\triangledown(F)) = \triangledown(F)$.
3. If $(\phi : \mathtt{m\_}\chi_1, ..., \mathtt{m\_}\chi_m / \psi) \in RD$, and $\phi \in \triangledown(F)$ and $\neg\chi_1, ..., \neg\chi_m \notin F$, then $\psi \in \triangledown(F)$.

Definition 40. We then say that a set E is an *extension* of the default theory Λ iff E is a fixed point of $\triangledown$, i.e. $\triangledown(E) = E$.

Some default theories have no extension, which makes it necessary to look for restricted default theories for which extensions can be proved. In any case, extensions are in fact the sets that contain the non-monotonic consequences, and we thus have provided an example of how the failure of the monotonicity condition M for the consequence relation can motivate a non-monotonic logic.

Besides supraclassicality and restricted monotonicity, as well as cut and consistency, there are other desirable properties for a non-monotonic consequence relation. In fact, note that any non-monotonic logic is expected to satisfy reflexivity, cut, and monotonicity, i.e. it should be endowed with a paraclassical consequence relation.

Definition 41. A consequence relation is said to be *paraclassical* if it satisfies the following conditions:

(R˜)	If $\phi \in \mathit{\Gamma}$, then $\mathit{\Gamma} \Vert\!\sim \phi$
(CUT˜)	If $\mathit{\Gamma} \Vert\!\sim \phi$ and $\mathit{\Gamma}, \phi \Vert\!\sim \psi$, then $\mathit{\Gamma} \Vert\!\sim \psi$
(M˜)	If $\mathit{\Gamma} \Vert\!\sim \phi$ and $\mathit{\Gamma} \subseteq \mathit{\Delta}$, then $\mathit{\Delta} \Vert\!\sim \phi$

Moreover, and reinforcing our decision to consider non-monotonic consequence relations as classical deductive, we have it that:

Proposition 42. *A consequence relation $\Vert\!\sim$ that is supraclassical, monotonic˜, and closed under substitution is identical to the classical deductive consequence relation, i.e.* $\Vert\!\sim\, = \,\Vdash$.

Definition 43. In order to be able to add formulae to the set of premises without for that changing the inference result we require a property that combines CUT$^{\sim}$ and RM; this is the property known as *cumulativity*:

$$\text{If } \Gamma \Vert\!\sim \phi, \text{ then } \Gamma \Vert\!\sim \psi \text{ iff } \Gamma, \phi \Vert\!\sim \psi.$$

Definition 44. Besides supraclassicality, the following properties characterize the relation between the classical consequence relation $\Vdash$ and the non-monotonic consequence relation $\Vert\!\sim$:

1. *Subclassical cumulativity*: If $\Delta \subseteq \Gamma$ and $\Delta \Vdash \phi$ for every $\phi \in \Gamma$ and $\Delta \Vert\!\sim \psi$, then $\Gamma \Vert\!\sim \psi$.
2. *Consistency preservation*: If $\Gamma \Vert\!\sim \phi$, then $\Delta \Vdash \phi$.
3. *Weak transitivity*: If $\Gamma \Vert\!\sim \Delta \Vdash \phi$, then $\Gamma \Vert\!\sim \phi$.
4. *Distribution*: If $\Gamma \Vert\!\sim \phi$ and $\Delta \Vert\!\sim \phi$, then $\Theta_\Gamma \cap \Theta_\Delta \Vert\!\sim \phi$.
5. *Left logical equivalence* (LLeq$^{\sim}$): If $\Theta_\Gamma = \Theta_\Delta$, then $\Gamma \Vert\!\sim \phi$ iff $\Delta \Vert\!\sim \phi$.
6. *Right weakening* (RW$^{\sim}$): If $\Gamma \Vert\!\sim \phi$ and $\phi \Vdash \psi$, then $\Gamma \Vert\!\sim \psi$.
7. *Left absorption*: $\Gamma \Vert\!\sim \psi$ iff $\Delta \Vdash \psi$, when $\Gamma \Vert\!\sim \phi$ for every $\phi \in \Delta$.
8. *Right absorption*: $\Gamma \Vert\!\sim \psi$ iff $\Delta \Vert\!\sim \psi$, when $\Gamma \Vdash \phi$ for every $\phi \in \Delta$.
9. *Adjunction* : If $\Gamma \Vert\!\sim \phi$ and $\Gamma \Vert\!\sim \psi$, then $\Gamma \Vert\!\sim \phi \wedge \psi$.
10. *Summation* : If $\Gamma, \phi \Vert\!\sim \chi$ and $\Gamma, \psi \Vert\!\sim \chi$, then $\Gamma, \phi \vee \psi \Vert\!\sim \chi$.
11. DT$^{\sim}$: If $\Gamma, \phi \Vert\!\sim \psi$, then $\Gamma \Vert\!\sim \phi \rightarrow \psi$.

Combinations of the above produce further relational properties (e.g., left and right absorption combine into *full absorption*). Makinson (2003) provides a comprehensive elaboration on what he calls "bridges between classical and non-monotonic logic" and, incidentally, the reader interested in "going non-monotonic" could profit from Makinson (2005).

3 Supraclassical Consequence and Commonsense Reasoning

As seen above, although we can approach it from a purely proof-theoretical perspective, deduction is intuitively and typically thought of as *truth-preservation.* Non-deductive logical consequence is characterized by inference that is not truth-preserving. But this can be in two ways, at least.

Definition 45. *Non-deductive inference:*

- *v. 1* – A consequence relation $\Vdash$ is not truth-preserving (or non-deductive) if there is a valid argument Γ/ϕ where the premises in Γ are true and the conclusion ϕ may be false.

- *v. 2* – A consequence relation $\Vdash$ is not truth-preserving (or non-deductive) when the relation between the premises and the conclusion is not one of truth-preservation, but rather of plausibility, strength, or probability.

Our interest falls on version 2. This can be the case when, for instance, ϕ appears to be a plausible explanation for the premises in Γ, or when Γ appears to account for the conclusion ϕ in some degree of strength or probability. This type of inference is in fact what Minsky saw as commonsense reasoning (see Introduction): for instance, the simple observation that the lawn is wet together with the knowledge that (i.e. that we know of) no one has watered it motivates the non-deductive inference that it must have rained as the best explanation. But the applications of such logical consequence relations or operations are vast and surpass the realm of everyday life; as a matter of fact, this kind of non-deductive reasoning has proved itself essential for such diverse functions as natural language processing, program verification and debugging, medical diagnosis, artificial vision, etc., all of which can be computerized. For example, programs for medical diagnosis are often based on plausible explanations for the symptoms shown by patients.

Commonsense reasoning requires the above kind of logical consequence because, firstly, reality is such that it makes deductive inference useless in many circumstances. For instance, patients with oily skin and spots can have acne, but they can – instead or also – have folliculitis; in fact, acne may be considered a "special" case of the latter. On the other hand, patients can have acne but no oily skin, and as a matter of fact, no spots; the less typical acne sufferers may instead have furuncles and cysts.

The fact is that if we had very fine-grained ontologies, we would perhaps be able to do with deductive reasoning most of the times. But we are satisficers – this is the second aspect –,which means that large-grained ontologies work just fine for us. And this is true of scientific ontologies, too: a medical doctor may diagnose a patient as having folliculitis just because the hair follicles are inflamed and/or infected; but this is not of much help by itself, for *this* folliculitis, indeed if it is so, may be of bacterial (several bacteria), fungal (several fungi), etc., origin. On the other hand, testing for all these possible causes would make medical consultation and treatment far more expensive and time-consuming.

We summarize and expand on the above with respect to commonsense, or everyday, inference:

1. Commonsense reasoning tries to make sense of observations, i.e. facts and events, not all of which are fully or even just adequately accessible to the observer (e.g., other people's mental states; causes behind an effect);

2. In order to do so, we need knowledge bases, i.e. sets of propositions about the world;

3. Knowledge bases are prone to a plethora of problems such as inconsistency, vagueness, redundancy, ambiguity, lacunae, etc.;

4. The observations, together with the available knowledge bases, constitute systems of reasoning generally characterized by uncertainty;

5. As satisficers, we want to keep the reasoning process simple, and tend to work with as few premises as possible and to draw one single conclusion at a time.[9]

So, we are fully justified in carrying out a formal account of this kind of reasoning. The question now is: Is this still *consequence*, in the logical sense of the term? Recall that classical logical consequence can be extended by means of the supraclassicality condition SCL (cf. Def. 31).

This actually means that there is a logical consequence operation $Cn^{>}$ s.t.

$$\text{(SPR)} \qquad Cn\,(\Gamma) \subseteq Cn^{>}\,(\Gamma)$$

or, more correctly,

$$\text{(SPR}^{*}\text{)} \qquad Cn\,(\Gamma) \preceq Cn^{>}\,(\Gamma)$$

where $Cn\,(\Gamma)$ denotes the classical consequence operation on a set of formulae Γ, and $Cn^{>}\,(\Gamma)$ is a consequence operation on Γ that generates *more* consequences from Γ than $Cn\,(\Gamma)$.

This might surprise the reader, given that $Cn\,(\Gamma)$ is defined as being a *closure* operation on a set Γ, which entails that no more formulae are derivable or entailable from $Cn\,(\Gamma)$ than those already in $Cn\,(\Gamma)$ (recall condition C2*, or C2, for that matter). How can a consequence operation be stronger (cf. Def. 5) than this? The justification is simple and has to do with SUB: this rule does not hold for $Cn^{>}$. That is to say that a supraclassical consequence is in principle not structural (cf. Def. 2.2).[10] As seen above, this entails non-monotonicity. More specifically, a supraclassical consequence relation may be obtained by adding background assumptions, by restricting the set of possible valuations, and/or by augmenting the set of rules. This enables us to work with default assumptions, default valuations, and default rules.[11]

But supraclassicality is not non-monotonicity simpliciter, as a non-monotonic consequence relation may generate more, as well as fewer, consequences than the classical consequence relation. To begin with, recall that a non-monotonic consequence relation $\Vert\!\sim$ is "at heart" classical; put formally, it is *paraclassical* (cf. Def. 41), a feature that allows for a plethora of relations – *monotonic bridges*, as Makinson (2005) puts it – with classical logic, the first of which is precisely supraclassicality, or the property that $\Vert\!\sim\, \subseteq \left(2^{F} \times F\right)$ is an extension of $\Vdash\, \subseteq \left(2^{F} \times F\right)$ (cf. Def. 31), so that we have

$$\Vdash \;\subseteq\; \Vert\!\sim$$

for a set of formulae $F \subseteq \mathsf{L}$.[12] It is precisely this generation of more consequences by means of the relation $\Vert\!\sim$ (or the operation $Cn^{\sim}$) that we see as supraclassicality. For

[9]As a matter of fact, it is interesting to mention that while this kind of reasoning appears to be effortless and very successful, deductive reasoning is difficult and error-prone, as some empirical investigations into logical reasoning have shown. See, e.g., Stenning & van Lambalgen (2008).

[10]Structurality is not to be identified *tout court* with rule SUB, but the idea to be conveyed is that substitution, whether as an inference rule or as a function, does not hold in principle for supraclassical consequence. In other words, supraclassical consequence is not closed under substitution. In effect, recall from Proposition 42 that a consequence relation $\Vert\!\sim$ (which will feature ubiquitously in this Section) is identical to the classical consequence relation $\Vdash$ if it is both supraclassical and closed under substitution.

[11]See Makinson (2005) for a comprehensive discussion.

[12]We then say that a consequence relation (or operation) is paraclassical iff it is both a closure relation (or operation) and supraclassical.

instance, if we add a set of background assumptions A to a knowledge base Δ, then we will have consequences of the form $A \cup \Delta \Vert\!\sim \phi$, or $\Delta \Vdash_A \phi$, where it might be the case that $\Delta \nVdash \phi$. Also, if we restrict the valuation set $W = \{v_0, v_1, ..., v_{n-1}\}$ as $W' \subseteq W$, we may actually generate more logical consequences by considering individual consequences for each and every valuation $val_i \in W'$ s.t. we have $\Delta \models_{W'} \psi$ where otherwise we might have $\Delta \nvDash \psi$. The same result can be obtained by considering a set of new, additional rules RI^+ s.t. we have $\Delta \vdash_{RI^+} \chi$ where otherwise we might have $\Delta \nvdash \chi$. Let us denote the first case by $Cn^{\sim}_A$, the consequence operation with relation to (or modulo) the set A of background assumptions, and the second case as $Cn^{\sim}_{W'}$, the consequence operation with relation to (or modulo) the set of restricted valuations W'; similarly, we write $Cn^{\sim}_{RI^+}$ to denote the consequence operation with relation to (or modulo) the set of additional rules RI^+. In any case, we have $Cn\,(\Gamma) \subseteq Cn^{\sim}_A\,(\Gamma)$, $Cn\,(\Gamma) \subseteq Cn^{\sim}_{W'}\,(\Gamma)$, and $Cn\,(\Gamma) \subseteq Cn^{\sim}_{RI^+}\,(\Gamma)$. Then, supraclassicality is in fact defined by the inequality:

$$(\text{SPR}^{**}) \qquad Cn\,(\Gamma) \preceq Cn^{\sim}\,(\Gamma)$$

The logical consequences that we approach next are all supraclassical in one way or another, largely so because the systems to be elaborated on have *new* rules, namely non-deductive rules. We elaborate on logical consequence for abductive, inductive, and probabilistic logics. In this process, we shall discuss some of the relations among these non-deductive logics; for a more comprehensive treatment of these relations see, e.g., Flach & Kakas (2000).

3.1 Abductive Consequence

Although abductive reasoning can be considered a kind of defeasible reasoning (Section 3.3), the retraction of inferences is not its main feature. Besides, it is an ampliative kind of reasoning, which means that DT does not *necessarily* hold in an abductive logical system. In particular, DDT does not hold because we are now working with an "inverted" MP rule.

Definition 46. Given an observation ω, we can infer that ϕ (is an explanation for ω) according to the following *abduction inference rule*

$$(\leftarrow \text{R}) \qquad \frac{\omega, \phi \rightarrow \omega}{\phi}$$

where the material implication connective can be interpreted as "is a good explanation for" or "is a cause for."

Then, DT holds iff $\Gamma \Vdash \phi \rightarrow \omega$, but this in turn holds only if $\leftarrow$ R holds, i.e.

$$\Gamma, \phi \Vdash \omega \text{ iff } \Gamma \Vdash \phi \rightarrow \omega \text{ iff } \Gamma, (\omega, \phi \rightarrow \omega)\,/\phi.$$

It should be obvious that we left the terrain of classicality and are now in the domain of supraclassicality, not the least because $\leftarrow$ R is not a rule of classical logic.

This type of reasoning is *ampliative* in the sense that, given a theory Θ (a set of propositions about the world Δ, i.e. a database with facts, laws, etc.), once we select ϕ as a good explanation for an observation ω, we can then incorporate it in our theory (knowledge base), and we end up with an extended theory $\Theta^{>}$ (an enlarged knowledge base $\Delta^{>}$, respectively). Thus, formalizing this kind of reasoning is essential for theory and knowledge base updates, which might entail belief revision and default reasoning, to name but a few applications.

Despite the many decidability problems of abduction for all but the most trivial problems (e.g., Bylander et al., 1991),[13] it has been applied, with different degrees of success, to fields (sometimes combined) such as diagnosis (e.g., Cox & Pietrzykowsky, 1987; de Kleer & Williams, 1987; de Kleer et al., 1992), human and robot perception (e.g., Shanahan, 2005), natural language processing (e.g., Hobbs et al., 1993), and plan recognition (e.g., Ng & Mooney, 1992).

The literature on abduction is vast, and we refer the reader to Gabbay & Woods (2005) for a comprehensive elaboration on this subject from the viewpoint of applications. It is common practice to separate abductive approaches into set-based and logical approaches, but we see no need to do this; on the contrary, we present an abductive system that is a hybrid of both. For a survey of these two approaches, as well as of a third, knowledge-based, approach, see Paul (1993).

Abduction is often, and following its coiner, C. S. Peirce, defined as the inference to the best explanation. In this very general perspective, abduction can be approached from two components, to wit, hypothesis assembly/selection and inference proper. Virtually any logical language will do to formalize abduction, and we shall be using $\mathsf{L}^{(*)}$; we remain at the propositional level for the sake of simplicity. Because hypothesis assembly/selection is a fundamental part of abduction, we need to provide our logical system with an abduction framework.

Definition 47. Given language $\mathsf{L}^{(*)}$,[14] we say that there is an *abduction problem* when, given a theory $\Theta \subseteq F_{\mathsf{L}^{(*)}}$ and an observation $\omega \in F_{\mathsf{L}^{(*)}}$, it is the case that $\Theta \nvDash \omega$ and $\Theta \nvDash \neg\omega$. We denote this problem by the pair $A? = (\Theta, \omega)$.

Note that – especially, but by no means only – in the case of scientific theories, it is often the case that actually there is a set of observations that pose abduction problems. In that case, we establish that $A? = (\Theta, \Omega)$, where Ω is a set of observations.[15]

Definition 48. An *abduction framework* is a pair $\mathcal{A} = \left(A?, \overleftarrow{A}\right)$, where $A?$ is an abduction problem and $\overleftarrow{A} \subseteq Vp_{\mathsf{L}}$ is a set of distinguished symbols (literals) from Vp_{L} called *abducibles*. An arbitrary formula built solely with symbols from $\overleftarrow{A}$ belongs to the set $F_{\overleftarrow{A}}$ of *abducible formulae*.

We identify $F_{\overleftarrow{A}}$ with $\overleftarrow{A}$, for the sake of simplicity. In practical terms, abducibles just are hypotheses.

Example 49. We shall be working with the following set Θ_{lawn}:

[13]Which should become apparent in the discussion below. In particular, the consistency condition makes first-order abductive logic essentially undecidable. Propositional abductive logic is considered intractable.

[14]We need a predicate language only to allow us to work with predicates, even if only 0-ary predicates. Although not necessary, this increases readability.

[15]Recall that theories are sets of formulae closed under logical consequence. However, whatever we say with respect to theories can be loosely extended to knowledge bases, and we will be working with these whenever the notion of closure can be relaxed. Thus, we can denote a knowledge base by Θ when a folk theory is meant.

- water_on_lawn→lawn_is_wet
- no_water_on_lawn→lawn_is_dry
- water_on_lawn→lawn_is_green
- no_water_on_lawn→lawn_is_yellow
- I_didn't_water_lawn
- there_was_no_tsunami
- tap_not_dripping
- is_rainy_season

Our set of abducibles is $\overleftarrow{A}_{wetlawn} = \{r, w, d, t\}$, where r stands for it_rained, w for watered_lawn, d for dripping_tap, t for tsunami. Let also o stand for our observation lawn_is_wet.

Note that we could work directly with the 0-ary predicates of Θ, but this would make the whole formalization more cumbersome. We keep them for readability. As usually, we shall be working mostly at the metalanguage level, and for that we denote our observation by ω and an arbitrary element of $\overleftarrow{A}$ by ϕ, ψ, etc.

Given the observation o that the lawn is wet, any of the elements of $\overleftarrow{A}_{wetlawn}$ is a plausible hypothesis. If on waking up and opening the window one sees that the lawn is wet, any of the hypotheses above is a candidate for an explanation for this observation. One might first infer r, as it is indeed the most plausible explanation, as well as the most parsimonious. However, on remembering that the lawn was watered late in the evening the day before, one should now consider w, even if this is indeed less parsimonious than r; in fact, it could be the case that the lawn had dried overnight. Hypothesis t really is the least plausible and least parsimonious of all, but it could be a true explanation, and thus it cannot be discarded without further consideration.

As we can see, the whole process of inferring a proposition ϕ to explain the observation ω is a *defeasible* process. Moreover, it is a parsimonious process, as it could well have been the case that the lawn had been watered and later on it rained, but one will typically not infer – at least before giving the case some further consideration – that both it rained and one watered the lawn.

It can also be the case that our set of abducibles does not contain a plausible explanation for the observation at hand. We need to look for another set, and in doing so we might end up with inconsistency if we do not eliminate (i.e. retract) our previous set of abducibles. On the other hand, this new set may actually cause the initial theory to change: we might have to reduce it or even enlarge it with the new additions. All this shows that the consequence relation at play in abduction is essentially *non-monotonic.*

All this also means that we have some set-theoretical work to do, namely with respect to Ω and $\overleftarrow{A}$. Recall that we are interested in the explanatory power of any $\phi \in \overleftarrow{A}$ as far as $\omega \in \Omega$ is concerned. In order for this explanatory power to be assured we need a domain for hypothesis assembly.

Definition 50. A *domain for hypothesis assembly* is a triple $\mathcal{H} = \left(\Omega, \overleftarrow{A}, e\right)$ where Ω is the set of observations, $\overleftarrow{A}$ is the set of hypotheses, and e is a function $e : 2^{\overleftarrow{A}} \longrightarrow 2^{\Omega}$ s.t. $e_{\Omega}\left(\overleftarrow{A}\right)$ is the *explanatory power* of the elements of $\overleftarrow{A}$ with respect to the observations in Ω. $e_{\Omega}\left(\overleftarrow{A}\right)$ is subject to the following conditions:

1. For any $\Omega' \subseteq \Omega$ and any $\overleftarrow{A}' \subseteq \overleftarrow{A}$, $e_{\Omega'}\left(\overleftarrow{A}'\right)$ is computable;
2. For $\overleftarrow{A}_1, \overleftarrow{A}_2 \subseteq \overleftarrow{A}$, if $\overleftarrow{A}_1 \subseteq \overleftarrow{A}_2$, then $e_{\Omega}\left(\overleftarrow{A}_1\right) \subseteq e_{\Omega}\left(\overleftarrow{A}_2\right)$;
3. If $\overleftarrow{A}_1, \overleftarrow{A}_2 \subseteq \overleftarrow{A}$, then $e_{\Omega}\left(\left(\overleftarrow{A}_1\right) \cup \left(\overleftarrow{A}_2\right)\right) = e_{\Omega}\left(\overleftarrow{A}_1\right) \cup e_{\Omega}\left(\overleftarrow{A}_2\right)$; and
4. The function $\alpha : \overleftarrow{A} \rightarrow \Omega$ defined for $\overleftarrow{A}' \subseteq \overleftarrow{A}$ and $\phi \in \overleftarrow{A}'$ s.t.

   $$\alpha(\phi) = \left\{\omega \in \Omega \,|\, \omega \in e_{\Omega}\left(\overleftarrow{A}'\right) \text{ and } \omega \notin e_{\Omega}\left(\overleftarrow{A}' - \{\phi\}\right)\right\}$$

   is computable.

More intuitively, according to condition 1, we know exactly which observations are explainable by which hypotheses; 2 is the *monotonicity* condition; 3 is a condition stronger than the previous one and that has to do with the *independence* of hypotheses, which allows us to generate $e_{\Omega}\left(\overleftarrow{A}_2\right) = e_{\Omega}\left(\overleftarrow{A}_1 \cup \{\phi\}\right)$ once we know $e_{\Omega}\left(\overleftarrow{A}_1\right)$; and 4 is the condition of *accountability*, which guarantees that each hypothesis ϕ accounts precisely for $e_{\Omega}(\phi)$ if we have $\alpha(\phi) = e_{\Omega}(\phi)$.

Given these conditions or assumptions, the triple $\mathcal{H} = \left(\Omega, \overleftarrow{A}, e\right)$ should then be subjected to algorithmic procedures to determine the plausibility of the hypotheses (the *screening phase*), the different hypotheses accounting for each $\omega \in \Omega'$ (the *collection phase*), the uniqueness of each $\overleftarrow{A}'$ (the *parsimony phase*), and the essentiality of each $\phi \in \overleftarrow{A}$ (the *critique phase*). Later on we shall see how this equates with conditions on the consequence relation. These algorithms can be found in Allemang et al. (1987).

As far as our abductive problem $A?_{wetlawn}$ goes, by respecting these conditions and phases we should end up with $\overleftarrow{A}'_{wetlawn} = \{r\}$. Note that some of the facts in Θ_{lawn} are irrelevant for the conclusion (e.g., `water_on_lawn`$\rightarrow$`lawn_is_green`), but that just is the nature of knowledge bases about the world, i.e. they often contain irrelevant information for the problem at hand. Moreover, note that the last four facts listed in our knowledge base helped us to eliminate the other elements of $\overleftarrow{A}_{wetlawn}$ and helped us to decide for r.

Definition 51. An *abduction system* is a pair $\mathrm{A} = (\mathcal{A}, \|\!\sim)$, where $\mathcal{A}$ is an abductive framework and $\|\!\sim$ is a non-monotonic consequence relation s.t., for a framework $\mathcal{A} = \left(A?, \overleftarrow{A}\right)$ and an observation $\omega \in \Omega$,

1. $\Theta, \phi \|\!\sim \omega$;

2. $\Theta \cup \phi$ is consistent or satisfiable;

3. $\phi \in \overleftarrow{A}$.

If $\Theta \cup \phi$ is a solution to an abductive problem $A? = (\Theta, \omega)$, we denote it as $A! = ((\Theta, \phi), \omega)$.

Clearly, $A!$ has a corresponding ternary consequence relation (Θ, ϕ, ω). In fact, it just is the consequence relation $\|\!\sim$. Given the above set-based conditions and assumptions, we now know that $\|\!\sim$ is required to satisfy the following conditions:

Definition 52. Given an abduction system $\mathrm{A} = (\mathcal{A}, \|\!\sim)$ as above where the negation of $\|\!\sim$ is denoted by $\not\|\!\sim$, the following *constraints* are imposed on the abductive consequence relation $\|\!\sim$ in $\Theta, \phi \|\!\sim \omega$:

1. *Minimality*: ϕ is a literal, or, if ϕ is a formula, then there is no formula ψ s.t. $Lit(\phi) \subset Lit(\psi)$, where $Lit(\phi)$ is the set of literals of ϕ.

2. *Consistency*: $\Theta \not\|\!\sim \neg\phi$, i.e. $\Theta, \phi \not\|\!\sim \bot$.

3. *Explanatoriness*: $\Theta \not\|\!\sim \omega$ and $\phi \not\|\!\sim \omega$.

Suppose now that we have $A! = \left(\left(\Theta, \overleftarrow{A}'\right), \omega\right)$, i.e. we managed to reduce $\overleftarrow{A}$ to a subset $\overleftarrow{A}'$ that is not a singleton. We can still proceed with an abductive inference, but clearly this requires caution now, as we do not really know which of the $\phi_i \in \overleftarrow{A}'$ is the best explanation – if there is one, that is.

Definition 53. We say that a subset $\overleftarrow{A}' = \{\phi_1 \vee ... \vee \phi_n\}$ is a *cautious explanation* of an observation ω, denoted by $\overleftarrow{A}^{\omega}_{\vee}$, if there is a function $\tilde{e}_\omega : 2^{\overleftarrow{A}} \longrightarrow \omega$ s.t.

1. $\Theta, \overleftarrow{A}^{\omega}_{\vee} \|\!\sim \omega$;

2. $\Theta \cup \overleftarrow{A}^{\omega}_{\vee}$ is consistent or satisfiable;

3. For every abducible formula $\phi \in \overleftarrow{A}^{\omega}_{\vee}$, we have $\Theta, \phi \|\!\sim \omega$ iff $\overleftarrow{A}^{\omega}_{\vee} = \{\phi\}$;

4. If there is no $\phi \in \overleftarrow{A}^{\omega}_{\vee}$ s.t. condition 1 is satisfied, then $\overleftarrow{A}^{\omega}_{\vee} = \bot$.

Condition 4 above tells us that $\tilde{e}_\omega$ is a *selection* function. Note how this notion of cautious explanation means a radical departure from classical deductive inference: we know that in the argument $\Gamma \Vdash \phi$ the set Γ is actually a conjunction of formulae, so that we have $\Gamma \Vdash \phi$ only if it is the case that $(\bigwedge_{i=1}^{n} \psi_i) \Vdash \phi$ for all $\psi_i \in \Gamma$. If we add another set to the premises, say Δ, then we must have it that, for all $\chi_i \in \Delta$:

$$\Gamma, \Delta \Vdash \phi \equiv \left[\left(\bigwedge_{i=1}^{n} \psi_i \in \Gamma\right) \cup \left(\bigwedge_{i=1}^{n} \chi_i \in \Delta\right)\right] \Vdash \phi$$

If we now appeal to DT, then we have:

$$\text{If } \Gamma, \Delta \Vdash \phi, \text{ then } \Gamma \Vdash (\chi_1 \wedge ... \wedge \chi_n) \rightarrow \phi$$

We cannot proceed like this here because we reached the situation where we are working with a set of possible explanations precisely because we could not reduce it any further; but we do not know if their conjunction is the explanation for the observation(s) at hand. For instance, let us suppose that in the evening we had watered the lawn, this had dried overnight, and then early in the morning, before we woke up, it had started to rain; we see that the lawn is wet, and we might think that the lawn is now wet because *both* we watered it *and* it rained; but in fact this conjunction is false as an explanation. The conclusion is: DT fails here. Thus we resort to disjunction as a means of "preserving" rule $\leftarrow$ R; in our particular abduction problem, *either* the fact that we watered the lawn *or* the fact that it rained, *or both*, is a solution, i.e.

$$\frac{o, (r \vee w) \rightarrow o}{r \vee w}.$$

Compare with how preservationism (see Augusto, 2020a, 4.5.2), too, resorts to a similar strategy with the aim of obtaining consistent inferences with inconsistent supersets. In any case, we can work with disjunctions of conjunctions of literals, what is known as a DNF; by the equivalence of this with a CNF, we can actually "trick" abductive inference to respect DT.[16] On the one hand, it may be the case that we are actually working with DNFs. For instance, suppose that, in our example, we wake up to see that there was a minor tsunami and that it is raining; because we also had watered the lawn in the previous evening, we have the disjunction of abducibles $r \vee (w \wedge t)$, which is obviously a DNF. Note, however, how this compromises the minimality condition. On the other hand, this may actually be useful if we are working with an abductive program, as DNFs and CNFs are often required for computerization.

Definition 54. Now let there be given the function $\tilde{e}_{\Omega} : 2^{\overleftarrow{A}} \longrightarrow 2^{\Omega}$. For observations $\omega_i \in \Omega$, $i = 1, 2, ..., n$, an abductive inference is called *cautious* if we have a cautious explanation $\overleftarrow{A}_{\vee}^{\omega_i}$ satisfying the following conditions:

1. If $\Vdash \omega_1 \leftrightarrow \omega_2$, then $\Vdash \left(\overleftarrow{A}_{\vee}^{\omega_1}\right) \leftrightarrow \left(\overleftarrow{A}_{\vee}^{\omega_2}\right)$.
2. $\overleftarrow{A}_{\vee}^{\omega_1 \wedge \omega_2} = \left(\overleftarrow{A}_{\vee}^{\omega_1}\right) \wedge \left(\overleftarrow{A}_{\vee}^{\omega_2}\right)$.

[16] A formula ϕ is said to be in a *disjunctive normal form* (DNF) iff ϕ has the form $\phi = \phi_1 \vee ... \vee \phi_n$, $n \geq 1$, where each of $\phi_1, ..., \phi_n$ is a conjunction of literals, i.e.

$$\phi = \bigvee_{i=1}^{n} \left(\bigwedge_{j=1}^{m} L_{i,j} \right).$$

A formula ϕ is said to be in a *conjunctive normal form* (CNF) iff ϕ has the form $\phi = \phi_1 \wedge ... \wedge \phi_n$, $n \geq 1$, where each of $\phi_1, ..., \phi_n$ is a disjunction of literals, i.e.

$$\phi = \bigwedge_{i=1}^{n} \left(\bigvee_{j=1}^{m} L_{i,j} \right).$$

For ϕ in CNF we have $\neg \bigwedge_{i=1}^{n} \left(\bigvee_{j=1}^{n} L_{i,j} \right) \equiv \bigvee_{i=1}^{n} \left(\bigwedge_{j=1}^{n} \neg L_{i,j} \right)$ and for ϕ in DNF we have $\neg \bigvee_{i=1}^{n} \left(\bigwedge_{j=1}^{n} \neg L_{i,j} \right) \equiv \bigwedge_{i=1}^{n} \left(\bigvee_{j=1}^{n} \neg L_{i,j} \right)$, i.e. every CNF (DNF) formula has an equivalent DNF (CNF, respectively) formula.

3. $\left(\overleftarrow{A}_{\vee}^{\omega_1}\right) \vee \left(\overleftarrow{A}_{\vee}^{\omega_2}\right) \Vdash \overleftarrow{A}_{\vee}^{\omega_1 \vee \omega_2}$.

How can we guarantee that $\overleftarrow{A}_{\vee}^{\omega}$ is indeed the best explanation for ω? Simply by generalizing the conditions above on the abductive consequence relation to sets that work as cautious explanations:

Definition 55. Given an abduction system $\mathrm{A} = (\mathcal{A}, \Vert\!\sim)$ and a selection function $\tilde{e}_{\omega} : 2^{\overleftarrow{A}} \longrightarrow \omega = \overleftarrow{A}_{\vee}^{\omega}$ as above, the following constraints are imposed on the abductive consequence relation $\Vert\!\sim$ in $\Theta, \overleftarrow{A}_{\vee}^{\omega} \Vert\!\sim \omega$:

1. *Minimality*: $\overleftarrow{A}_{\vee}^{\omega}$ is the minimal elementary explanation for ω, i.e.
   (a) there is no $\overleftarrow{B}_{\vee}^{\omega}$ s.t. $\overleftarrow{B}_{\vee}^{\omega} \subseteq \overleftarrow{A}_{\vee}^{\omega}$ and $\overleftarrow{A}_{\vee}^{\omega} \Vdash \overleftarrow{B}_{\vee}^{\omega}$ ($\Theta \cup \overleftarrow{A}_{\vee}^{\omega} \Vdash \overleftarrow{B}_{\vee}^{\omega}$) implies that $\overleftarrow{B}_{\vee}^{\omega} \Vdash \overleftarrow{A}_{\vee}^{\omega}$ ($\Theta \cup \overleftarrow{B}_{\vee}^{\omega} \Vdash \overleftarrow{A}_{\vee}^{\omega}$, respectively), and
   (b) there are no literals $\phi, \psi \in \overleftarrow{A}_{\vee}^{\omega}$ s.t. $\phi = \psi$.
2. *Consistency*: $\Theta \not\Vert\!\sim \neg\left(\overleftarrow{A}_{\vee}^{\omega}\right)$, i.e. $\Theta, \overleftarrow{A}_{\vee}^{\omega} \not\Vert\!\sim \bot$; in particular, for an arbitrary $\phi_i \in \overleftarrow{A}_{\vee}^{\omega}$, $\Theta, \phi_i \not\Vert\!\sim \bot$.
3. *Explanatoriness*: $\Theta \not\Vert\!\sim \omega$ and $\overleftarrow{A}_{\vee}^{\omega} \not\Vert\!\sim \omega$.

Definition 56. Let now the set $ME(\omega)$ be the set of minimal elementary explanations of ω. Then, the *explanation closure* of ω is the disjunction explanation given by

$$\bigvee_{\overleftarrow{A} \in ME(\omega)} \overleftarrow{A} = \overleftarrow{A}_{\vee}^{\omega}.$$

In this way, the cautious explanation of ω, given by the function $\tilde{e}_{\omega}$, assures us that we have the most parsimonious (or the minimal) explanation for our observation. That is to say that among the selected explanations one or more will be the best explanation(s). This is how assured we can be of our inference process. But then that is what disjunctive closure gives us.

Definition 57. The *disjunctive closure* of a set A, denoted by $\overline{A_{\vee}}$, is defined as

$$\overline{A_{\vee}} = A \cup \{(\phi_1 \vee ... \vee \phi_n) \mid \phi_i \in A\}.$$

Above, we made no commitment to whether the abductive consequence relation $\Vert\!\sim$ is defined syntactically or semantically. As seen above (cf. Example 18), one way to provide a logic with a proof system that is at the same time somehow a semantics is by means of semantic or analytic tableaux. In the case of abduction, this might actually be the ideal procedure, as it is not a trivial task to provide an abductive logic with a semantics (or with a proof system, for that matter). Aliseda (2006) has proposed a semantic tableaux proof system for abductive logic. The rules and interpretation are exactly the same as above (cf. Example 29), with the following additional interpretation:

Definition 58. Let $\mathfrak{T}_\Theta$ be a tableau for a theory Θ. A proof of $\Theta \vdash \omega$ is a closed tableau for $\Theta \cup \{\neg\omega\}$, where ω is the fact to be explained.

Suppose now that the tableau $\mathfrak{T}_\Theta$ does not close for ω. Then, abduction is concretized in the tableau by adding new formulae to the open branches until the tableau is eventually closed; in particular, ω may be derivable by means of a minimal extension of Θ, i.e. the set of abducibles (which must be in the vocabulary of Θ). Aliseda (2006) considers then three conditions for a formula ϕ to fulfill this task:

Definition 59. Given a theory Θ, a fact to be explained ω, and an abducible ϕ, ϕ is an explanation of ω if:

1. $\mathfrak{T}_{\Theta\cup\{\neg\omega\}\cup\{\phi\}}$ is closed. We have $\Theta, \phi \mid\sim \omega$. (*Plain abduction*)
2. Plain abduction + $\mathfrak{T}_{\Theta\cup\{\phi\}}$ is open. We have $\Theta \not\mid\sim \phi$. (*Consistent abduction*)
3. Plain abduction + $\mathfrak{T}_{\Theta\cup\{\neg\omega\}}$ is open and $\mathfrak{T}_{\{\phi\}\cup\{\neg\omega\}}$ is open. Respectively, we have $\Theta \not\mid\sim \omega$ and $\phi \not\mid\sim \omega$. (*Explanatory abduction*)

We apply here the syntactical version of the consequence relation $\|\sim$ as defined above, as it is compatible with Aliseda's approach. Obviously, the objective is to find algorithms to provide us with a ϕ that *closes* a $\mid\sim$-consequence between a theory Θ and an observation ω. Aliseda (2006) does provide such algorithms; Aliseda (2007) defends her choice of a tableaux approach to abduction with the claim that in order to "operate a logical system," inference is not enough: a search strategy (a heuristics) is required. According to her, her *Extended Semantic Tableaux* approach fulfills this requirement.

3.2 Inductive Consequence

Suppose that in our theory Θ_{lawn} we have `it_rains`→`lawn_is_wet` as an additional fact (cf. Example 49), representable in terms of propositional variables as $r \to l$. We wake up to see that it is raining, and we infer that the lawn is wet. We appear to be epistemically justified in doing so, as knowledge bases are often composed of repeated observations for which no counterexamples have been observed – though they may well exist, or they are known to exist but the reasoners do not share that knowledge. Also, we appear to infer deductively, as we have an argument of the form

$$\frac{r \to l, r}{l}$$

which is of course the form of MP. But in fact DT does not hold, for, suppose that the lawn was covered with a waterproof canvas; then it would not be wet, despite the fact that it is raining. As a matter of fact, we cannot even apply MP, because verification of empirical data makes rule SUB inapplicable. We are now in the domain of supraclassicality.

Suppose now that, still armed with our theory Θ_{lawn}, on waking up we verify that the lawn is wet, i.e. l holds. We can consult our theory, and by doing so we select r as a plausible explanation for l. We are not only entitled to do so by our theory, we actually save up a lot of time in thinking of other plausible hypotheses. But we are also no longer in the domain of abduction. Recall that this is an ampliative kind of reasoning in the sense that at the end of the reasoning process we have the augmented theory $\Theta^{>}_{lawn}$. If $r \rightarrow l$ is already contained in our theory, then at the end of the reasoning process our theory will be unchanged in cardinality.

But this does not mean that it remains exactly the same theory, because now the plausibility of $r \rightarrow l$ is so to say strengthened. Every time we verify this conditional to be true, we increase its "strength" in terms of plausibility, even though we cannot here speak of truth-preservation. We are in the realm of induction.

Abduction and induction are related in many ways, not the least of which is that induction can be seen as a special case of abduction, in the sense that it is a special form of inference to the best explanation. This is a rather philosophical question (e.g., Harman, 1965), but a formal approach can actually support this view. On the other hand, in the field of inductive logic programming, abduction, together with justification, is seen as a component of induction (e.g., Muggleton & de Raedt, 1994). Also, inductive logic is commonly associated with conditional probability (call this Bayesian approaches), but here we consider this aspect to belong more properly to probabilistic logic (see next Subsection).

The literature on induction is, if not larger, older than that for abduction, as induction has been studied since at least Aristotle first called it *epagoge*.[17] Besides the literature cited below, we also refer the reader to Gabbay et al. (2011) for a comprehensive treatment of this topic.

Induction is typically introduced to the lay reader with an example with ravens.

Example 60. The *raven example* for an instance of inductive logical reasoning:

```
raven 1 is black
raven 2 is black
⋮
raven n is black
----------------
(1) all ravens are black
(2) raven n+1 is black
```

Note that there are two possible conclusions given the same set of premises: (1) is known as *generalization*, while (2) is called *prediction*. As a matter of fact, we can combine both types of conclusion in a single argument (see below). A look at the arguments above shows that induction, too, is a form of ampliative reasoning, but now in the sense that the conclusion contains information that goes beyond the information provided by the premises. In the example above, conclusion 1 states that if a (sufficiently large) sample of As are all found to be Bs, then it is reasonable to infer that *all* As are Bs; in conclusion 2, from a (sufficiently large) sample of observed As as being Bs, we have the prediction that the next observed A will also be a B. What the premises do is: they give us a degree of likelihood, or strength, to our conclusions.

Definition 61. An argument of the form

$$\frac{\phi_1 \wedge ... \wedge \phi_n}{\varphi}$$

[17] Tellingly, he contrasted it with *syllogismos*.

is said to be *inductive* if

1. $\phi_1 = \phi_i = \phi_n \approx \varphi$,
2. $(\bigwedge_{i=1}^{n} \phi_i) \succ_\varphi \left(\bigwedge_{i=1}^{n-1} \phi_i\right) \succ_\varphi ... \succ_\varphi (\phi_1 \wedge \phi_2) \succ_\varphi (\phi_1) \succ_\varphi \emptyset$, where $\succeq_\varphi$ denotes a binary relation (a partial order) of strength with respect to φ, and
3. $\emptyset \nVdash \varphi$.

Condition 1 states that φ is approximately the same (denoted by the symbol $\approx$) as all the basically identical ϕ_i;[18] indeed, a conclusion of type 1 generalizes ϕ universally, and a conclusion of type 2 generalizes ϕ_n to ϕ_{n+1}. Condition 2 states that the higher the number of similar premises ϕ_i, the stronger the inference relation is between the ϕ_i and φ; this is so to the point where there are no tautologies or theorems in this inductive framework, as condition 3 states.

Whereas now most authors would invoke probability to account for conditions 1-3, we leave that for probability logic proper and invoke instead the concept of *limit*.[19]

Definition 62. The *strength* of support (or confirmation) of a set of premises $\Gamma = \{\phi_1, ..., \phi_n\}$ where $\phi_1 = \phi_i = \phi_n$ with respect to a conclusion $\varphi \approx \phi_i$ in an argument of the inductive form above is a function of the limit of the $\phi_i \in \Gamma$, i.e.

$$(\mathrm{LIM}_{\Vdash}^{\infty}) \qquad \left(\lim_{n\to\infty} \bigwedge_{i=1}^{n} \phi_i\right) \cup \{\varphi\} \Vdash \top.$$

An argument satisfying $\mathrm{LIM}_{\Vdash}^{\infty}$ is said to be *inductively strong*. Otherwise, it is said to be *inductively weak*, according to the condition

$$\left(\mathrm{LIM}_{\Vdash}^{0}\right) \qquad \left(\lim_{n\to 0} \bigwedge_{i=1}^{n} \phi_i\right) \cup \{\varphi\} \text{ is inconsistent.}$$

We are here evidently at the metalogical level. Intuitively, $\mathrm{LIM}_{\Vdash}^{\infty}$ states that as the number (the conjunction) of the premises approaches infinity, the stronger the consistency (or satisfiability) relation between the premises and the conclusion is. On the contrary, according to $\mathrm{LIM}_{\Vdash}^{0}$, as the number of premises approaches zero, the strength of this relation vanishes, and the inconsistency (or unsatisfiability) increases.

This can be interpreted as a mathematical answer to *Hume's Problem (of Induction)*. David Hume (1748/1999) famously claimed that we are not justified to believe that the sun will rise tomorrow on the grounds that it has done so until today. According to him, this is just a psychological (survival) strategy based on association and habit. The problem with this is, according to Hume, that we believe the future to be identical with the past, a belief we are perhaps not justified in holding. This is supposed, among other things, to have made induction inadequate as a scientific

[18] This latter identity is better expressed in FOL: the premises of the raven argument are all identical in the sense that they all express formally the existential proposition $\exists x\, (Raven\,(x) \wedge Black\,(x))$.

[19] Clearly, this concept is compatible with a probability viewpoint of inductive logic, but here we keep them separate. See Paris & Vencovská (2015) for one such viewpoint in a comprehensive monograph. See Woods (2002) for a critical exposition of a probability calculus in relation to inductive logic.

methodology, as it is often the case in a scientific context that we believe a conclusion to be confirmed or justified by the repeated verifications expressed in the premises. The philosophical discussion on this problem is prolific, but we consider that the LIM conditions do, if not solve the problem, strongly reduce its impact.

In any case, what $\mathrm{LIM}^{\infty}_{\Vdash}$ gives us is that φ remains a *hypothesis*, not the least because enumeration up to infinity is simply not feasible. To illustrate this, it is possible indeed that the sun will not rise tomorrow, but the number of instances in which the sun rose every next day is now so large that we are very much sure (though not absolutely so) that it will do so tomorrow, and this is sureness enough to make us plan for tomorrow. To be sure, there might be a huge volcano eruption tonight that will cause the sun rays not to be able to penetrate earth's atmosphere tomorrow because of an impenetrable smoke and ash cloud, but this does not count as the sun not rising; this has to do with the rotation of the earth and its position with relation to the sun. In turn, $\mathrm{LIM}^{0}_{\Vdash}$ provides us with a means *not* to attempt an inductive inference, that is, not to hypothesize φ from an insufficient number of premises $\phi_i \approx \varphi$.

We can now make the LIM conditions more precise from the viewpoint of inductive inference in the following way:

Proposition 63. *Let* $\Vert\!\sim$ *denote an* inductive consequence relation[20] *and let the symbol* $\not\Vert\!\sim$ *denote the negation of the relation. Then, we have the conditions:*

$$\left(\mathrm{LIM}^{\infty}_{\Vert\sim}\right) \qquad \left(\lim_{n\to\infty}\bigwedge_{i=1}^{n}\phi_i\right) \Vert\!\sim \varphi$$

and

$$\left(\mathrm{LIM}^{0}_{\Vert\sim}\right) \qquad \left(\lim_{n\to 0}\bigwedge_{i=1}^{n}\phi_i\right) \not\Vert\!\sim \varphi$$

The $\mathrm{LIM}_{\Vert\sim}$ conditions guarantee the non-deductive character of inductive logical systems in the sense that, say, an inductive proof system could be believed to "emulate" (or aspire to do so) a deductive proof system.

Recall that φ is a *possible hypothesis* iff $\mathrm{LIM}^{\infty}_{\Vert\sim}$ holds. In what follows, we simplify the set of premises (the evidence) ϕ_i as ϕ, which is justified by the similarity of the ϕ_i; we are also justified in this because we aim at a general characterization of inductive consequence relations. We follow here Flach (2000) closely. We shall be working with the language L and the formulae $\phi, \varphi, \chi \in F^{\mu}_{\mathsf{L}}$ (we shall omit the superscript μ).

Definition 64. Given language L, an *inductive consequence relation* is a relation $\Vert\!\sim\subseteq F_{\mathsf{L}} \times F_{\mathsf{L}}$ (or $\Vert\!\sim\subseteq 2^{F_{\mathsf{L}}} \times F_{\mathsf{L}}$).[21] Given formulae $\phi, \varphi \in F_{\mathsf{L}}$, the *closure of* $\Vert\!\sim$ is a function $\overline{\Vert\!\sim} : 2^{F_{\mathsf{L}}} \times F_{\mathsf{L}}$ s.t. $\overline{\Vert\!\sim_{\phi}} = \{\varphi \mid \phi \Vert\!\sim \varphi\}$ is the *closure of* ϕ *under* $\Vert\!\sim$.

Recall that $\overline{\Vert\!\sim_{\phi}}$ enjoys the SCL condition, in which case it has to be distinguished from classical closure (see above). In any case, $\overline{\Vert\!\sim_{\phi}}$ is subject to restrictions:

[20] It should now be obvious why an inductive consequence relation is non-monotonic. As above, $\Vert\!\sim$ denotes a non-differentiated consequence relation in terms of proof- and model-theory.

[21] This is for the sake of simplicity, i.e. we want to work with *one* possible hypothesis, rather than with multiple hypotheses.

Definition 65. Let $\overline{\|\!\sim_{\phi}}$ be the closure of ϕ under $\|\!\sim$; we say that $\|\!\sim'$ is (at least) as restrictive as $\|\!\sim$ if $\overline{\|\!\sim'_{\phi}} \subseteq \overline{\|\!\sim_{\phi}}$ and it is more restrictive than $\|\!\sim$ if $\overline{\|\!\sim'_{\phi}} \subsetneq \overline{\|\!\sim_{\phi}}$. A set of rules RI' is said to be (at least) as restrictive as a set of rules RI if for every $\|\!\sim$ satisfying RI there is a unique least restrictive $\|\!\sim'$ satisfying RI' s.t. $\|\!\sim'$ is as restrictive as $\|\!\sim$; then we say that $\|\!\sim'$ is the RI'-restriction of $\|\!\sim$.

Definition 66. The following are the properties of a *general* proof-theoretical inductive consequence relation $|\sim$:

1. *Consistency*:
$$(\vdash) \qquad \frac{\phi \,|\sim \varphi}{\nvdash \varphi \to \neg\phi}$$
2. *Left reflexivity*:
$$(\text{LR}) \qquad \frac{\phi \,|\sim \varphi}{\phi \,|\sim \phi}$$
3. *Right reflexivity*:
$$(\text{RR}) \qquad \frac{\phi \,|\sim \varphi}{\varphi \,|\sim \varphi}$$
4. *Left equivalence*:
$$(\text{LEq}) \qquad \frac{\vdash \phi \leftrightarrow \varphi, \phi \,|\sim \chi}{\varphi \,|\sim \chi}$$
5. *Right equivalence*:
$$(\text{REq}) \qquad \frac{\vdash \varphi \leftrightarrow \chi, \phi \,|\sim \varphi}{\phi \,|\sim \chi}$$
6. *Verification* (χ denotes a prediction):
$$(\text{VER}) \qquad \frac{\vdash (\phi \wedge \varphi) \to \chi, \phi \,|\sim \varphi}{\phi \wedge \chi \,|\sim \varphi}$$
7. *Falsification* (χ denotes a prediction):
$$(\text{FALSF}) \qquad \frac{\vdash (\phi \wedge \varphi) \to \chi, \phi \,|\sim \varphi}{\phi \wedge \neg\chi \not|\!\sim \varphi}$$
8. *Right extension* (χ denotes a prediction):
$$(\text{RExt}) \qquad \frac{\vdash (\phi \wedge \varphi) \to \chi, \phi \,|\sim \varphi}{\phi \,|\sim \varphi \wedge \chi}$$

Intuitively, property or rule 1 assures us that the possible hypothesis φ is not inconsistent with the evidence ϕ. Given property 1, we have it that a formula is consistent in an inductive argument either as a hypothesis or as evidence; $\phi \,|\sim \phi$ (property 2) indicates that ϕ is consistent with the theory of the reasoning agent, and $\varphi \,|\sim \varphi$ (property 3) denotes the consistency of the hypothesis with itself, in both cases iff φ is consistent with ϕ. Properties 4 and 5 state that the logical form of evidence and hypothesis is irrelevant. Property 6 states that if the prediction χ is actually observed to hold, then φ remains a possible hypothesis; otherwise, if χ is verified not to hold, then φ is considered to have been refuted (property 7). Finally, property 8 allows us to add any prediction to the possible hypothesis (rather than to the evidence).

In abductive logic, we are interested in *explanatory consequence relations*, whereas in inductive logic our interest falls on *confirmatory consequence relations*.

Definition 67. Let φ be a hypothesis for evidence ϕ. We denote by $\phi \mid\sim \varphi$ the consequence relation between ϕ and φ s.t. ϕ *confirms* φ. This confirmatory consequence relation satisfies the following conditions (or follows the following rules) iff (a) ϕ is consistent and, if there is more than one hypothesis, φ_i are consistent, and (b) the properties of a general inductive consequence relation are satisfied:

1. *Admissible entailment*:

$$\text{(AE)} \qquad \frac{\vdash \phi \to \varphi, \phi \mid\sim \phi}{\phi \mid\sim \varphi}$$

2. *Confirmatory reflexivity*:

$$\text{(CR)} \qquad \frac{\phi \mid\sim \phi, \phi \not\mid\sim \neg\varphi}{\varphi \mid\sim \varphi}$$

3. *(Inductive) Right weakening*:

$$\text{(iRW)} \qquad \frac{\vdash \varphi \to \chi, \phi \mid\sim \varphi}{\phi \mid\sim \chi}$$

4. *Right* $\wedge$:

$$\text{(R}\wedge\text{)} \qquad \frac{\phi \mid\sim \varphi, \phi \mid\sim \chi}{\phi \mid\sim \varphi \wedge \chi}$$

5. *Predictive right weakening*:

$$\text{(PRW)} \qquad \frac{\vdash (\phi \wedge \varphi) \to \chi, \phi \mid\sim \varphi}{\phi \mid\sim \chi}$$

6. *Right consistency*:

$$\text{(R}\vdash\text{)} \qquad \frac{\phi \mid\sim \varphi}{\phi \not\mid\sim \neg\varphi}$$

7. *Left logical equivalence*:

$$\text{(LLEq)} \qquad \frac{\vdash \phi \leftrightarrow \chi, \phi \mid\sim \varphi}{\chi \mid\sim \varphi}$$

We provide a brief explanation for the conditions/rules above. Condition 1 expresses the fact that consistent evidence confirms any of its consequences.

Condition 2 states that a hypothesis that is consistent with the evidence confirms itself. It appears also evident that if a hypothesis is confirmed by the evidence, then any hypothesis entailed by this hypothesis is also confirmed by the evidence (condition 3).

Condition 4 is evident, too: if evidence ϕ confirms hypotheses φ and χ, then it confirms the set $\{\varphi, \chi\} = (\varphi \wedge \chi)$.

Condition 5 is equivalent to a combination of conditions 3 and 8 in Def. 66; in effect, PRW implies right weakening, since $\vdash \varphi \to \chi$ implies $\vdash (\phi \wedge \varphi) \to \chi$, and given that this implies $\vdash (\phi \wedge \varphi) \to (\varphi \wedge \chi)$, PRW also implies right extension. Then, given a confirmatory argument, PRW states that any prediction is confirmed by the same evidence.

Condition 6 expresses the obvious fact that if φ belongs to the set of hypotheses confirmed by the evidence ϕ, then $\neg\varphi$ is not a member of this set.

Finally, condition 7 tells us that if evidence ϕ confirms hypothesis φ, then any evidence χ logically equivalent to ϕ also confirms hypothesis φ.

These conditions or rules were largely – but more or less loosely – based on Hempel's (1945) adequacy conditions that any definition of (inductive) confirmation should satisfy. Although Hempel's adequacy conditions are established with the scientific method in mind, they are interesting to our study because they are actually conceived from the viewpoint of logical consequence. We thus leave them here for the comparison with the above conditions.[22] It is easy to see how these conditions can extrapolate the scientific method taken in a stricter sense and be applied to other contexts where inductive confirmation is a concern, once we make sure that Hempel's notion of consequence fits, or can be made to fit, our logical definition of consequence.[23]

Definition 68. (Hempel, 1945) Any adequate definition of *confirmation* must insure the fulfillment of the following conditions:

1. *Entailment condition*: Any sentence that is entailed by the evidence[24] is confirmed by it.

2. *Consequence condition*: If the evidence confirms every one of a set of sentences K,[25] then it also confirms any sentence which is a logical consequence of K.

   (a) *Special consequence condition*: If the evidence confirms a hypothesis h, then it also confirms any consequence of h.

   (b) *Equivalence condition*: If the evidence confirms a hypothesis h, then it also confirms any hypothesis that is logically equivalent with h.

3. *Consistency condition*: Every logically consistent evidence is logically equivalent with the set of all the hypotheses which it confirms.

   (a) Unless the evidence is self-contradictory, it does not confirm any hypothesis with which it is not logically compatible.

[22] See Flach (2000).

[23] This caveat is relevant, because Hempel distinguishes the formal from the experimental sciences in that the former require no empirical test (*in principle!*). However, recall that in inductive logic we are in the domain of supraclassicality, which admits of empirical components in a way that deductive logic does not (cf. rule SUB).

[24] *An observation report*, in the original jargon. As seen above, by "evidence" we mean a set of sentences whose members are $\phi_i = \phi_j$ for $i \neq j$.

[25] Originally, *a class* of sentences.

(b) Unless the evidence is self-contradictory, it does not confirm any hypotheses which contradict each other.

Hempel (1945) complemented the above conditions with satisfaction, disconfirmation, and neutrality criteria.

Definition 69. Hempel (1945):

1. *Satisfaction criterion of confirmation*: evidence b
   (a) directly confirms a hypothesis h if b entails the development of h for the class of objects that are mentioned in b.[26]
   (b) confirms a hypothesis h if h is entailed by a set of sentences each of which is directly confirmed by b.
2. *Disconfirmation criterion*: evidence b disconfirms a hypothesis h if it confirms the negation of h.
3. *Neutrality criterion*: evidence b is neutral with respect to a hypothesis h if b neither confirms nor disconfirms h.

The above conditions for an inductive confirmatory logical consequence relation (Def.s 66-7) actually provide us with a Gentzen-like proof system for inductive logic that we shall call $\mathcal{CS}$. The logical rules are those of $\mathcal{LK}$ and all the above can be seen as the structural rules that together with these constitute the proof system $\mathcal{CS}$.

In order to prove the soundness of $\mathcal{CS}$ Flach (2000) elaborated on a confirmatory semantics. In terms of a semantics for the confirmatory inductive consequence relation, it makes sense to consider a hypothesis to be confirmed iff it is true in all models of some given semantics. This is the starting point of the *confirmatory semantics* proposed by Flach (2000), of which we provide the main points (in an adapted notation).

Definition 70. Given a set of formulae $F_{\mathsf{L}} \subseteq \mathsf{L}$, $\phi, \varphi \in F_{\mathsf{L}}$, a *confirmatory structure* is a triple $\mathcal{C} = (S, |\cdot|, \|\cdot\|)$, where S is a set of semantic objects, and $|\cdot|, \|\cdot\| : F_{\mathsf{L}} \longrightarrow 2^S$ are functions s.t. $|\phi|$ denotes the set of regular models constructed from the premises ϕ and $\|\varphi\|$ denotes the models for the conclusion (the hypothesis). $\mathcal{C}$ defines a *closed confirmatory consequence relation* $|\approx$ as $\phi |\approx \varphi$ iff

1. $|\phi| \neq \emptyset$, and
2. $|\phi| \subseteq \|\varphi\|$.

Definition 71. A confirmatory structure $\mathcal{C}$ is said to be *simple* if it satisfies:

1. Reflexivity

[26] By *the development of hypothesis h for a finite class of individuals C*, Hempel (1945, p. 109) provides the following explanation: "the development of h for C states what h would assert if there existed exclusively those objects which are elements of C. Thus, e.g., the development of the hypothesis $h_1 = \forall x\,(P\,(x) \vee Q\,(x))$ for the class $\{a, b\}$ is $(P\,(a) \vee Q\,(a)) \wedge (P\,(b) \vee Q\,(b))$." Hempel (1945, p. 109) gives further examples.

2. iRW
3. R$\wedge$
4. $|\phi| \subseteq \|\phi\|$
5. $\|\phi \wedge \varphi\| = \|\phi\| \cap \|\varphi\|$
6. $\|\neg\phi\| = S - \|\phi\|$
7. $\|\phi\| = S$ iff $\models \phi$

Simple confirmatory structures are axiomatizable by the proof system $\mathcal{CS}$. Details of the rules of inference are as follows:

1. The rules of $\mathcal{CS}$ are PRW, R$\wedge$, and R$\vdash$.
2. The derived rules of $\mathcal{CS}$ include iRW, RExt, LR, AE, and $\vdash$.[27]

A simple confirmatory consequence relation $\|\!\sim$ is sound if it satisfies the rules of inference of $\mathcal{CS}$. We now turn to the confirmatory semantics, in order to prove the completeness of the simple confirmatory consequence relation $|\approx$.

Definition 72. Let $|\approx$ be a confirmatory consequence relation for a confirmatory structure $\mathcal{C} = (U, |\cdot|, \|\cdot\|)$, where U is the set of models of F_{L} in consideration. A *model* $\mathcal{M} \in U$ is said to be *normal* for ϕ, denoted by $\hat{\mathcal{M}}_\phi$, iff for all $\varphi \in F_{\mathsf{L}}$ s.t. $\phi \,|\!\approx \varphi$ we have $\models_{\mathcal{M}} \varphi$.

Proposition 73. *Suppose that the consequence relation $|\approx$ satisfies iRW and R$\wedge$. Let ϕ be an admissible formula. Then, all normal models for ϕ satisfy φ iff $\phi \,|\!\approx \varphi$.*

Definition 74. Given $\mathcal{C}$, we can now make the following further definitions with respect to the confirmatory consequence relation $|\approx$ for arbitrary ϕ:

1. $|\phi| = \begin{cases} \left\{\mathcal{M} \in U \,|\, \hat{\mathcal{M}}_\phi\right\} & \text{if } \phi \text{ is admissible} \\ \emptyset & \text{otherwise} \end{cases}$;
2. $\|\varphi\| = \{\mathcal{M} \in U \,|\, \models_{\mathcal{M}} \varphi\}$.

Theorem 75. *A consequence relation $|\approx$ is simple confirmatory iff it satisfies the rules of $\mathcal{CS}$.*

See Flach (2000) for the relevant proofs. The strong assumptions of the confirmatory semantics above (i.e. the evidence is subject to completeness assumptions) are not always satisfiable, reason why one may need to work with a *weaker* notion of confirmation, in which a confirmed hypothesis is required to be true only in some of the regular models.

Definition 76. Let $\mathcal{C}$ be a confirmatory structure as above. We say that $\mathcal{C}$ defines an *open confirmatory consequence relation* $\phi \,|\!\approx \varphi$ iff $(|\phi| \cap \|\varphi\|) \neq \emptyset$.

[27] Rules LLEq and CR pose additional restrictions on $\mathcal{C}$, to wit, and respectively, (i) if $\models \phi \leftrightarrow \varphi$, then $|\phi| = |\varphi|$, and (ii) if $\|\varphi\| \neq \emptyset$, then $|\varphi| \neq \emptyset$.

This *openness* can be radicalized by an identity between $|\cdot|$ and $\|\cdot\|$, which entails dropping the condition that regular models be models of the premises:

Definition 77. A simple confirmatory structure $\mathcal{C}^{\diamond} = (S, \|\cdot\|, \|\cdot\|)$ is called a *classical confirmatory structure.* A consequence relation $\mid\approx_{\diamond}$ is said to be *weak confirmatory* iff it is the open confirmatory consequence relation defined by $\mathcal{C}^{\diamond}$.

Flach (2000) then presents a proof system for $\mid\approx_{\diamond}$, which we shall call $\mathcal{C}^{\diamond}\mathcal{S}$.

Definition 78. A weak confirmatory consequence relation $\mid\approx_{\diamond}$ satisfies the rules of $\mathcal{C}^{\diamond}\mathcal{S}$, to wit, $\vdash$, PRW, and the following two further rules:

1. *Predictive convergence*:[28]

$$\text{(PC)} \qquad \frac{\vdash (\phi \wedge \chi) \rightarrow \varphi, \phi \mid\sim \chi}{\varphi \mid\sim \chi}$$

2. *Disjunctive rationality*:

$$\text{(DR)} \qquad \frac{\phi \vee \chi \mid\sim \varphi, \chi \not\mid\sim \varphi}{\phi \mid\sim \varphi}$$

Finally, Flach (2000) also provides an account of *explanatory consequence relations* in the context of inductive logic. Our treatment of these relations was carried out above in the Subsection dedicated to abduction.

3.3 Probabilistic Consequence

As seen above, if we are dealing with uncertainty, then the deduction theorem (DT) is bound to be of no (much) use to us in a formal approach to reasoning. In the so-called probability logics,[29] DT may fail, because we are now interested in the probability that, given $\phi \Vdash \psi$ (or $\Gamma \Vdash \Delta$), it is the case that $\Vdash \phi \rightarrow \psi$ (or $\Vdash \Gamma \rightarrow \Delta$, respectively). For example, let ϕ stand for "Patient has acne" and ψ for "Patient has spots". Clearly, $\phi \| \sim \psi$ holds, as $\psi \in Cn^{\sim}(\{\phi\})$, but this neither allows us to deduce that $\phi \rightarrow \psi$, nor does it give us a degree of certainty with respect to $\Vdash \phi \rightarrow \psi$. Indeed, the patient with acne may instead have deep-buried cysts.

This involves in fact the notion of conditional probability (see below), namely in the sense that the conditional connective requires or is amenable to a probabilistic interpretation (e.g., Adams, 1975). From this viewpoint, the introduction of probabilities for sentences generalizes deductive logic (see Carnap, 1950), but it may simply be seen as a degree of implication (or confirmation), in which case we are in fact in the field of inductive logic.

[28] Note that this is a combination of VER and the rule for convergence CON, which expresses a *monotonicity* property for induction, i.e. the fact that rejecting a hypothesis is not a defeasible process, but must be based solely on the evidence:

$$\text{(CON)} \qquad \frac{\vdash \phi \rightarrow \varphi, \phi \mid\sim \chi}{\varphi \mid\sim \chi}$$

PC is thus a strengthening of convergence in the sense that φ can be any set of predictions.

[29] We carry out here a treatment of probability logic highly circumscribed to the notion of consequence; more comprehensive treatments are to be found in, e.g., Roeper & Leblanc (1999) and Hailperin (1996).

Definition 79. A probability function $Pr : F \rightarrow [0,1]$ is said to satisfy the classical probability axioms[30] on a set of sentences $F \subseteq \mathsf{L}$ iff Pr satisfies the following conditions for $\phi, \psi \in F$:

1. $Pr(\phi) \geq 0$.
2. If $\Vdash \phi$, then $Pr(\phi) = 1$.
3. If $\Vdash \neg(\phi \wedge \psi)$, then $Pr(\phi \vee \psi) = Pr(\phi) + Pr(\psi)$.

The above definition has two immediate consequences:

The first is well known from probability theory:

$$Pr(\neg\phi) = 1 - Pr(\phi)$$

The second consequence is important from the viewpoint of the Lindenbaum-Tarski algebras:

$$Pr(\phi) = Pr(\psi) \text{ if } \Vdash \phi \leftrightarrow \psi$$

Conditions 2 and 3 above are parasitic on deductive logic: condition 2 uses the notion of tautology and condition 3 that of logical incompatibility. In effect, Def. 79.2 above can be reformulated as:

Theorem 80. $\Vdash \phi$ *iff* $Pr(\phi) = 1$.

Given this, if we define 0 as falsity and stipulate the remaining truth values $\mathtt{v}_i$ for $0 < i < 1$, then probability logic can be seen as reducing to a many-valued logic. (See Augusto, 2020b.)

5.3.4. A probability function $Pr : F \rightarrow [0,1]$ is a *conditional probability* if, whenever $Pr(\psi) > 0$,

$$Pr(\phi|\psi) := \frac{Pr(\phi \wedge \psi)}{Pr(\psi)}$$

or, whenever $Pr(\Delta) > 0$,

$$Pr(\Gamma|\Delta) := \frac{Pr(\Gamma \wedge \Delta)}{Pr(\Delta)}$$

where both Γ and Δ can be singletons.

[30]Recall the Kolmogorov axioms for a probability space $(\mathbf{U}, \mathscr{F}, Pr)$ where $\mathbf{U}$ is the universal set, $\mathscr{F}$ is a σ-field on $\mathbf{U}$, and $A, B \in \mathscr{F}$:

1. $Pr(A) \geq 0$. (*Non-negativity*)
2. $Pr(\mathbf{U}) = 1$. (*Normalization*)
3. $Pr(A \cup B) = Pr(A) + Pr(B)$ if $A \cap B = \emptyset$. (*Finite additivity*)

See Woods (2002) for a basic logical probability calculus. We remark that the Kolmogorov axioms are so for a probability measure defined on a Boolean set algebra; thus, their application in propositional logic poses no problems. Moreover, their extension to FOL can be carried out by resorting to Lindenbaum-Tarski algebras and first-order models thereof.

Note that Pr does not necessarily entail a semantical interpretation: the interval $[0,1]$ just is the interval of the probabilities that can be assigned to a formula or a set of formulae. This does not require that a truth-value assignment be expanded to the truth-value set $[0,1]$, i.e. $val_{Pr} : F \longrightarrow [0,1]$. In effect, we can envisage a probabilistic consequence relation as a purely quantitative relation in which at play is merely *probability-preservation* (vs. truth-preservation). However, if we choose to extend the semantics of $\mathsf{L}^{(*)}$ by considering the interval $[0,1]$ as our new truth-value set, then we are extending classical logic.

Be it as it may, Pr just did not come out of the blue, and we have to specify where it belongs, i.e. to the proof-theoretical or to the model-theoretical components of $\mathsf{L}^{(*)}$. This motivates complications that can be eschewed if we opt to remain at the metalogical level, i.e. if we work with a consequence relation $\Vdash_x$, where x is a probability. In this case, a metalogical expression such as $\Gamma \Vdash_x \phi$ is read as "Γ partially entails ϕ to degree x." The "rule" for *partial entailment* is formulated as follows.

Definition 81. We say that a set of formulae $\Gamma = \{\psi_1, ..., \psi_n\}$ *partially entails* a formula ϕ *to degree* x if the following rule is satisfied:

$$\text{(PEnt)} \qquad \Gamma \Vdash_x \phi \quad \text{iff} \quad Pr\left(\phi | \bigwedge_{i=1}^{n} \psi_i\right) = x$$

We say that Γ *entails* ϕ if $x = 1$. If $\Gamma = \emptyset$, we speak of *unconditional probability.*

It is easy to see that partial entailment makes of a probabilistic logic a deductive logic, and a monotonic one, for that matter. Evidently, PEnt is basically not acceptable in relevance logic, as not all the $\psi_i \in \Gamma$ may be relevant for the conditional probability of ϕ given Γ (see Augusto, 2020a, 4.5.1). Moreover, PEnt poses the problem of the uniqueness of the probability function depending only on Γ and ϕ, i.e. a distinguished probability function Pr' such that $\Gamma \Vdash_x \phi$ iff $Pr'(\phi|\Gamma) = x$ and x is a degree of confirmation; it appears, however, that no such function can in fact be unique, especially so in the case of FOL and of infinite sets. At play is here what for many makes of probabilistic logics either minor logics or non-logics simpliciter, to wit, the introduction of factual – i.e. empirical – considerations so as to be able to attempt to determine the uniqueness of x in terms of degree of confirmation. The problem is aggravated if one sees the degree of confirmation as the de-facto *degree of consequence*, as is expressed by PEnt (e.g., Roeper & Leblanc, 1999). In any case, we are here already in the domain of inductive logic, and it is not certain when the transition from deductive logic did in fact occur.[31]

One can drop the uniqueness requirement by, instead, conceiving of a probabilistic consequence relation defined as

$$\Gamma \Vdash_A \phi \quad \text{iff} \quad A = \left\{x \,|\, x = Pr\left(\phi | \bigwedge \Gamma\right)\right\}$$

but this *generalized partial entailment* is weaker than PEnt in that, are Γ and ϕ logically unrelated, then $A = [0,1]$, and the entailment relation $\Vdash_A$ is trivial. Clearly,

[31] For an elaboration on the problems of partial entailment and some remedying proposals thereto, see, e.g., Williamson (2002).

restrictions have to be set upon A, the set of probability functions, but this again takes us to the realm of the empirical – notoriously so if Γ is, say, a knowledge base.

Hawthorne and Makinson (Hawthorne, 2014; Hawthorne & Makinson, 2007) have come up with a family of rules for consequence relations – the family $\mathscr{O}$ – by "revising" the notion of partial entailment with a *threshold* associated to the consequence relation. Given any probability function Pr defined on sentences of a propositional language, Pr satisfies conditions 1-3 above, they define a probabilistic consequence relation for $0 < t \leq 1$ by the following rule:

Definition 82. The pair (Pr, t) generates a *probabilistic consequence relation* $\|\!\sim_{Pr,t}$ according to the rule:

$$\phi \|\!\sim_{Pr,t} x \text{ iff } \begin{cases} Pr(\phi) = 0 \\ Pr(x|\phi) \geq t \end{cases}$$

Definition 83. The family $\mathscr{O}$ of rules for consequence relations is defined as follows:

$(\mathcal{O}1)$	$\phi \mid\sim \phi$	*Reflexivity*
$(\mathcal{O}2)$	If $\phi \mid\sim x$ and $x \vdash y$, then $\phi \mid\sim y$	*Right weakening*
$(\mathcal{O}3)$	If $\phi \mid\sim x$ and $\phi \dashv\vdash \psi$, then $\psi \mid\sim x$	*Left classical equivalence*
$(\mathcal{O}4)$	If $\phi \mid\sim x \wedge y$, then $\phi \wedge x \mid\sim y$	*Very cautious monotonicity*
$(\mathcal{O}5)$	If $\phi \wedge \psi \mid\sim x$ and $\phi \wedge \neg\psi \mid\sim x$, then $\phi \mid\sim x$	*Weak* $\vee$
$(\mathcal{O}6)$	If $\phi \mid\sim x$ and $\phi \wedge \neg y \mid\sim y$, then $\phi \mid\sim x \wedge y$	*Weak* $\wedge$

The main limitations of this family of rules for consequence relations are that they are solely for consequence relations that are in Horn rule form, i.e. in the form "If $\phi_1 \mid\sim x_1, ..., \phi_n \mid\sim x_n$, then $\psi \mid\sim y$" for a finite number of premise conditions $\phi_n \mid\sim x_n$, and the family $\mathscr{O}$ does not appear to make for a sound set of rules. In any case, probability logic is plagued with issues to do with spatial and temporal complexity, and is basically undecidable.

To finish our short approach to probabilistic consequence relations, we wanted to show that the transition from a propositional to a first-order system is not necessarily over-complicated, despite the essentially undecidable nature of the latter. Moreover, we wanted to give an example of a model-theoretical probabilistic consequence relation – mostly because a purely proof-theoretically based probabilistic consequence relation is hard to come by.

The probabilistic logic ε-logic, where ε is a fixed error parameter (Terwijn, 2005; Kuyper & Terwijn, 2013), satisfies both wishes. In particular, it is an inductive probabilistic logic, though the parameter ε is conceived as ε-*truth*, and thus points to a deductive preservation of truth: in fact, the focus here is on *probabilistic truth*.

We shall be working with the language L^* with, if necessary, the equality symbol added.

Definition 84. Let $\mathcal{M}$ be a classical first-order model for L^* and let $\mathcal{D}$ be a probability measure on $\mathcal{M}$. Given $\varepsilon \in [0, 1]$, we inductively define the relation $\models_{(\mathcal{M},\mathcal{D}),\varepsilon}$ (we omit the subscript $(\mathcal{M}, \mathcal{D})$ henceforth) as follows:

1. For every atomic formula ϕ,
$$\models_{\varepsilon} \phi \text{ if } \models_{\mathcal{M}} \phi.$$
2. For $\wedge, \vee$ it holds that
   (a)
   $$\models_{\varepsilon} \phi \wedge \psi \text{ if } \models_{\varepsilon} \phi \text{ and } \models_{\varepsilon} \psi;$$
   (b)
   $$\models_{\varepsilon} \phi \vee \psi \text{ if } \models_{\varepsilon} \phi \text{ or } \models_{\varepsilon} \psi.$$
3. For $\exists$ it holds that
$$\models_{\varepsilon} \exists x \phi(x)$$
if there exists an $a \in \mathcal{M}$ such that $\models_{\varepsilon} \phi(a)$.
4. For $\forall$ it holds that
$$\models_{\varepsilon} \forall x \phi(x) \text{ if } Pr[a \in \mathcal{M} \mid \models_{\varepsilon} \phi(a)] \geq 1 - \varepsilon.$$
5. For $\neg$ it holds that
   (a) for atomic ϕ, $\models_{\varepsilon} \neg\phi$ if $\not\models_{\varepsilon} \phi$;
   (b) $\models_{\varepsilon} \neg(\phi \wedge \psi)$ if $\models_{\varepsilon} \neg\phi \vee \neg\psi$, and
   (c) $\models_{\varepsilon} \neg(\phi \vee \psi)$ if $\models_{\varepsilon} \neg\phi \wedge \neg\psi$;
   (d) $\models_{\varepsilon} \neg\neg\phi$ if $\models_{\varepsilon} \phi$;
   (e) $\models_{\varepsilon} \neg(\phi \to \psi)$ if $\models_{\varepsilon} \phi \wedge \neg\psi$;
   (f) $\models_{\varepsilon} \neg\exists x \phi(x)$ if $\models_{\varepsilon} \forall x \neg\phi(x)$;
   (g) $\models_{\varepsilon} \neg\forall x \phi(x)$ if $\models_{\varepsilon} \exists x \neg\phi(x)$.
6. For $\to$ it holds that
$$\models_{\varepsilon} \phi \to \psi \text{ if } \models_{\varepsilon} \neg\phi \vee \psi.$$

As is evident, the approach is thoroughly classical except for the universal quantifier. In effect, we cannot even call it "universal," as now the consequence relation $\models_{\varepsilon} \phi(a)$ defined in terms of $\forall$ means that for many/most $a \in \mathcal{M}$, where "many/most" depends on ε, it holds that $\models_{\varepsilon} \phi(a)$. This is believed to solve the undecidability of sentences with the universal quantifier given a finite amount of information. Importantly, because both $\models_{\varepsilon} \exists x \phi(x)$ and, given this interpretation of the "universal" quantifier, $\models_{\varepsilon} \forall x \neg\phi(x)$ may hold simultaneously, this is a *probabilistic paraconsistent* logic.[32] This necessarily motivates a non-classical definition of satisfiability and validity:

[32] Paraconsistency is the view that one can work, from the logical point of view, with information characterized by inconsistency and contradiction in a non-trivial way. More formally: a logic is said to be paraconsistent iff it is not the case that for all formulae ϕ, ψ we have $\phi, \neg\phi \Vdash \psi$ (*v. 1*) or iff it is the case that for formulae ϕ, ψ we have $\Vdash \phi$ and $\Vdash \neg\phi$, but not $\Vdash \psi$ (*v. 2*). Note that these are in fact definitions of the paraconsistent consequence relation as a *non-explosive consequence relation.*

Definition 85. A formula $\phi(x_1, ..., x_n)$ is *ε-satisfiable* if there are an ε-model $(\mathcal{M}, \mathcal{D})$ and $a_1, ..., a_n \in \mathcal{M}$ such that $\models_{(\mathcal{M},\mathcal{D}),\varepsilon} \phi$. We write $\models_\varepsilon \phi$ when ϕ is an *ε-tautology* or is *ε-valid*, i.e. when $\models_\varepsilon \phi(a_1, ..., a_n)$ holds for all ε-models $(\mathcal{M}, \mathcal{D})$ and all $a_1, ..., a_n \in \mathcal{M}$.

In the ε-logic, the model $\mathcal{M}$ and the probability distribution tango to provide the desirable property that every set is measurable:

Proposition 86. *For an ε-model $(\mathcal{M}, \mathcal{D})$, it holds that*

1. *For all formulae $\phi(x_1, ..., x_n)$ and all $a_1, ..., a_{n-1} \in \mathcal{M}$, the set*

$$\{a_n \in \mathcal{M} \mid \models_\varepsilon \phi(a_1, ..., a_n)\}$$

   is $\mathcal{D}$-measurable.

2. *All n-ary relations, including equality, are $\mathcal{D}^n$-measurable and all n-ary functions are measurable as mappings $(\mathcal{M}^n, \mathcal{D}^n) \longrightarrow (\mathcal{M}, \mathcal{D})$, where $\mathcal{D}^n$ denotes the n-fold product measure. All constants are $\mathcal{D}$-measurable.*

Finally,

Proposition 87. *The set $\{a_1, ..., a_n \in \mathcal{M}^n \mid \models_\varepsilon \phi(a_1, ..., a_n)\}$ is $\mathcal{D}^n$-measurable.*

All this accounts for the (desired) *algorithmic, finitist* character of proofs in ε-logic, though the computational complexity is basically very hard. (See Kuyper, 2014, for a more comprehensive elaboration on both further requirements for proofs and the issue of complexity in ε-logic.)

Acknowledgments

This article is essentially a reprint of contents first published in Augusto (2020a). My thanks to College Publications, London, for allowing authors to keep the copyright of their books and to reprint parts of them.

References

Adams, E. (1975). *The logic of conditionals.* Dordrecht: Reidel.

Aliseda, A. (2006). *Abductive reasoning. Logical investigations into discovery and explanation.* Berlin: Springer.

Aliseda, A. (2007). Abductive reasoning: Challenges ahead. *Theoria*, *60*, 261-270.

Allemang, D., Tanner, M., Bylander, T., & Josephson, J. (1987). Computational complexity of hypothesis assembly. *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, 1112-1117.

Augusto, L. M. (2019). *Formal logic: Classical problems and proofs.* London: College Publications.

Augusto, L. M. (2020a). *Logical consequences: Theory and applications. An introduction.* 2nd ed. London: College Publications.

Augusto, L. M. (2020b). *Many-valued logics: A mathematical and computational introduction*. 2nd ed. London: College Publications.

Augusto, L. M. (2020c). Toward a general theory of knowledge. *Journal of Knowledge Structures & Systems, 1*(1), 63-97.

Beth, E. W. (1955). Semantic entailment and formal derivability. *Mededelingen van de Koninklijke Nederlandse Akademie van Wetenschappen, Afdeling Letterkunde*, *18*, 309-42.

Bylander, T., Allemang, D., Tanner, M. C., & Josephson, J. R. (1991). The computational complexity of abduction. *Artificial Intelligence*, *49*, 25-60.

Carnap, R. (1950). *Logical foundations of probability.* Chicago: University of Chicago Press.

Cox, P. T. & Pietrzykowsky, T. (1987). General diagnosis by abductive inference. *Proceedings IEEE Symposium on Logic Programming*, 183-189.

de Kleer, J. & Williams, B. (1987). Diagnosing multiple faults. *Artificial Intelligence*, *32*, 97-130.

de Kleer, J., Mackworth, A. K., & Reiter, R. (1992). Characterizing diagnoses and systems. *Artificial Intelligence*, *56*, 197-222.

Flach, P. A. (2000). Logical characterizations of inductive learning. In D. M. Gabbay & R. Kruse (eds.), *Abductive reasoning and learning* (pp. 155-196), Norwell, MA: Kluwer.

Flach, P. A. & Kakas, A. C. (eds.) (2000). *Abduction and induction: Essays on their relation and integration.* Dordrecht: Kluwer.

Gabbay, D. M. (2014). What is a logical system? An evolutionary view: 1964-2014. In D. M. Gabbay, J. H. Siekmann, & J. Woods (eds.), *Handbook of the history of logic. Vol. 9: Computational logic* (pp. 41-132). Amsterdam, etc.: Elsevier.

Gabbay, D. M. & Woods, J. (2001). The new logic. *Logic Journal of the IGPL*, *9*, 141-174.

Gabbay, D. M. & Woods, J. (2003). *A practical logic of cognitive systems. Vol. 1: Agenda relevance. A study in formal pragmatics.* Amsterdam, etc.: Elsevier.

Gabbay, D. M. & Woods, J. (2005). *A practical logic of cognitive systems. Vol. 2: The reach of abduction: Insight and trial.* Oxford, etc.: Elsevier.

Gabbay, D. M., Hartmann, S. & Woods, J. (eds.) (2011). *Handbook of the History of Logic. Vol. 10: Inductive Logic.* Amsterdam/Oxford: North-Holland.

Gentzen, G. (1934-5). Untersuchungen über das logische Schliessen. *Mathematische Zeitschrift*, *39*, 176-210, 405-431. (Engl. trans.: Investigations into logical deduction. In M. E. Szabo (ed.), *The Collected Papers of Gerhard Gentzen* (pp. 68-131). Amsterdam: North-Holland.)

Hailperin, T. (1996). *Sentential probability logic. Origins, development, current status, and practical applications.* London: Associated University Presses.

Harman, G. H. (1965). The inference to the best explanation. *Philosophical Review*, *74*, 88-95.

Hawthorne, J. (2014). New Horn rules for probabilistic consequence: Is O+ enough? In S. O. Hansson (ed.), *David Makinson on classical methods for non-classical problems* (pp. 119-141). Dordrecht, etc.: Springer.

Hawthorne, J. & Makinson, D. C. (2007). The quantitative/ qualitative watershed for rules of uncertain inference. *Studia Logica*, *86*, 247-297.

Hempel, C. G. (1945). Studies in the logic of confirmation. *Mind*, *54*, 1-26 (Part I), 97-121 (Part II).

Hobbs, J. R., Stickel, M. E., Appelt, D. E., & Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence*, *63*, 69-142.

Hume, D. (1999). *An enquiry concerning human understanding.* Ed. Tom L. Beauchamp. Oxford & New York: OUP. (Originally published in 1748.)

Kant, I. (1787). *Kritik der reinen Vernunft.* 2. Auflage. Riga: Johan Friedrich Hartknoch. (Engl. trans.: *Critique of pure reason.* Trans. by N. Kemp Smith. London: MacMillan & Co., 1929.)

Kuyper, R. (2014). Computational aspects of satisfiability in probability logic. *Mathematical Logic Quarterly*, *60*, 444-470.

Kuyper, R. & S. A. Terwijn (2013). Model theory of measure spaces and probability logic. *The Review of Symbolic Logic*, *6*, 367-393.

Makinson, D. C. (2003). Bridges between classical and nonmonotonic logic. *Logic Journal of the IGPL*, *11*, 69-96.

Makinson, D. C. (2005). How to go nonmonotonic. In D. M. Gabbay & F. Guenthner (eds.), *Handbook of philosophical logic.* 2nd ed. Vol. 12 (pp. 175-278), Amsterdam: Springer.

McCarthy, J. (1959). Programs with common sense. *Proceedings of the Symposium on Mechanization of Thought Processes*, 77-84.

McCarthy, J. & Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (eds.), *Machine Intelligence* (pp. 463–502), Edinburgh: Edinburgh University Press.

Minsky, M. (1961). Steps toward artificial intelligence. *Proceedings of the IRE*, *49*, 8-30.

Minsky, M. (1974). A framework for representing knowledge. Report AIM, 306, Artificial Intelligence Laboratory, MIT.

Minsky, M. (1991). Logical versus analogical or symbolic versus connectionist or neat versus scruffy. *AI Magazine*, *12*, 34-51.

Muggleton, S. & de Raedt, L. (1994). Inductive logic programming: Theory and methods. *Journal of Logic Programming*, *19*, 629-679.

Ng, H. T. & Mooney, R. J. (1992). Abductive plan recognition and diagnosis: A comprehensive empirical evaluation. *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, Cambridge, MA, 499-508.

Paris, J. & Vencovská, A. (2015). *Pure inductive logic.* Cambridge: CUP.

Paul, G. (1993). Approaches to abductive reasoning: An overview. *Artificial Intelligence Review*, *7*, 109-152.

Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, *13*, 81-132.

Roeper, P. & Leblanc, H. (1999). *Probability theory and probability logic.* Toronto, etc.: University of Toronto Press.

Shanahan, M. (2005). Perception as abduction: Turning sensor data into meaningful representation. *Cognitive Science*, *29*, 103-134.

Smullyan, R. M. (1968). *First-order logic.* Mineola, NY: Dover.

Stenning, K., & van Lambalgen, M. (2008). *Human reasoning and cognitive science.* Cambridge, MA: MIT Press.

Tarski, A. (1930). Fundamentale Begriffe der Methodologie der deduktiven Wissenschaften. I. *Monatshefte für Mathematik und Physik*, *37*, 361-404. (Engl. trans.: Fundamental concepts of the methodology of the deductive sciences. In A. Tarski, *Logic, semantics, metamathematics: Papers from 1923 to 1938* (pp. 60-109). Oxford: Clarendon Press, 1956.)

Terwijn, S. A. (2005). Probabilistic logic and induction. *Journal of Logic and Computation*, *15*, 507-515.

Williamson, J. (2002). Probability logic. In D. M. Gabbay et al. (eds.), *Handbook of the logic of argument and inference. The turn towards the practical* (pp. 397-424), Amsterdam: North-Holland.

Wójcicki, R. (1988). *Theory of logical calculi: Basic theory of consequence operations.* Dordrecht: Kluwer.

Woods, J. (2002). Standard logics as theories of argument and inference: Induction. In D. M. Gabbay et al. (eds.), *Handbook of the logic of argument and inference. The turn towards the practical* (pp. 105-170), Amsterdam: North-Holland.

Cite this article as:

Augusto, L. M. (2023). Supraclassical consequence: Abduction, induction, and probability for commonsense reasoning. *Journal of Knowledge Structures & Systems*, *4*(1), 1-46.

EDITORIAL INFORMATION

Editor-in-chief: Luis M. Augusto
Reprint (Source: Augusto, 2020a)