Extending Depth-Bounded Reasoning to First-Order Logic

Marcello D'Agostino
University of Milan

Costanza Larese
University of Milan

Sanjay Modgil
King's College, London

Abstract

This paper extends the framework of Depth-Bounded Boolean Logics to first-order logic, developing an intuitive "informational" semantics for quantifiers, a corresponding hierarchy of approximations, and a proof-theoretical characterization via suitable natural deduction rules.

1 Introduction

Classical logic models *logically omniscient agents*, that is, agents—whether human, artificial, or hybrid — who can recognize all valid consequences of their assumptions. This normative ideal, however, clashes with the reality of reasoning under limited cognitive and computational resources. As Gabbay and Woods observe:

> A logic is an idealization of certain sorts of real-life phenomena. By their very nature, idealizations misdescribe the behaviour of actual agents. This is to be tolerated when two conditions are met. One is that the actual behaviour of actual agents can defensibly be made out to approximate to the behaviour of the ideal agents of the logician's idealization. The other is the idealization's facilitation of the logician's discovery and demonstration of deep laws [7, p. 158].

Thank the referees (or not), and anyone else for reasons relating to the paper as a whole, by adding \titlethanks{text} *to the preamble* (or \titlethanks{} to delete this reminder).

For acknowledgements specific to the author (e.g. thanking a particular author's funding body), use \thanks{text} after the author name in the \addauthor command (or do not add the command at all for no such addition).

The *approximation problem* is therefore to define, for a given logic L, a hierarchy of systems that (i) converge to L, (ii) preserve rational standards at each stage, and (iii) provide realistic models of the deductive capacities of resource-bounded agents.

The *depth-bounded approach* offers a systematic solution. Its central insight is the distinction between two kinds of information:

- **Actual information:** information explicitly possessed by an agent or feasibly accessible;
- **Virtual information:** information not actually possessed but temporarily assumed *as if* it were possessed, in order to advance inference. Footnote Intuition Pump

The aim is to minimize recourse to virtual information, thereby controlling combinatorial explosion[1].

This approach, which stems from [6], has been thoroughly investigated for classical propositional logic in a series of papers (e.g. [2, 5, 4]) and in a recent book by the present authors with Dov Gabbay [8]. At depth 0, reasoning is restricted to propositional inferences that can be justified in terms of actual information alone. This yields a tractable subsystem whose semantics can be captured in several equivalent ways, including a non-deterministic three-valued matrix first noted (but never properly investigated) by Quine [9] to represent what he called "the primitive meaning of the logical operators" (see [3]) and later rediscovered by Crawford and Etherington [1] in connection with unit resolution. The 0-depth logic validates exactly those inferences that can be justified without recourse to virtual assumptions.

Higher layers of the hierarchy are generated by controlling the nested uses of virtual information. Increasing the permitted depth of ~~this nesting~~ nested applications of this rule yields a sequence of sound and paracomplete systems of growing inferential power and computational complexity, which in the limit converge to full classical logic [4]. The depth parameter thus measures the cognitive and computational effort required to *dig out* implicit consequences from the premises.

From the semantic distinction between actual and virtual information, one obtains a parallel proof-theoretical distinction, captured by *intelim rules* (introduction and elimination rules for Boolean operators) together with a single branching rule (RB). This is a structural rule—essentially a restricted cut, related to the principle of bivalence of classical logic—that governs the introduction of virtual information. Depth-0 reasoning is captured by intelim sequences, while deeper approximations are represented by intelim trees. Each system satisfies a subformula property

[1] In proof-theory the use of virtual information is apparent in the "discharge rules" of Gentzen-style natural deduction. As argued in [8], these rules are not suited to provide an adequate measure of depth for deductive reasoning

and supports analytic proof search, so that tractability arises as a consequence of the informational perspective rather than from an *ad hoc* syntactic restriction.

Conceptually, the depth-bounded framework provides a fresh account of the informational content of deduction. An inference is of depth 0 when possessing actual information about its premises suffices, without supplementation, to possess actual information about its conclusion. At higher depths, inference becomes conditional on simulating virtual extensions of the actual information state. This perspective dissolves the traditional "scandal of deduction" (see [6]), since it recognizes that logical inference is rarely trivial even at the propositional level: it typically involves a non-trivial deployment of virtual information, which incurs real cognitive and computational costs.

The aim of the present paper is to extend the depth-bounded approach, originally developed for Boolean logic, to the domain of *first-order reasoning*. We discuss an intuitive informal semantics for the quantifiers that mirrors the distinction between actual and virtual information, and combine it with a depth-bounded proof theory to obtain a hierarchy of approximations to classical first-order logic. This is an exploratory Festschrift piece, emphasizing ideas over formal details or technical results, that continues Gabbay's long-standing interest in resource-sensitive and controlled derivation. Formal aspects and results will be investigated in a subsequent paper.

Section 2 reviews the propositional framework of depth-bounded Boolean logics (DBBL). Section 3 discusses the informational meaning of the quantifiers focusing on the information we actually possess about *objects*, rather than propositions. The exposition is informal and technical details are omitted. Section 4 presents the intelim rules and the RB rule for first-order inference and Section 5 integrates the Boolean and quantificational components into a bivariate measure of depth, proposing a unified view of bounded reasoning. Section 6 concludes with remarks on conceptual implications and directions for further research.

2 An overview of Depth-bounded Boolean Logics

In DBBL the underlying semantic notions are not classical truth and falsity, but *informational* truth and falsity. A formula P is *informationally true* if we actually possess the information that P is true, *informationally false* if we actually possess the information that P is false, and *informationally indeterminate* otherwise. By "actually possessing" a piece of information we mean that either we possess it explicitly as part of our data or it is feasibly accessible to us by a chain of simple inferences each of which descends immediately from the *informational meaning* of the Boolean operators—that is, that part of their meaning that can be specified solely in terms of informational truth and informational falsity (see [8, Chapters 1-2]). For conceptual clarity we use *signed formulas*, i.e., expressions of the form $T\,P$ and $F\,P$ where the former means that P is informationally true and the latter that P is informationally false. Information states are typically partial and therefore it cannot be assumed that for every information state and for every formula P either $T\,P$ or $F\,P$.

The *intelim rules* for DBBL express the most basic inferences that can be immediately justified solely in terms of informational truth and falsity. While all the rules that are sound for informational truth and falsity are also sound for their classical "metaphysical" counterparts that obey the principle of bivalence, some of the rules that are sound for classical truth and falsity are obviously unsound for their informational counterparts. For example, under the informational reading of the signs T and F, from $T\,P \vee Q$ it does not follow that either $T\,P$ or $T\,Q$. Indeed, very often we are informed that a certain disjunction is true without holding any specific information about either disjunct ("either the password or the username is incorrect"). Similarly, we may well be informed that a conjunction is false without holding any specific information about either conjunct; hence $F\,P \wedge Q$ does not imply that either $F\,P$ or $F\,Q$. On the other hand, the intelim rules can be easily seen as sound under the informational interpretation of the signs.

Footnote To be clear: classical semantics would stipulate that if in fact T P∨Q (the world makes it so that P∨Q) then in fact either TP or TQ. However one may possess the information that T P∨Q without.... That said, from an informational ... we can achieve ... we have the information that TP or TQ. This informational move can be effectively be formalised in ~~...~~ ~~...~~ deploying an informational move ~~that~~ such that for any P we can assume TP or FP. Bivalence — we can assume we are in possession of TP or FP

The intelim rules for the four standard Boolean operators are illustrated in Table 1. Their unsigned version in Table 2 is sufficient for all practical purposes, albeit semantically less transparent. (See [5] for a generalization to arbitrary Boolean operators.)

The intelim rules—or equivalently, the informational semantics mentioned in the introduction and discussed in [8]—characterize the basic layer of 0-depth inference. Inferences based only on the intelim rules can be arranged in sequences such that each signed formula in the sequent is either a premise or follows from previous formulas by applying one of the rules.

To characterize classes of increasingly deeper inferences and, asymptotically, full classical propositional logic we only need to add a single branching rule:

$$T\,P \qquad F\,P \qquad \text{or} \qquad P \qquad \neg P \qquad \text{for unsigned formulas.}$$

RB governs the use of virtual information, i.e., information that we do not actually possess, but must be simulated in order to reach deeper conclusions: in each branch we suppose a jump to a hypothetical refinement of the current information state, in which we actually possess either the information that P is true or the information that P is false.

$$\frac{F\,A}{T\,\neg A}\,F\neg I \qquad \frac{T\,A}{F\,\neg A}\,T\neg I \qquad \frac{T\,\neg A}{F\,A}\,T\neg E \qquad \frac{F\,\neg A}{T\,A}\,F\neg E$$

$$\frac{T\,A}{T\,A \vee B}\,T \vee I_1 \qquad \frac{T\,B}{T\,A \vee B}\,T \vee I_2 \qquad \frac{\begin{array}{c}F\,A\\ F\,B\end{array}}{F\,A \vee B}\,F \vee I$$

$$\frac{\begin{array}{c}T\,A \vee B\\ F\,A\end{array}}{T\,B}\,T \vee E_1 \qquad \frac{\begin{array}{c}T\,A \vee B\\ F\,B\end{array}}{T\,A}\,T \vee E_2 \qquad \frac{F\,A \vee B}{F\,A}\,F \vee E_1 \qquad \frac{F\,A \vee B}{F\,B}\,F \vee E_2$$

$$\frac{F\,A}{F\,A \wedge B}\,F \wedge I_1 \qquad \frac{F\,B}{F\,A \wedge B}\,F \wedge I_2 \qquad \frac{\begin{array}{c}T\,A\\ T\,B\end{array}}{T\,A \wedge B}\,T \wedge I$$

$$\frac{\begin{array}{c}F\,A \wedge B\\ T\,A\end{array}}{F\,B}\,F \wedge E_1 \qquad \frac{\begin{array}{c}F\,A \wedge B\\ T\,B\end{array}}{F\,A}\,F \wedge E_2 \qquad \frac{T\,A \wedge B}{T\,A}\,T \wedge E_1 \qquad \frac{T\,A \wedge B}{T\,B}\,T \wedge E_2$$

$$\frac{F\,A}{T\,A \to B}\,T\!\to I_1 \qquad \frac{T\,B}{T\,A \to B}\,T\!\to I_2 \qquad \frac{\begin{array}{c}T\,A\\ F\,B\end{array}}{F\,A \to B}\,F\!\to I$$

$$\frac{\begin{array}{c}T\,A \to B\\ T\,A\end{array}}{T\,B}\,T\!\to E_1 \qquad \frac{\begin{array}{c}T\,A \to B\\ F\,B\end{array}}{F\,A}\,T\!\to E_2 \qquad \frac{F\,A \to B}{T\,A}\,F\!\to E_1 \qquad \frac{F\,A \to B}{F\,B}\,F\!\to E_2$$

Table 1: Intelim rules for the four standard Boolean operators.

Intelim rules can be combined with RB to yield *intelim trees*. These are downward-growing trees in which each branch corresponds to a possible information state. Branching occurs only through RB, which introduces the virtual assumptions $T\,P$ and $F\,P$ on two sibling branches.

Intelim trees An *intelim tree* for a set X of signed formulae is a finite tree such that each node is either: (i) an assumption from X, (ii) the conclusion of an application of an intelim rule to earlier nodes on the same branch, or (iii) a virtual assumption introduced by an application of RB. The *depth* of an intelim tree is the maximum number of nested applications of RB.

A branch is *closed* if it contains both $T\,P$ and $F\,P$ for some formula A; otherwise it is *open*. A signed formula sP (where s is either T or F) is *deducible* at depth k from a set X of signed formulas if there is an intelim tree of depth k whose open branches all end with Y.

$$\frac{A}{\neg\neg A}\,\neg I \qquad \frac{\neg\neg A}{A}\,\neg\neg E$$

$$\frac{A}{A \lor B}\,\lor I_1 \qquad \frac{B}{A \lor B}\,\lor I_2 \qquad \frac{\begin{array}{c}\neg A\\ \not{B}\end{array}}{\neg(A \lor B)}\,\neg\lor I$$

$$\frac{\begin{array}{c}A \lor B\\ \neg A\end{array}}{B}\,\lor E_1 \qquad \frac{\begin{array}{c}A \lor B\\ \neg B\end{array}}{A}\,\lor E_2 \qquad \frac{\neg(A \lor B)}{\neg A}\,\neg\lor E_1 \qquad \frac{\neg A \lor B}{\neg B}\,\neg\lor E_2$$

$$\frac{\neg A}{\neg(A \land B)}\,\neg\land I_1 \qquad \frac{\neg B}{\neg(A \land B)}\,\neg\land I_2 \qquad \frac{\begin{array}{c}A\\ B\end{array}}{A \land B}\,\land I$$

$$\frac{\begin{array}{c}\neg(A \land B)\\ A\end{array}}{\neg B}\,\neg\land E_1 \qquad \frac{\begin{array}{c}\neg(A \land B)\\ B\end{array}}{\neg A}\,\neg\land E_2 \qquad \frac{A \land B}{A}\,\land E_1 \qquad \frac{A \land B}{B}\,\land E_2$$

$$\frac{\neg A}{A \to B} \to I_1 \qquad \frac{B}{A \to B} \to I_2 \qquad \frac{\begin{array}{c}A\\ \neg B\end{array}}{\neg(A \to B)}\,\neg\!\to I$$

$$\frac{\begin{array}{c}A \to B\\ A\end{array}}{B} \to E_1 \qquad \frac{\begin{array}{c}A \to B\\ \neg B\end{array}}{\neg A} \to E_2 \qquad \frac{\neg(A \to B)}{A}\,\neg\!\to E_1 \qquad \frac{\neg(A \to B)}{\neg B}\,\neg\!\to E_2$$

Table 2: Intelim rules for the four standard Boolean operators (unsigned version).

Note that according to the above definition, if an intelim tree of depth k is closed (i.e., every branch is closed), then any sP is deducible from the premises. Hence, k-depth consequence is *explosive*. However, for each k there are classically inconsistent sets of formulae that are not inconsistent at depth k.

Example: The classically inconsistent set

$$T\,p \lor q, T\,p \lor \neg q, T\,\neg p \lor q, T\,\neg p \lor \neg q$$

is consistent at depth 0.

Hence, depth-bounded consequence remains monotonic and reflexive but not complete with respect to classical entailment.

Figure 1 shows three examples of intelim trees of depth 0, 1 and 2 respectively. Formulas are labelled with the name of the rule from which they are obtained; unlabelled formulas are premises or virtual information introduced via RB. Note that in the third tree the conclusion u follows from the premises because it occurs in all open branches. Note also that intelim trees can also be used to prove tautologies directly from no premise. An example of a derivation of the excluded middle law is shown in Figure 2.

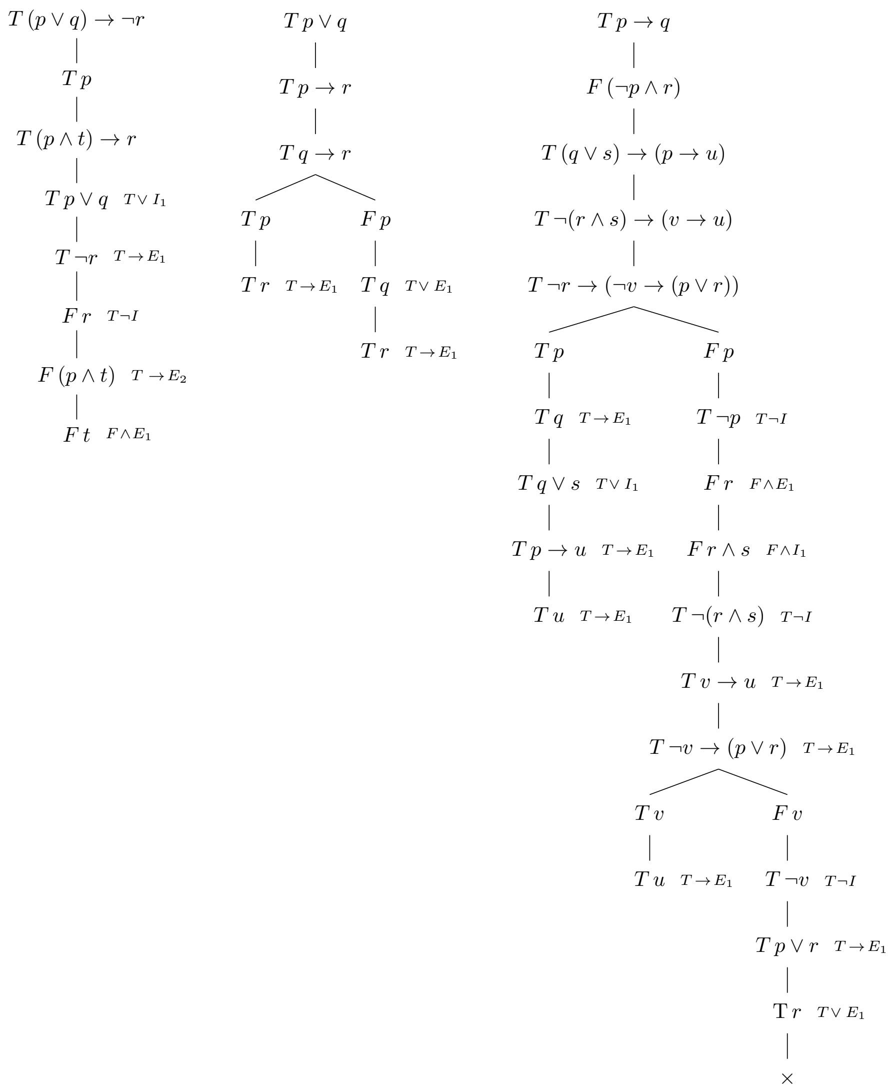

Figure 1: intelim trees of depth 0, 1 and 2 respectively.

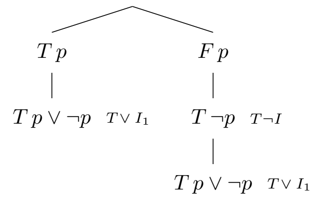

Figure 2: A direct proof of the excluded middle law.

3 The informational meaning of quantifiers

When the informational perspective is extended from propositional to quantified reasoning, the focus shifts from information about *propositions* to information about *objects*. From this standpoint, the objects that are introduced by quantifier elimination are not metaphysical individuals but *informational objects*: bundles of properties providing a partial description of unknown objects in the domain. For ease of exposition, we shall call them *informons*.

An informon carries information about how such unknown objects are partially characterized within an information state, and thus represents the objectual dimension of information. Each informon represents an agent's current information about a certain object: the properties that are known to hold of it and those that remain undetermined. A property expressed by a signed formula in one free variable $T\,A(x)$ or $F\,A(x)$ holds of an informon ι if $T\,A(x)$, respectively $F\,A(x)$, belongs to the bundle of properties associated to ι. Informons then provide *approximations* to unknown objects that are typically refined as new information is acquired either from an external source or through (depth-bounded) inference. From In a purely informational perspective, the unknown objects themselves could be identified with *ideal informons*: complete (possibly infinite) descriptions in which all properties that can be expressed in the given language are either informationally true or informationally false. This is the limit to which the approximations converge. A "real" informon, on the other hand, is always an approximation of some ideal informon, consisting of all its properties that can be determined in a given information state.

To emphasize the approximation process, we use the same fictitious name (a *parameter*) for the successive approximations of the *same* ideal informon. A parameter thus persists as the identifier of a single ideal informon, while its associated "real" informon refines over time as new information is acquired. For example, if we are informed that some creature has a heart, we may choose a parameter c to denote an unknown ideal informon that provides a complete description of one of the individual creatures which have a heart. At this stage we have only a minimal dispositional account ✓

approximation for c, namely the informon ι_1 whose bundle contains $T\,\mathsf{HasHeart}(x)$. Learning that all creatures with a heart have kidneys puts us in a disposition to assent to $\mathsf{HasKidneys}(c)$; the bundle for c therefore expands to the informon ι_2 containing both $T\,\mathsf{HasHeart}(x)$ and $T\,\mathsf{HasKidneys}(x)$, which represent the best approximation to c available in the current information state. Hence we are in a disposition to assent to $\mathsf{HasKidneys}(c)$.

To summarize:

idealized

- parameters statically denote largely unknown objects (ideal informons) that live in *complete* information states, purely *ideal* entities that are highly inaccessible;
- the information content actually carried by a parameter c is dynamically determined by the current information state and consists of an informon ι associated to c that refines as our information grows[2].

We say that $T\,Ax$ $(F\,A(x))$ *holds* of an informon ι to mean that the property expressed by the signed formula in one free variable $T\,Ax$ $(F\,A(x))$ belongs to the bundle of properties that characterize ι. We also write $\iota_1 \sqsubseteq \iota_2$ to mean that ι_1 is an *approximation* of ι_2, that is, $\iota_1 = \iota_2$ or is less defined than it [3].

In general, the following holds:

Informon inheritance property Given any pair of informons ι_1 and ι_2, and any property expressed by the signed formula $T\,Ax$ $(F\,A(x))$,

$$\iota_1 \sqsubseteq \iota_2 \quad \Longrightarrow \quad \text{if } T\,Ax\ (F\,Ax) \text{ holds of } \iota_1, \text{ then } T\,Ax\ (F\,A(x)) \text{ holds of } \iota_2.$$

A remarkable feature of depth-bounded reasoning is that an informon which is consistent at a given depth, may turn out to be inconsistent at a greater depth, although we shall not elaborate on this aspect in the context of this preliminary informal exposition. An example will suffice: if our information state is such that the following set of signed formulas

$$\{T\,A(c) \vee B(c), T\,A(c) \vee \neg B(c), T\,\neg A \vee B(c), T\,\neg A \vee \neg B(c)\}$$

[2]This picture fits the classical *monotonic* view of object-level information. In a non-monotonic view, informons can also *shrink*.

[3]There is a structural analogy with Scott domains [10, 11], where elements are partially defined objects ordered by information content and $\bot$ denotes the least defined element. Here, informons form a partially ordered set under inclusion of properties. Distinct minimal informons carrying independent properties need not have a common refinement, although their intersection—and hence a meet—is always defined. Informational reasoning thus unfolds along directed subsets whose suprema correspond to limits of consistent information growth rather than to computational approximations.

hold in it, the informon currently associated with c is always consistent if the allowed propositional depth is 0, but becomes inconsistent at depth 1.

At the propositional level, on the other hand, information is carried by discrete *information tokens* (or, more briefly, *veritons*), which record the status of propositions (informationally true, informationally false, or indeterminate). Informons and veritons therefore play similar roles in the informational hierarchy: the former track how properties attach to (otherwise undetermined) partial descriptions of ideal informons, while the latter track how informational truth and falsity attach to (otherwise undetermined) propositions. Their interaction will underpin the extension of depth-bounded reasoning to first order logic.

To summarize, informons are partial specifications of objects that approximate their ideal complete descriptions. Two kinds of informons must be distinguished.

Arbitrary (least defined) informons. These correspond to the so-called *arbitrary objects* introduced by universal quantifiers and by the negation of existentials. Each of them represents an object of which no information is available besides the properties that are known to hold of all informons in the domain. We may denote the arbitrary informon associated to a given information state by $\bot_{\mathsf{o}}$[4]. In the limit, i.e., in a complete information state, $\bot_{\mathsf{o}}$ is the *intersection of all ideal informons* that live in its domain. This limit is itself approximated by "real" informons that live in the domains of partial information states.

At each stage of the inferential process, an arbitrary informon can therefore be identified with the bundle of properties representing the general information currently held about *all* (unknown) objects in the domain. If no general information is available, the bundle of properties initially associated to $\bot_{\mathsf{o}}$ is empty, but may expand at later stages. For example, if we are informed that the domain of discourse consists only of (Euclidean) triangles, the property of being a triangle holds of $\bot_{\mathsf{o}}$, but the property of being a right triangle does not. Later, we may discover new properties that hold of all informons in this domain. For example, that the sum of the internal angles of any triangle equals 180°. Each such discovery enlarges the bundle of properties belonging to $\bot_{\mathsf{o}}$. In this sense, learning and reasoning can refine the *information content carried by* $\bot_{\mathsf{o}}$ without altering its structural role as the least defined element of the ordering. Note that by the informon inheritance principle, whatever property holds of an informon, holds also of any refinement of it.

Reasoning with arbitrary informons therefore preserves maximal generality with respect to the domain of discourse, while yielding no object-specific information.

[4]This usage parallels that adopted for the undefined truth value of a proposition in the informational semantics of DBBL (see [8]). In both cases, $\bot$ marks the least informative element of the corresponding domain—the starting point of informational growth.

Specific informons. These are introduced by existential quantifiers and by the negation of universals. A specific informon embodies the assumption that the properties that hold of it, hold of some unknown object (ideal informon), though most of its properties are not yet determined. Through successive inferences the information attached to it can be refined, producing progressively more defined informons. Such refinement mirrors the way witnesses are used in classical proof theory, but here it is viewed as a controlled expansion of the agent's information state. For example, in the domain of Euclidean triangles, the informon characterized by the property of having two equal sides may later be refined by adding the property that two of its angles are equal. Note that a specific informon is always a refinement of some arbitrary informon which is the bottom element of the information ordering. Hence, by the informon inheritance principle, any property that holds of an arbitrary informon holds also of any specific informon.

For future reference, the principal symbols employed in the informational semantics are summarized in Table 3.

Symbol	**Informal meaning**
$T\,A$	the information that the proposition A is true
$F\,A$	the information that the proposition A is false
$\bot$	undefined truth value (no information about A)
ι	a generic informon (bundle of properties, object–level information)
$\bot_{\circ}$	arbitrary (least defined) informon
$\iota_1 \sqsubseteq \iota_2$	ι_1 approximates ι_2 (is equal or less defined than ι_2)

Table 3: Basic symbols for informational semantics.

What do we mean when we say that we *actually possess the information* that a sentence of the form $\forall x\, A$ or $\exists x\, A$ is true, or that it is false?

Let us start with the notion of *actually possessing* the information that $\forall x\, A$ is true. The answer cannot be that we literally possess the information that $A[x/a]$ is true for *all* the unknown objects of the domain. A more feasible and information-driven answer is the following: we are in the disposition to assent to *any* sentence of the form $A[x/a]$ where a stands for any ideal informon that is approximated by the real one currently associated to $\bot_{\circ}$. Since any informon is a refinement of the arbitrary informon $\bot_{\circ}$, being in a disposition to assent to $A(c)$ for *any* ideal informon c is tantamount to saying that $T\,A[x/a]$ holds for the arbitrary (least defined) informon $\bot_{\circ}$.

Let us now turn to the notion of *actually possessing* the information that $\exists x\, Ax$ is true. This means that there is a specific ideal informon that satisfies $T\, A(x)$. We may not yet know which further properties distinguish b from other, compatible ideal informons: the commitment is only that some (as yet largely undetermined) ideal informon satisfies $T\, A(x)$ as well as all the properties that hold of $\perp_{\circ}$. Accordingly, we are in the disposition to assent to $A[x/b]$, where b stands for some ideal informon currently approximated by a real informon ι that represents the approximation of b in the current information state; subsequent inference may *refine* the informon associated to b by adding further properties.

A similar account applies to falsity. To *actually possess* the information that $\forall x\, Ax$ is false is to hold that there is a specific ideal informon that satisfies $F\, A(x)$. We are therefore in the disposition to dissent from $A[x/b]$, where b stands for some ideal informon currently approximated by the minimal informon ι such that $F\, A(x)$ holds of ι; as new information is acquired, the informon associated to b may be refined by adding properties that provide a better approximation. On the other hand, to actually possess the information that $\exists x A(x)$ is false means to be in a disposition to dissent from $A[x/a]$ where a names the arbitrary, least defined, informon and, therefore, from $A[x/b]$ for any specific ideal informon b.

A simple example will help clarify. Suppose we are informed that

$$\textit{Something is either a blick or a gleem.} \tag{1}$$

At this stage, we are in a disposition to assent to "b is a blick or a gleem" for some unkown object b about which this is the only information we possess. So the informon associated to b at this stage is just the minimal one whose only characterizing property is $T\, \mathit{blick}(x) \vee \mathit{gleem}(x)$. Next, we learn that

$$\textit{All blicks are plocks.} \tag{2}$$

Now, the property $T\, \mathit{blick}(x) \rightarrow \mathit{plock}(x)$ holds of $\perp_{\circ}$ and hence also of the informon associated to b. Finally we learn that

$$\textit{All gleems are plocks.} \tag{3}$$

The current approximation of $\perp_{\circ}$ thus has both implications in its bundle, and the informon associated with b is updated accordingly.

If no virtual information is allowed, the informon associated to b at depth 0 is stable. At depth 1, on the other hand, we can distinguish the cases $T\, \mathit{blick}(b)$ and $T\, \mathit{gleem}(b)$, each of which corresponds to a better approximation to the ideal informon b. Since in either case we infer $\mathit{plock}(b)$, the informon associated to b now supports also $T\, \mathit{plock}(x)$.

?

This example shows how the semantics of informational objects mirrors the growth of information brought about by additional premises and by logical inference. Moreover, it shows how the best approximations of the unknown objects in the domain that is available at depth $k > j$ can be more accurate than the best available at depth j.

4 Intelim and structural rules for FOL

In this section we propose intelim rules for quantifiers. These rules are formulated in accordance with their informational meaning explained in the previous section. We also introduce a version of RB that is adequate for quantificational reasoning.

Recall that in our framework a *parameter* is a syntactic symbol that identifies an ideal informon introduced in the course of inference. Parameters serve as persistent names for these ideal informons and may take the form of a functional term to track dependencies among them.

We use three sorts of parameters:

- **γ-parameters (arbitrary).** Stand for the same arbitrary (least defined) informon $\bot_{\mathsf{o}}$. They behave like free variables and carry no dependency arguments. Note that γ-parameters cannot occur in the formula to which the branching rule RB is applied, since branching would contradict their role as names of the arbitrary informon $\bot_{\mathsf{o}}$, which by definition remains minimally informative and cannot split into alternative informational developments. ?

- **δ-parameters (specific).** Stand for specific ideal informons introduced by $T\,\exists$ or $F\,\forall$ elimination. They may carry a dependency vector of currently active γ's, written $d(\vec{\gamma})$ (pure bookkeeping).

- **ε-parameters (neutral).** Introduced only by the branching rule RB. They may later be *declared arbitrary-for-$T\,A$* (or *arbitrary-for-$F\,A$*) in exactly one of the sibling subtrees generated by the branching; thereafter they behave as γ's in that subtree and as δ's in the sibling subtree.

Typographic convention. We distinguish parameter sorts by their initial letter (indices are used only when needed):

$$\underbrace{g, g_1, g_2, \dots}_{\gamma\ \text{(arbitrary)}} \qquad \underbrace{d, d_1, d_2, \dots}_{\delta\ \text{(specific)}} \qquad \underbrace{e, e_1, e_2, \dots}_{\varepsilon\ \text{(neutral)}}$$

A δ-parameter may carry a *dependency vector* of the currently active γ's, written $d(g_1, \dots, g_k)$, as a bookkeeping device.

Global discipline. (i) No parameter occurs in the *premises.* (ii) Declarations for ε's are *subtree-local.*

We use A^x_t to denote the result of substituting all occurrences of x in A with the parameter t. Conversely, we use A^t_x to denote the result of substituting all occurrences of t in A with the variable x.

$T\,\forall$ and $F\,\exists$ elimination.

$$\frac{T\,\forall x\, A}{T\, A^x_g}\ (T\,\forall E_\gamma) \qquad \frac{F\,\exists x\, A}{F\, A^x_g}\ (F\,\exists E_\gamma)$$

Side condition: g is a γ-parameter that does not occur in $A(x)$. It need not be globally fresh; the same g may be reused in other applications of $(\forall E_\gamma)$ provided it satisfies this local condition.

Although all γ-parameters represent the same arbitrary informon $\bot_{\mathsf{o}}$, they are treated syntactically as distinct placeholders. This allows simultaneous reasoning about several occurrences of an arbitrary object—for instance in $R(g_1, g_2)$—without attributing informational differences to the objects themselves. The differentiation among γ's is purely syntactic, reflecting distinct *uses* of the same informational constant rather than distinct entities. In this respect, γ-parameters play the role of *eigenvariables*: temporary names for an arbitrary ideal informon, semantically identical to all others but locally distinguished so that reasoning steps remain independent. The need for several γ-parameters becomes apparent when the rules that introduce them are applied in connection with the following substitution rule.

γ-Substitution.

$$\frac{T\, A(g)}{T\, A^g_t}\ (\gamma\text{-Subst}) \qquad \text{if } g \text{ is a } \gamma \text{ and } t \text{ is any term free for } g \text{ in } A.$$

Reading: γ's behave like free variables; once $T\,A(g)$ is obtained, g may be instantiated by any term t. Therefore, in $R(g_1, g_2)$ the two arbitrary parameters can be instantiated by (possibly) distinct terms.

$T\,\forall$ and $F\,\exists$ introduction.

$$\frac{T\, A(g)}{T\,\forall x\, A^x_g}\ (T\,\forall I_\gamma) \qquad \frac{F\, A(g)}{F\,\exists x\, A^g_x}\ (F\,\exists I_\gamma)$$

Side condition: g is a γ-parameter.

When the truth of an existential or the falsity of a universal enters our information state, reasoning creates a new *specific informon* representing the (unknown) object whose existence is informationally asserted. In the syntax, this object is represented by a *new parameter*. Parameters are persistent identifiers of *ideal* informons: each parameter continues to denote the same ideal informon while its associated *real* informon—the current approximation available in the information state—is progressively refined through inference. When the *ideal informon* introduced by $T\,\exists A$ or $F\,\forall A$ depends on previously introduced arbitrary informons, its parameter takes the functional form $d(\vec{\gamma})$, where the elements of $\vec{\gamma}$ are the γ-parameters occurring in A. The functional notation merely records informational dependencies; it has no semantic import. Introducing a term such as $d(\vec{\gamma})$ to replace $T\,\exists A$ or $F\,\forall A$ should therefore not be read as postulating a function of $\vec{\gamma}$, but as introducing a new specific informon that depends on the terms that replace the γ's when the γ-substitution rule is applied.

$T\,\exists$ and $F\,\forall$ elimination.

$$\frac{T\,\exists x\,A}{T\,A^x_{d(\vec{\gamma})}}\;(T\,\exists E_\delta) \qquad\qquad \frac{F\,\forall x\,A}{F\,A^x_{d(\vec{\gamma})}}\;(F\,\forall E_\delta)$$

Side condition: $d(\vec{\gamma})$ is a fresh δ-parameter whose argument list consists exactly of the γ's active at the introduction site.

$T\,\exists$ and $F\,\forall$ introduction.

$$\frac{T\,A(t)}{T\,\exists x\,A^t_x}\;(T\,\exists I_\delta) \qquad\qquad \frac{F\,A(t)}{F\,\forall x\,A^t_x}\;(F\,\forall I_\delta)$$

Micro-example From $T\,\forall x\,(A(x) \to B(x))$ and $T\,\exists x\,A(x)$, infer $T\,\exists x\,B(x)$:

$$\begin{array}{ll}
T\,\forall x\,(A(x) \to B(x)) & \text{premise} \\
T\,\exists x\,A(x) & \text{premise} \\
\hline
T\,A(d) & T\,\exists E_\delta \\
T\,(A(g) \to B(g)) & T\,\forall E_\gamma \\
T\,(A(d) \to B(d)) & \gamma\text{-Subst} \\
T\,B(d) & T\,\to E \\
T\,\exists x\,B(x) & T\,\exists I_\delta
\end{array}$$

We call a *Q-formula* any formula whose main logical operator is a quantifier. We say that a Q-formula is *fully instantiated* in a branch r of an intelim tree when all its initial quantifiers have been eliminated (replaced by γ's and δ's) in r.

For example: the formula $\forall x \exists y\, R(x, y)$ is fully instantiated in r if r contains $R(g, d)$ for some γ-parameter g and some δ-parameter d.

Remark 4.1. *If a Q-formula of the form $Q_1x_1 \cdots Q_nx_n\ A(x_1, \ldots, x_n)$ is fully instantiated in a branch while the main operator of A is Boolean and some subformula of A is itself a Q-formula, that inner Q-formula may not be fully instantiated in that branch, depending on the upper bound on the propositional depth of the intelim tree. In the micro-example above, all Q-formulas are fully instantiated at depth 0. If we add the premise*

$$T\,\forall x \forall y\ \big(\exists z\,(R(x, y) \wedge R(y, z))\ \vee\ \exists z\,(S(x, y) \wedge S(y, z))\big),$$

it becomes fully instantiated (without branching) as

$$T\ \exists z\,(R(g_1, g_2) \wedge R(g_2, z))\ \vee\ \exists z\,(S(g_1, g_2) \wedge S(g_2, z)),$$

but the existential subformulas inside the disjunction cannot be fully instantiated at depth 0, because the disjunction may not be unpacked unless we are informed of the falsity of one disjunct. At depth 1, via RB, both existentials are fully instantiated in at least one branch. In general, there is always a propositional depth at which all the Q-formulas occurring in the premises become fully instantiated.

RB and neutral parameters The RB rule takes the same branching form as in the propositional case:

$$\begin{array}{ccc} & \wedge & \\ T\,A & & F\,A \end{array}$$

with the following provisos: (i) the formula A contains no γ-parameter; and (ii) at the split one may optionally introduce fresh ε-parameters as neutral names on the child branches; ε's have no arguments.

Applications of RB may introduce *fresh* neutral ε-parameters. As long as it stays neutral, an ε-parameter e represents the *minimal virtual informon* for which $A(x)$ is decided: along each child subtree, the current virtual information state contains either $T\,A(e)$ or $F\,A(e)$. No special informational role is fixed yet. Such a neutral ε may later be declared (as needed) to play the role of a γ-parameter on exactly one of the sibling subtrees generated by the branching; it then behaves like a δ on the other:

$$\begin{array}{ccc} & \wedge & \\ T\,A(e)\ \boxed{e : \gamma} & & F\,A(e) \end{array} \qquad \begin{array}{ccc} & \wedge & \\ T\,A(e) & & F\,A(e)\ \boxed{e : \gamma} \end{array}$$

In the first case we say that e has been declared *arbitrary-for-$T\,A$*, and this licenses an application of $T\,\forall$-introduction. In the second, e has been declared *arbitrary-for-$F\,A$*, which licenses an application of $F\,\exists$-introduction. For example:

$$\begin{array}{cc}
T\,A(e)\ \boxed{e:\gamma} & F\,A(e) \\
| & | \\
T\,\forall x A(x) & F\,\forall x A(x) \\
| & | \\
T\,\forall x A(x) \vee \neg\forall x A(x) & T\,\neg\forall x A(x) \\
 & | \\
 & T\,\forall x A(x) \vee \neg\forall x A(x)
\end{array}$$

$$\begin{array}{cc}
T\,A(e) & F\,A(e)\ \boxed{e:\gamma} \\
| & | \\
T\,\exists x A(x) & F\,\exists x A(x) \\
| & | \\
T\,\exists x A(x) \vee \neg\exists x A(x) & T\,\neg\exists x A(x) \\
 & | \\
 & T\,\exists x A(x) \vee \neg\exists x A(x)
\end{array}$$

Declaration discipline

- Declarations are subtree-local and cannot be made with contradictory polarities on sibling subtrees (the same e cannot be simultaneously arbitrary-for-$T\,A$ and arbitrary-for-$F\,A$ in sibling subtrees).
- Just as γ-parameters cannot occur in the formula to which RB is applied, an ε declared to be a γ in a subtree cannot occur in any application of RB *below* the node containing the declaration. Equivalently, RB can be applied to formulas containing only δ's or neutral ε's.

Quantifier introduction via neutral ε-parameters.

$$\frac{T\,B(e)}{T\,\forall x\,B^{e}_{x}}\ (T\,\forall I_{\varepsilon})$$

Side condition: $T\,B(e)$ occurs on every open branch that contains e. For example:

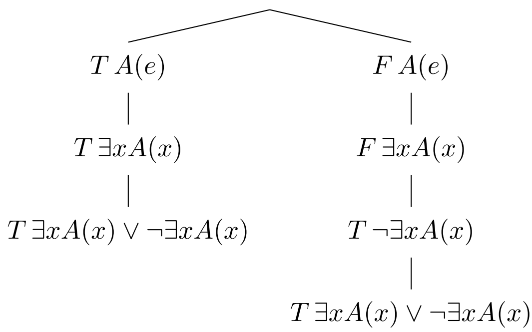

Remark 4.2 (Relation to Hilbert's ε)**.** *ε-parameters are reminiscent of Hilbert's ε-terms in postponing existential commitment via symbolic placeholders. In our signed setting, however, ε's are not term-forming operators: they are informational names introduced by RB that may remain neutral (minimal informons with either $T\,A(e)$ or $F\,A(e)$ per branch) or be declared arbitrary-for-$T\,A$ / arbitrary-for-$F\,A$ subtree-locally, thereby adopting γ-like / δ-like behaviour under the corresponding polarity.*

5 A bivariate measure of depth

6 Conclusions and future work

References

[1] J.M. Crawford and D.W. Etherington. A non-deterministic semantics for tractable inference. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI '98)*, pages 286–291. AAAI Press, 1998.

[2] M. D'Agostino. Tractable depth-bounded logics and the problem of logical omniscience. In F. Montagna and H. Hosni, editors, *Probability, Uncertainty and Rationality*, CRM series, pages 245–275. Springer, 2010.

[3] M. D'Agostino. Analytic inference and the informational meaning of the logical operators. *Logique et Analyse*, 227:407–437, 2014.

[4] M. D'Agostino. An informational view of classical logic. *Theoretical Computer Science*, 606:79–97, 2015.

[5] M. D'Agostino, M. Finger, and D.M. Gabbay. Semantics and proof-theory of depth-bounded Boolean logics. *Theoretical Computer Science*, 480:43–68, 2013.

[6] M. D'Agostino and L. Floridi. The enduring scandal of deduction: Is propositional logic really uninformative? *Synthese*, 167(2):271–315, 2009.

[7] D.M. Gabbay and J. Woods. The new logic. *Logic Journal of the IGPL*, 9(2):141–174, 2001.

[8] C. Larese M. D'Agostino, D.M. Gabbay and S. Modgil. *Depth-bounded Reasoning. Volume 1: Classical Propositional Logic.* Number 102 in Studies in Logic. College Publications, London, 2024.

[9] W.V.O. Quine. *The Roots of Reference.* Open Court, 1974.

[10] Dana S. Scott. Outline of a mathematical theory of computation. In *Proceedings of the 4th Annual Princeton Conference on Information Sciences and Systems*, pages 169–176, Princeton, NJ, 1970. Princeton University. Reprinted in *Lecture Notes in Computer Science*, Vol. 188, Springer, 1985.

[11] Dana S. Scott. Domains for denotational semantics. In M. Nielsen and E. M. Schmidt, editors, *Automata, Languages and Programming*, volume 140 of *Lecture Notes in Computer Science*, pages 577–613. Springer, Berlin, Heidelberg, 1982.

Received \jreceived