L&PS – Logic & Philosophy of Science
Vol. XI, No. 1, 2013, pp. 3-57

Depth-Bounded Logic for Realistic Agents

Marcello D'Agostino
Department of Economics and Management
University of Ferrara (Italy)
marcello.dagostino@unife.it

ABSTRACT. In this paper we survey the "informational view" of classical propositional logic that has been outlined in (D'Agostino & Floridi, 2009; D'Agostino, 2010, 2013; D'Agostino *et al.*, 2013). This view is based on a kind of "informational semantics" for the logical operators and on a non-standard proof-theory. The latter is a system of classical natural deduction (Mondadori, 1989; D'Agostino, 2005) that, unlike Gentzen's and Prawitz's systems, provides natural means for measuring the "depth" of inferences in terms of the minimum number of nested applications of a single (non-eliminable) structural rule, which is an informational version of the Principle of Bivalence and is closely related to classical (analytic) cut. This leads to defining, in a natural way, hierarchies of tractable depth-bounded logical systems that indefinitely approximate Boolean logic. We argue that this approach may be apt to provide more realistic prescriptive models of resource-bounded logical agents and, at the same time, solve the most disturbing anomalies that affect the received view in classical semantics and proof-theory. We also suggest that this informational view of classical logic can partially vindicate the old Kantian idea of synthetic a priori knowledge.

KEYWORDS: Classical propositional logic, Semantic Information, Natural Deduction, Informational Semantics, Computational Complexity, Analytic/Synthetic

Published by L&PS – Logic and Philosophy of Science
http://www2.units.it/˜episteme/ – ISSN: 1826-1043

1. Introduction: what role for logic in the theory of rationality?

In his book *Minimal Rationality* (1986), Christopher Cherniak maintained that a theory of feasible and appropriate logical inferences constitutes and important component of the theory of rationality, but an *unconstrained* logic is by itself irrelevant to the psychology and epistemology of reasoners. The main problem is that standard logical systems model a *logically omniscient* agent and provide no means to account for the "cost of thinking":[1]

> Until recently, philosophy has uncritically accepted highly idealized conceptions of rationality. But cognition, computation and information have costs; they do not just subsist in some immaterial effluvium (Cherniak, 1986, p. 3).

In Cherniak's view, the primary task for a logic that aims to play a prescriptive role is to provide a theory of "a *minimal agent* who has limits on cognitive resources such as time and memory" so that, "according to such a more realistic account, an agent can have a less than perfect deductive ability" (Cherniak, 1986, p. 3).

However, logic cannot content itself with constructing models of the logical competence of minimal agents. In fact, it is uncontroversial that the capability of correctly recognizing inconsistency or logical entailment is a matter of degree. While any agent who understands the meaning of "$\to$" can recognize that $A \to B$ and A logically entail B, fewer are able to perform a logical inference involving a complex pattern of *case reasoning* and very few are able to prove a theorem from the axioms of a mathematical theory. How can we measure and empirically test the degree of difficulty or "logical depth" of an inference? To quote Cherniak again:

> Philosophy seems to have largely overlooked the point that without an extensive theory of the difficulty of different inferences, which provides

[1] This is essentially the same problem raised by Leonard Savage in his *Difficulties in the theory of personal probability*: "A person required to risk money on a remote digit of π in order to comply fully with the theory [of personal probability] would have to compute that digit, though this would really be wasteful if the cost of computation were more than the prize involved. For the postulates of the theory imply that you should behave in accordance with the logical implications of all that you know. Is it possible to improve the theory in this respect, making allowance within it for the cost of thinking, or would that entail paradox, as I am inclined to believe but unable to demonstrate?" (Savage, 1967, p. 308).

> information on which ones will be accomplished under given conditions, the predictive value of any attribution of a cognitive system of beliefs and desires would be severely limited (Cherniak, 1986, p. 29).

This naturally leads to the following

> *Approximation problem*: Can we define *in a natural way* a hierarchy of logical systems that indefinitely approximate a given ideal Logic in such a way that these approximations provide useful formal models of the logical competence of different resource-bounded agents?

Robust solutions to this problem are likely to have a significant practical impact in all research areas — from economics, to philosophy and cognitive science — where there is an urgent need for more realistic models of deduction, and require an imaginative re-examination of logical systems as they are usually presented in the literature. From this point of view, it would make sense to require that a logical system consist not only in an algorithmic or semantic characterization of a logic L, but also in a definition of how this logic L can be *approximated in practice* by realistic (not logically omniscient) agents, no matter whether human or artificial.

Despite its practical and theoretical significance, the approximation problem has been surprisingly neglected in the logical and philosophical literature.[2] A first reason is the difficulty of finding robust solutions which are independent of the choice of a specific formalism. A second reason is that the most popular human-oriented formalizations of classical logic — Natural deduction, Tableaux, Sequent Calculus — are structurally inadequate to define a measure of the difficulty of inferences that satisfy some sensible requirements.[3] A third reason is deeply entrenched in our philosophical tradition and constitutes a huge

[2] It has indeed received some attention in Computer Science and Artificial Intelligence (Cadoli & Schaerf, 1992; Dalal, 1996, 1998; Crawford & Etherington, 1998; Massacci, 1998; Sheeran & Stålmarck, 2000; Finger, 2004a,b; Finger & Wassermann, 2004, 2006; Finger & Gabbay, 2006), but comparatively little attention has been devoted to embedding such efforts in a systematic proof-theoretical and semantic framework.

[3] For example, the requirement that the meaning of the logical operators (as given by the operational rules) remains the same throughout the sequence, so that an agent may still be credited with grasping the meaning of a finite set of sentences, e.g., the axioms of a theory, even if (s)he is unable to recognize all its logical consequences.

stumbling block in the way of any sensible solution of the approximation problem.

According to the received view logic is *informationally trivial*. Deductive inference is commonly described as being "tautological", in the sense of being *uninformative* or *non-ampliative*: the information carried by the conclusion is (in some sense) contained in the information carried by the premises. This is related to the claim that the validity of an inference depends solely on the meaning of the "logical words" and convey no factual information. In this sense logic is also said to be "analytic". On the other hand, there is a tension between this view and the well-known results showing that most interesting logics are either undecidable or (very likely to be) intractable. How can logic be, at the same time, informationally trivial and computationally hard? And if it really were informationally trivial, how could we avoid logical omniscience and suitably *grade* the difficulty of deductive inference for agents with bounded cognitive resources? The fundamental question is:

(1) do we *actually* possess the information that the conclusion of a valid inference is true *whenever* we possess the information that its premises are true?

Here by "actually possessing" a piece of information we mean that it is information that is accessible to us *in practice*, and not only in principle, and with which we can *operate* as opposed to information that is only *potentially* available to us.

The lack of a general decision procedure implies that the intuitive answer is a loud "NO" in the domain of classical first-order logic. Since first-order logic admits of no decision procedure, there is no guarantee that we are always in a position to recognize the truth of a valid consequence A of a set Γ of sentences whenever we recognize the truth of the sentences in Γ.

But even restricting our attention to propositional logic, the intuitive answer is in the negative. The theory of computational complexity[4] tells us that the decision problem for Boolean logic is co-NP complete (Cook, 1971), that is, among the most difficult problems in co-NP. Even if yet unproven, the conjecture that there exists no feasible decision procedure for such problems is widely believed to be true. This means that we cannot expect a real agent, even if equipped with

[4] See (Stockmeyer, 1987) for an introduction.

an up-to-date computer running a decision procedure for Boolean logic, to be always able to recognize *in practice* that a certain conclusion follows from a given set of premises.

Things do not improve if we consider the best known subclassical propositional logics, such as intuitionistic, relevance and linear logic. As for intuitionistic logic, which *prima facie* appears to be more appropriate than classical logic to characterize logical consequence in terms of our knowledge or information states, Richard Statman (1979) proved that it is PSPACE complete (as is its pure implication fragment),[5] and so — given that NP $\subseteq$ PSPACE and that the inclusion is believed to be strict — its computational complexity is likely to be worse than that of classical logic. As for Relevance Logics, Alasdair Urquhart proved that the main systems *E* (Entailment), *R* (Relevance Logic) and *T* (Ticket Entailment) (Anderson & Belnap Jr, 1975) are all undecidable (Urquhart, 1984) and that the computational complexity of the main decidable subsystems is not better than that of Boolean logic (Urquhart, 1990). For example, the implication fragment of *R* is ESPACE complete and the fragment known as *first-degree entailment* — characterized by Belnap's 4-valued logic (Belnap Jr., 1976, 1977) — is co-NP complete. Finally, the full system of Linear Logic is undecidable (Lincoln, 1995), its multiplicative-additive fragment is PSPACE complete (Lincoln *et al.*, 1992) and the multiplicative fragment is NP complete (Kanovich, 1992).[6] Thus, logic is informationally trivial only for *ideal agents* and we cannot realistically assume that a rational, but resource-bounded agent, be informed of all the logical consequences of his or her beliefs. This strongly suggests that the conclusion of a complex inference may convey information that is *not* contained in the premises in the objective — not merely psychological — sense that there is (and probably there will never be) a feasible procedure for extracting this information from the information conveyed by the premises.

In the sequel we shall survey a new "informational view" of classical propositional logic that has been outlined in (D'Agostino & Floridi, 2009; D'Agostino,

[5] More precisely (Statman, 1979) shows that intuitionistic propositional logic can be reduced to its implication fragment and that the latter is PSPACE-hard. It then follows from (Ladner, 1977), where it is shown that S4 in in PSPACE and from the well-known polynomial translation of intuitionistic propositional logic into S4, that the decision problem for intuitionistic propositional logic is PSPACE complete.

[6] Indeed, even the constant-only fragment is NP complete (Lincoln & Winkler, 1994).

2010, 2013; D'Agostino *et al.*, 2013).[7] This view is based on a kind of "informational semantics" for the logical operators and on a non-standard proof-theory. The latter is a system of classical natural deduction (Mondadori, 1989; D'Agostino, 2005) that, unlike Gentzen's and Prawitz's systems, provides natural means for measuring the "depth" of inferences in terms of the minimum number of nested applications of a single (non-eliminable) structural rule, which is an informational version of the Principle of Bivalence and is closely related to classical (analytic) cut. We argue that this approach may be apt to provide an adequate solution to the approximation problem and, at the same time, solve the most disturbing anomalies that surround the received view in classical semantics and proof-theory. It leads to defining, in a natural way, a sequence of tractable depth-bounded deduction systems that appear to be a plausible model for representing rational agents with increasing, albeit bounded, cognitive resources. We also suggest how this "informational view" of classical logic can partially vindicate the old Kantian idea of synthetic a priori knowledge.

2. The received view and its anomalies

According to the received view purely deductive reasoning is "analytic" and, therefore, "tautological". Deductive inferences are valid solely by virtue of the meaning of the logical operators and can be recognized as such by pure conceptual analysis. Hence, the information carried by their conclusion is already contained, albeit implicitly, in their premises. In this section we give an exposition of this view and discuss its main anomalies.

2.1. A persistent dogma of empiricism

The idea that deductive reasoning conveys no new information was one of the trademarks of logical empiricism and originated in their rejection of any "synthetic a priori" knowledge:

> The scientific world-conception knows no unconditionally valid knowledge derived from pure reason, no "synthetic judgments a priori" of the

[7] Clearly there is a considerable amount of overlapping between the present exposition and these papers, but here ideas and results that were scattered through them all are combined together for the first time to provide an overall picture of our approach of and of its philosophical underpinnings.

> kind that lie at the basis of Kantian epistemology [...] It is precisely in the rejection of the possibility of synthetic knowledge a priori that the basic thesis of modern empiricism lies. The scientific world-conception knows only empirical statements about things of all kinds, and analytic statements of logic and mathematics (Hahn *et al.*, 1973, p. 308).

According to the logical empiricists, the truths of logic and mathematics are necessary and do not depend on experience. Given their rejection of any synthetic a priori knowledge, this position could be justified only by claiming that logical and mathematical statements are "analytic", i.e. true "by virtue of language". This means that their truth can be recognized, at least in principle, by means only of the meaning of the words that occur in them. Since information cannot be increased independent of experience, such analytic statements must also be "tautological", i.e., carry no information content. Hence:

> The conception of mathematics as tautological in character, which is based on the investigations of Russell and Wittgenstein, is also held by the Vienna Circle. It is to be noted that this conception is opposed not only to apriorism and intuitionism, but also to the older empiricism (for instance of J.S. Mill), which tried to derive mathematics and logic in an experimental-inductive manner as it were (Hahn *et al.*, 1973, p. 311).

The underlying idea is well represented in the following quotation from Hempel:

> It is typical of any purely logical deduction that the conclusion to which it leads simply re-asserts (a proper or improper) part of what has already been stated in the premises. Thus, to illustrate this point by a very elementary example, from the premise, "This figure is a right triangle", we can deduce the conclusion, "This figure is a triangle"; but this conclusion clearly reiterates part of the information already contained in the premise. [...] The same situation prevails in all other cases of logical deduction; and we may, therefore, say that logical deduction — which is the one and only method of mathematical proof — is a technique of conceptual analysis: it discloses what assertions are concealed in a given set of premises, and it makes us realize to what we committed ourselves in accepting those premises; but none of the results obtained by this technique ever goes by one iota beyond the information already contained in the initial assumptions (Hempel, 1945, p. 9).

This view is a persistent dogma of (logical) empiricism that has survived Quine's reservations on the very notion of analyticity (Quine, 1961) to become part of the philosophical folklore. After all, Quine's well-known arguments against the analytic/synthetic distinction spared the claim that the notion of analyticity had been sufficiently clarified in the restricted domain of logic. According to (Quine, 1961), statements that are analytic "by general philosophical acclaim" fall into two classes: those that may be called *logically true*, such as "no unmarried man is married" and those that may be turned into logical truths by replacing synonyms with synonyms, such as "no bachelor is married". Admittedly, Quine's problem is that "we lack a proper characterization of this second class of analytic statements" for, in his view, "the major difficulty lies not in the first class of analytic statements, the logical truths, but rather in the second class, which depends on the notion of synonymy" (Quine, 1951, pp. 22–23 of the 1961 edition).

Later on, in his *The Roots of Reference*, Quine clarified that the impossibility of tracing a sharp demarcation between analytic and synthetic sentences does not exclude that there may be undisputed cases of analytic sentences (typically the logical laws) (Quine, 1973, pp. 79–80) and in a 1993 interview, he made his position crystal-clear:

> Yes so, on this score I think of the truths of logic as analytic in the traditional sense of the word, that is to say true by virtue of the meaning of the words. Or as I would prefer to put it: they are learned or can be learned in the process of learning to use the words themselves, and involve nothing more (Bergström & Føllesdag, 1994, p. 199).[8]

2.2. Semantic information

Probably the idea that logical truths and inferences are "analytic" owes most of its philosophical appeal to the fact that it offers the strongest possible justification of deductive practice: logical deduction provides an infallible means of transmitting truth from the premises to the conclusion for the simple reason that the conclusion adds nothing to the information that was already contained in the premises. At the half of the 20th century Bar-Hillell and Carnap's notion of "semantic information" (Bar-Hillel & Carnap, 1953) closed the circle by providing a precise mathematical characterization of the "information content" of a sentence according to which all deductive inferences turn out to be "tautological", in the sense of being uninformative.

[8] Quoted in (Decock, 2006).

Although their effort was clearly inspired by the rising enthusiasm for Shannon and Weaver's new mathematical theory of information (Shannon & Weaver, 1949), their starting point was their dissatisfaction with the nonchalant tendency of fellows scientists to apply its concepts and results well beyond the "warranted areas". Shannon and Weaver's central problem was only how uninterpreted data can be efficiently encoded and transmitted. So the idea of applying their theory to contexts in which the *interpretation* of data plays an essential rôle was a major source of confusion and misunderstandings:

> The Mathematical Theory of Communication, often referred to also as Theory of (Transmission of) Information", as practised nowadays, is not interested in the content of the symbols whose information it measures. The measures, as defined, for instance, by Shannon, have nothing to do with what these symbols symbolyse, but only with the frequency of their occurrence. [...] This deliberate restriction of the scope of the Statistical Communication Theory was of great heuristic value and enabled this theory to reach important results in a short time. Unfortunately, however, it often turned out that impatient scientists in various fields applied the terminology and the theorems of Communication Theory to fields in which the term "information" was used, presystematically, in a semantic sense, that is, one involving contents or designata of symbols, or even in a pragmatic sense, that is, one involving the users of these symbols (Bar-Hillel & Carnap, 1953, p. 147).

By contrast, they put forward a theory of *semantic* information, in which the contents of symbols were "decisively involved in the definition of the basic concepts" and "an application of these concepts and of the theorems concerning them to fields involving semantics thereby warranted" (Bar-Hillel & Carnap, 1953, p. 148). The basic idea is simple and can be briefly explained as follows.

Suppose we are interested in the weather forecast for tomorrow and that we focus only on the possible truth values of the two sentences "tomorrow will rain" (R) and "tomorrow will be windy" (W). Then, there are four possible relevant states of the world, described by the following conjunctions:

$$R \wedge W \qquad R \wedge \neg W \qquad \neg R \wedge W \qquad \neg R \wedge \neg W.$$

Now, the sentence "tomorrow it will rain and will be windy" is intuitively more informative than the sentence "tomorrow it will rain". We can explain this by noticing that it excludes more possibilities, i.e, more possible (relevant) states of the world. On the other hand, the sentence "tomorrow it will rain or will not rain" conveys no information, since it does not exclude any possible state. So, it seems natural to identify the information conveyed by a sentence with the set of all "possible worlds" that are excluded by it, and to assume that its measure should be somehow related to the size of this set.

The same basic idea, identifying the information carried by a sentence with the set of the possible states that it excludes, had already made its appearance in Popper's *Logic of Scientific Discovery* (1934), where it played a crucial rôle in defining the "empirical content" of a theory and in supporting Popper's central claim, namely that the most interesting scientific theories are those that are highly falsifiable, while unfalsifiable theories are devoid of any empirical content:

> The amount of positive information about the world which is conveyed by a scientific statement is the greater the more likely it is to clash, because of its logical character, with possible singular statements. (Not for nothing do we call the laws of nature "laws": the more they prohibit the more they say.) (Popper, 1959, p. 19). [...]
>
> It might then be said, further, that if the class of potential falsifiers of one theory is "larger" than that of another, there will be more opportunities for the first theory to be refuted by experience; thus compared with the second theory, the first theory may be said to be "falsifiable in a higher degree". This also means that the first theory says more about the world of experience than the second theory, for it rules out a larger class of basic statements. [...] Thus it can be said that the amount of empirical information conveyed by a theory, or its empirical content, increases with its degree of falsifiability. (Popper, 1959, p. 96).

The theory of semantic information so provided what is, to date, the strongest justification for the thesis that deductive reasoning is "tautological". Indeed, an inevitable consequence of this theory is that *all* logical truths are equally uninformative (they exclude no possible world). But in classical logic a sentence B is deducible from a finite set of premises $A_1,\ldots,A_n$ if and only if the conditional $(A_1 \wedge \ldots \wedge A_n) \rightarrow B$ is a tautology. Accordingly, since tautologies carry no information at all, no logical inference can yield an increase of information. Therefore, if we identify the semantic information carried by a sentence with the set of all possible worlds it excludes, we must also accept the inevitable consequence that, in any valid deduction, the information carried by the conclusion is contained in the information carried by the (conjunction of) the premises. While this theory seems to justify the empiricist dogma that logic is "analytic", it appears to be at odds with our intuitions and clash with the commonsense notion of information. As Michael Dummett put it:

> Once the justification of deductive inference is perceived as philosophically problematic at all, the temptation to which most philosophers succumb is to offer too strong a justification: to say, for instance, that when we recognize the premises of a valid inference as true, we have thereby already recognized the truth of the conclusion (Dummett, 1991, p. 195).

In fact, such a definitive foundation for deductive practice is obtained at the price of its *trivialization*. Logic lies on a bedrock of platitude.

2.3. The enduring scandal of deduction

Cohen and Nagel were among the first to point out that the traditional tenet that logical deduction is devoid of any informational content sounds paradoxical:

> If in an inference the conclusion is not contained in the premises, it cannot be valid; and if the conclusion is not different from the premises, it is useless; but the conclusion cannot be contained in the premises and also possess novelty; hence inferences cannot be both valid and useful (Cohen & Nagel, 1934, p. 173)

A few decades later Jaakko Hintikka described this paradox as a true "scandal of deduction":

> C.D. Broad has called the unsolved problems concerning induction a scandal of philosophy. It seems to me that in addition to this scandal of induction there is an equally disquieting scandal of deduction. Its urgency can be brought home to each of us by any clever freshman who asks, upon being told that deductive reasoning is "tautological" or "analytical" and that logical truths have no "empirical content" and cannot be used to make "factual assertions": in what other sense, then, does deductive reasoning give us new information? Is it not perfectly obvious there is some such sense, for what point would there otherwise be to logic and mathematics? (Hintikka, 1973, p. 222).

For Ludwig Wittgenstein, the whole problem was, needless to say, a pseudo-problem, arising from our use of an imperfect language.[9] In his *Tractatus*,

[9] This topic is discussed in detail in (Carapezza & D'Agostino, 2010).

Wittgenstein raises the question of an "adequate notation" through which each sentence *shows* its meaning, where the latter is to be identified with the possibility of its being true or false: "The sense of a proposition is its agreement and disagreement with the possibilities of the existence and non-existence of the atomic facts." (T. 4.2). While the truth of an elementary propositions consists in the existence or non-existence of a certain fact about the world, the truth of complex propositions depends on the logical relations between the elementary propositions occurring in them: complex propositions are truth functions of the elementary propositions. Thus, the meaning of a proposition consists in the conditions under which it is true or false, and an adequate notation should be able to *show* these conditions explicitly: "a proposition shows its sense" (T. 4.022). Nevertheless, "[in common language] it is humanly impossible to deduce the logic of language" (T. 4.002), because the grammatical structure does not mirror the logical structure of the sentence itself. The logic underlying linguistic utterances could instead be made evident by a more appropriate symbolism, one capable of making it immediately visible without resorting to any "deductive process".

In a logically perfect language the recognition of tautologies should be immediate. Since the deducibility of a certain conclusion from a given set of premises is equivalent to the tautologyhood of the conditional whose antecedent is the conjunction of the premises and whose consequent is the conclusion of the inference, then the correctness of any inference would prove, in a symbolism of the kind, to be immediately visible. So, given a "suitable notation", logical deduction could actually be reduced to the mere *inspection* of propositions:

> When the truth of one proposition follows from the truth of others, we can see this from the structure of the propositions. [*Tractatus*, 5.13]
>
> In a suitable notation we can in fact recognize the formal properties of propositions by mere inspection of the propositions themselves. [6.122].
>
> Every tautology itself shows that it is a tautology. [6.127(b)]

In accordance with Wittgenstein's idea, one could specify a procedure that translates sentences into a "perfect notation" that fully brings out the information they convey, for instance by computing the whole truth-table for the conditional which represents the inference. Such a table displays all the relevant possible worlds and allows one to distinguish immediately those that make a sentence true from those that make it false, the latter representing (collectively) the "semantic information" carried by the sentence. Once the translation has been performed, logical consequence can be recognized by "mere inspection"

Thus, if information could be fully unfolded by means of some mechanical translation into a "perfect logical language", the scandal of deduction could be avoided without appealing to psychologism. Sometimes we fail to immediately "see" that a conclusion is implicit in the premises because we express both in a concise notation, a sort of stenography that prevents us from fully recognizing the formal properties of propositions until we decode it into an adequate notation. From this point of view, semantic information would be a perfectly good way of specifying the information carried by a sentence with reference to an algorithmic procedure of translation.

Although this idea may seem to work well for propositional logic, one can easily see how the Church-Turing undecidability theorem excludes the possibility of a perfect language, in Wittgenstein's sense, for first-order logic: since first-order logical truth is undecidable, we can never find an algorithm to translate every sentence into a perfect language in which its tautologyhood could be immediately decided by mere inspection. This negative result is also the main motivation for Hintikka's criticism of Bar-Hillell and Carnap's notion of semantic information.

> [...] measures of information which are not effectively calculable are well-nigh absurd. What realistic use can there be for measures of information which are such that we in principle cannot always know (and cannot have a method of finding out) how much information we possess? One of the purposes the concept of information is calculated to serve is surely to enable us to review what we know (have information about) and what we do not know. Such a review is in principle impossible, however, if our measures of information are non-recursive (Hintikka, 1973, p. 228).

Hintikka's positive proposal consists in distinguishing between two *objective* and *non-psychological* notions of information content: "surface information", which may be increased by deductive reasoning, and "depth information" (equivalent to Bar-Hillel and Carnap's "semantic information"), which may not. While the latter justifies the traditional claim that logical reasoning is tautological, the former vindicates the intuition underlying the opposite claim. In his view, *first-order* deductive reasoning may increase surface information, although it never increases depth information (the increase being related to deductive steps that introduce new individuals). Without going into details,[10] we observe here that Hintikka's proposal classifies as non-analytic only some inferences of the non-monadic predicate calculus so leaves the "scandal of deduction" unsettled in the domain of propositional logic:

> The truths of propositional logic are [...] tautologies, they do not carry any new information. Similarly, it is easily seen that in the logically valid inferences of propositional logic the information carried by the conclusion is smaller or at most equal to the information carried by the premisses. The term "tautology" thus characterizes very aptly the truths and inferences of propositional logic. One reason for its one-time appeal to philosophers was undoubtedly its success in this limited area (Hintikka, 1973, p. 154).

Hence, in Hintikka's view, for every finite set of Boolean sentences Γ and every Boolean sentence A,

(2) If $\Gamma \vdash A$ the information carried by A is included in the information carried by Γ.

As argued in the introduction, this is highly unsatisfactory given the likely intractability of Boolean logic. Thus, some degree of uncertainty about whether or not a certain conclusion follows from given premises cannot be, in general, completely eliminated even in the restricted and "simple" domain of propositional logic. So, if we take seriously the time-honoured and common-sense concept of information, according to which information consists in relieving us from uncertainty, we should conclude that in some cases learning that a certain conclusion logically follows from the premises *does* relieve us from uncertainty, and therefore increases our information, even at the propositional level. A well-known example is the solution of an "expert level" sudoku, which is surely informative for ordinary solvers even if it follows from the initial information by propositional logic only. The scandal of deduction has recently received renewed attention leading to a number of original contributions (e.g., (Primiero, 2008, Ch. 2), (Sequoiah-Grayson, 2008), (Sillari, 2008b), (D'Agostino & Floridi, 2009), (Duží, 2010), (Jago, 2012) that do not appear, however, to be reducible to a single conceptual paradigm.

[10] For a criticism of Hintikka's approach see (Sequoiah-Grayson, 2008).

2.4. The BHC paradox

Another straightforward consequence of Bar-Hillel and Carnap's theory of "semantic information" is that contradictions, like "tomorrow it will rain and it will not rain", carry the maximum amount of information, since they exclude all possible states. Bar-Hillel and Carnap were well aware that their theory of semantic information sounded counterintuitive in connection with contradictory (sets of) sentences, as shown by the near-apologetic remark they included in their (1953):

> It might perhaps, at first, seem strange that a self-contradictory sentence, hence one which no ideal receiver would accept, is regarded as carrying with it the most inclusive information. It should, however, be emphasized that semantic information is here not meant as implying truth. A false sentence which happens to say much is thereby highly informative in our sense. Whether the information it carries is true or false, scientifically valuable or not, and so forth, does not concern us. A self-contradictory sentence asserts too much; it is too informative to be true (Carnap & Bar-Hillel, 1953, p. 229).

Popper had also realized that his closely related notion of empirical content worked reasonably well only for *consistent* theories. For, *all* basic statements are potential falsifiers of all inconsistent theories, which would therefore, without this requirement, turn out to be the most scientific of all. So, for him, "the requirement of consistency plays a special rôle among the various requirements which a theoretical system, or an axiomatic system, must satisfy" and "can be regarded as the first of the requirements to be satisfied by every theoretical system, be it empirical or non-empirical" (Popper, 1959, p. 72). So, "whilst tautologies, purely existential statements and other unfalsifiable statements assert, as it were, *too little* about the class of possible basic statements, self-contradictory statements assert *too much*. From a self-contradictory statement, any statement whatsoever can be validly deduced" (Popper, 1959, p. 71). In fact, what Popper claimed was that the information content of inconsistent theories is null, and so his definition of information content as monotonically related to the set of potential falsifiers was intended only for consistent ones:

> But the importance of the requirement of consistency will be appreciated if one realizes that a self-contradictory system is uninformative. It is so because any conclusion we please can be derived from it. Thus no

> statement is singled out, either as incompatible or as derivable, since all are derivable. A consistent system, on the other hand, divides the set of all possible statements into two: those which it contradicts and those with which it is compatible. (Among the latter are the conclusions which can be derived from it.) This is why consistency is the most general requirement for a system, whether empirical or non-empirical, if it is to be of any use at all (Popper, 1959, p. 72).

2.5. The problem of logical omniscience in epistemic logic

The anomalies of the received view on classical logic carry over to the widely studied logics of knowledge (epistemic logic) and belief (doxastic logic), as well as to the more recent attempts to axiomatize the "logic of being informed" (information logic).[11] If an agent a knows (or believes, or is informed) that a sentence A is true, and B is a logical consequence of A, then a is supposed to know (or believe, or be informed) also that B is true. This is often described as paradoxical and labelled as "the problem of logical omniscience". Let $\Box_a$ express any of the propositional attitudes at issue, referred to the agent a. Then, the "logical omniscience" assumption can be expressed by saying that, for any finite set Γ of sentences,

(3) $$\text{if } \Box_a A \text{ for all } A \in \Gamma \text{ and } \Gamma \vdash B, \text{ then } \Box_a B,$$

where $\vdash$ stands for the relation of logical consequence. Observe that, letting $\Gamma = \emptyset$, it immediately follows from (3) that any rational agent a is supposed to be aware of the truth of all classical tautologies, that is, of all the sentences of a standard logical language that are "consequences of the empty set of assumptions". In most axiomatic systems of epistemic, doxastic and information logic assumption (3) emerges from the combined effect of the "distribution axiom", namely

(K) $\Box_a(A \to B) \to (\Box_a A \to \Box_a B)$

and the "necessitation rule":

(N) if $\vdash A$, then $\vdash \Box_a A$.

[11] For a survey on epistemic and doxastic logic see (Halpern, 1995; Meyer, 2003). For information logic, or "the logic of being informed", see (Floridi, 2006; Primiero, 2009).

On the other hand, despite its paradoxical flavour, (3) seems an inescapable consequence of the standard Kripke-style semantical characterization of the logics under consideration. The latter is carried out in terms of structures of the form $(S, \tau, R_1, \ldots, R_n)$, where S is a set of possible worlds, τ is a function that associates with each possible world s an assignment $\tau(s)$ of one of the two truth values (0 and 1) to *each* atomic sentence of the language, and each R_a is the "accessibility" relation for the agent a. Intuitively, if s is the actual world and sR_at, then t is a world that a would regard as a "possible" alternative to the actual one, i.e., compatible with what a knows (or believes, or is informed of). Then, the truth of complex sentences is defined, starting from the initial assignment τ, via a forcing relation $\vDash$. This incorporates the usual semantics of classical propositional logic and defines the truth of $\Box_a A$ as "A is true in all the worlds that a regards as possible". In this framework, given that the notion of truth in a possible world is an extension to the modal language of the classical truth-conditional semantics for the standard logical operators, (3) appears to be both compelling and, at the same time, counter-intuitive.

Now, under this reading of the consequence relation $\vdash$, which is based on classical propositional logic, (3) may perhaps be satisfied by an "idealized reasoner", in some sense to be made more precise,[12] but is not satisfied, and is not likely to ever be satisfiable, in practice. As mentioned in the introduction, even restricting ourselves to the domain of propositional logic, the theory of computational complexity tells us that the decision problem for Boolean logic is co-NP-complete. So, the clash between (3) and the classical notion of logical consequence, which arises in any real application context, may only be solved either by waiving the assumption stated in (3), or by waiving the consequence relation of classical logic in favour of a weaker one with respect to which it may be safely assumed that the modality $\Box_a$ is closed under logical consequence for any realistic agent.

[12] It should be noted that the appeal to an "idealized reasoner" has usually the effect of sweeping under the rug a good deal of interesting questions, including how idealized such a reasoner should be. Idealization may well be a matter of degree.

Both options have been discussed in the literature.[13] Observe that, according to the latter, the problem of logical omniscience does not lie in assumption (3) in itself, but rather in the standard (classical) characterization of logical consequence for a propositional language that is built in the possible-world semantics originally put forward by Jaakko Hintikka as a foundational framework for the investigation of epistemic and doxastic logic.

Frisch (1987) and Levesque (1988) were among the first authors to explore this route and argue for a notion of "limited inference" based on "a less idealized view of logic, one that takes very seriously the idea that certain computational tasks are relatively easy, and others more difficult" (Levesque, 1988), p. 355. A more recent (and related) proposal can be found in (Fagin *et al.*, 1995), where the authors suggest to replace classical logic with a non-standard one, deeply rooted in relevance logic and called NPL (for "Nonstandard Propositional Logic"), to mitigate the problem of logical omniscience. The mitigation consists mainly in the existence of a polynomial time decision procedure for the CNF fragment of the proposed logical system ((Fagin *et al.*, 1995, Theorem 7.4)). However, the decision problem for the unrestricted language of NPL is still co-NP-complete (Theorem 6.4). Moreover, NPL shares with relevance logic and with Levesque's notion of limited inference the invalidity of disjunctive syllogism (from $A \vee B$ and $\neg A$ one cannot infer B) which sounds disturbing to most classical ears. Finally, the NPL-based approach does not allow, in a natural way, for the possibility of defining *degrees* of logical omniscience, that may apply to *increasingly idealized* reasoning agents, in terms of correspondingly stronger consequence relations. On the other hand, the possibility of characterizing in a uniform way such a hierarchy of approximations to the "perfect reasoner" (which may well be a classical one) would certainly allow for all the flexibility needed by a suitable model of practical rationality.

---

[13] See (Meyer, 2003) (Section 4), (Halpern, 1995) (Section 4) and (Kwang, 1997) for a survey and proper references. See also: (Parikh, 2008) for an interesting third view that draws on the tradition of subjective probability, and (Artemov & Kuznets, 2009) for an approach based on proof size. A general semantic framework in which several different approaches can be usefully expressed is that based on "awareness structures", which draws on the distinction between "explicit" and "implicit" knowledge, to the effect that an agent may implicitly know that a sentence is a logical consequence of a set of assumptions, without being *aware* of it. See (Sillari, 2008a,b) for an insightful discussion of this framework and proper references to the literature.

In the next section we shall outline a sort of "informational view" of classical propositional logic that provides a natural solution to the the approximation problem raised in the introduction as well as paving the way for a solution of the anomalies of the received view discussed in this section.

3. An informational view of classical logic

In this section we outline the basic principles of an informational semantics for the logical operators. We survey two kinds of such semantics — constraint-based and modular — introduced, respectively, in (D'Agostino & Floridi, 2009) and (D'Agostino *et al.*, 2013), which admit of natural proof-theoretical characterizations. The deductive arguments that can be justified on the basis of this informational semantics are a specific subset of the classical valid arguments, namely those that no make use of any "virtual information", that is do not simulate information that is not even implicitly contained in the data. However, full Boolean logic can be indefinitely approximated by allowing the use of virtual information up to a given fixed depth. Finally we suggest how this informational approach can solve the anomalies of the received view.

3.1. Informational semantics

What is a "sensible" semantics for the logical operators? Classical logicians have a straightforward answer to this question: the time-honoured semantics based on the classical truth-tables that fix the meaning of each logical operator $\sharp$ by fixing the conditions under which a sentence containing $\sharp$ as main logical operator is true or false in terms of the truth or falsity of its immediate constituents. Such conditions provide an explanation of the meaning of the logical operators in terms of the two central notions of truth and falsity, which are assumed as understood. It is regarded as essential to the understanding of these notions that they obey the classical principles of *bivalence* (each sentence, in a given state of affairs, is either determinately true or determinately false) and *non-contradiction* (no sentence can be at the same time true and false in the same state of affairs). Both principles can be concisely expressed by assuming that a sentence is false if and only if it is not true.

This approach has been severely criticized as too "metaphysical". As Weir puts it:

> [Classical semantics] has come under a great deal of attack, especially from those who subscribe to the Wittgensteinian slogan that meaning is use and interpret it as requiring that all ingredients of meaning can be made manifest in our use of sentences, especially in teaching or communicating their senses, for it is often claimed that classical bivalent semantics, in ascribing truth-values to sentences regardless of whether these values are discoverable, violates this requirement (Weir, 1986, p. 459).

A less "metaphysical" approach might consist in replacing classical truth and falsity, as central notions of the theory of meaning, with other, more accessible ones. The standard answer is intuitionistic logic, where the meaning of the logical operators may indeed be explained (with some difficulty) in terms of a notion of truth as provability or verifiability that is not recognition-transcendent.[14] Indeed, the well-known Kripke semantics defines logical consequence as truth-preserving over states that Kripke himself intuitively described as "points in time (or 'evidential situations'), at which we may have various pieces of information" (Kripke, 1965, p. 100). However, given the complexity issues discussed in the introduction, intuitionistic validity cannot, in general, be recognized in practice by any realistic agent who is in a given partial information state. A problematic feature is that the truth of some complex sentences at an information state s cannot be established without "visiting" information states that are *essentially richer* than s. For example, in order to recognize that a conditional $A \to B$ is true at a state s in which A is not true, a reasoning agent must ideally transfer from s to a "virtual" state s^* in which the antecedent A is true and any other sentence has the same value as in s; that is, the agent reasons *as if* his state were s^*, observes that in s^* the consequent B must be true as well, and concludes that $A \to B$ must be true in his real information state s.

This use of "virtual information" is part of our common reasoning practice and is not too problematic as long as the structure of the sentence whose truth is being evaluated keeps simple. However, when recognizing the sentence as true requires weaving in and out of a complex recursive pattern of virtual information states, the situation may soon get out of control, as shown by the fact that the decision problem for the pure implicational fragment of intuitionistic logic is also PSPACE complete (Statman, 1979; Svejdar, 2003). The neces-

[14] These issues, and all the subtleties that they involve, have been thoroughly discussed in the logical literature, especially in the writings of Michael Dummett; the reader is referred to (Dummett, 1991) for an overall picture.

sity of venturing out of one's actual information state in order to recognize the truth of certain sentences is what makes such inference steps "non-analytic" in a sense very close to Kant's original sense:[15] we essentially need to go *beyond the data*, using "virtual information", i.e., simulating situations in which we hold information that we do not actually hold. Although all virtual information is eventually removed, to the effect that the truth of a the conditional sentence depends only on the information initially available, it remains true that such inference steps could not be performed at all without (temporarily) trespassing on richer information states.

A second problematic feature is the treatment of disjunction. In Kripke semantics a disjunction $A \vee B$ is true at an information state s if and only if either A is true at s or B is true at s. This reflects the intuitionistic notion of truth as (conclusive) verification, more precisely, the idea that the truth of a sentence coincides with the existence of a canonical proof for it, that is, a proof obtained "by the most direct means". In a natural deduction system this is a proof whose last step is the application of an introduction rule.[16] Indeed, in intuitionistic terms, we have a canonical proof of $A \vee B$ if and only if we have either a canonical proof of A or a canonical proof of B. However this does not seem to be a compelling feature of our understanding of $\vee$ in relation to a more ordinary notion of "information state", in which the truth of a sentence may be licensed by some weaker kind of epistemic condition.

It is not difficult to come up with intuitive examples in which we hold enough information to assert a disjunction as true, but we do not hold enough information to assert either of the two disjuncts as true. Suppose we put two bills of 50 and 100 euros in two separate envelopes and then we shuffle the envelopes so as to loose track of which contains which. If we pick up one of them, we certainly hold the information that it contains either a 50-euro bill or a 100-euro bill, but we do not hold the information that it contains a 50-euro bill, nor do we hold the information that it contains a 100-euro bill.[17]

These difficulties have been addressed in (D'Agostino & Floridi, 2009) and

[15] See (D'Agostino & Floridi, 2009) on this point.

[16] See (Dummett, 1991) (Chapter 11) and (Prawitz, 2006) for a thorough discussion.

[17] This example is particularly tricky in that we could claim that we have, in some sense, arrived at the disjunction in a canonical way, except that the information has decayed during the process of shuffling the envelopes.

(D'Agostino, 2013) on the basis of an *informational semantics* for classical propositional logic in which the meaning of the Boolean operators is fixed exclusively in terms of the information that we *actually possess*. The primary notions of this semantics are not truth and falsity, but *informational truth* and *informational falsity*, namely *holding the information* that a sentence is true, respectively false. Here, by saying that an agent x *holds* the information that A is true (respectively false) we mean that this is information that is *practically available* to x and with which x can operate. Clearly these notions do not obey the classical Principle of Bivalence (or, better, its informational version). We cannot assume that for every sentence A either we hold the information that A is true or we hold the information that A is false. However, we may assume that they satisfy the informational version of the Principle of Non-Contradiction: no agent can *actually* possess both the information that A is true and the information that A is false, as this would be deemed to be equivalent to possessing no definite information about A.

We use the values 1 and 0 to denote, respectively, informational truth and falsity. When a sentence takes neither of these two defined values, we say that it is *informationally indeterminate*.[18] We call *partial valuation* any partial mapping from the set of well-formed formulae of a standard propositional language $\mathscr{L}$ to the the set $\{1,0\}$ of the two determinate informational values. We write $v(A) = \bot$ as shorthand for "$v(A)$ is undefined", that is A is informationally indeterminate.[19]

Partial valuations can be interpreted in a variety of ways. According to one interpretation $v(A) = 1$ means intuitively that A is true, $v(A) = 0$ that A is false and $v(A) = \bot$ that A is neither-true-nor-false. According to another, which is the one we shall adopt in this paper, a partial valuation v represents the information held by a given agent a about the truth or falsity of sentences; $v(A) = 1$ means "a holds the information that A is true" , $v(A) = 0$ means "a holds the information that A is false" and $v(A) = \bot$ means "a holds no information about the truth or falsity of A".[20]

---

[18] This is the symbol for "undefined", the bottom element of the information ordering, not to be confused with the "falsum" logical constant.

[19] Here $\bot$ is the symbol for "undefined", the bottom element of the information ordering, not to be confused with the "falsum" logical constant.

[20] For other interpretations of partial valuations, see (Blamey, 1986).

3.2. Constraint-based semantics

In valuation-based approaches, the intended meaning of the logical operators is usually specified by defining, within the set of all possible valuations, those which are *admissible*, i.e. those that comply with this intended meaning. Admissible valuations are usually defined by specifying a set of necessary and sufficient conditions that a valuation should satisfy. The usual conditions for the Boolean operators are the following:

C1 $v(A) = 1$ if and only if $v(\neg A) = 0$;

C2 $v(A \wedge B) = 1$ if and only if $v(A) = 1$ and $v(B) = 1$;

C3 $v(A \vee B) = 1$ if and only if $v(A) = 1$ or $v(B) = 1$;

C4 $v(A \to B) = 1$ if and only if $v(A) = 0$ or $v(B) = 1$.

C5 $v(A) = 0$ if and only if $v(\neg A) = 1$;

C6 $v(A \wedge B) = 0$ if and only if $v(A) = 0$ or $v(B) = 0$;

C7 $v(A \vee B) = 0$ if and only if $v(A) = 0$ and $v(B) = 0$;

C8 $v(A \to B) = 0$ if and only if $v(A) = 1$ and $v(B) = 0$.

A valuation satisfying the above conditions is said to be *saturated*. More specifically, we say that a valuation v is *upward saturated*, if v satisfies the above conditions in the "only-if" direction, and *downward saturated* if it satisfies them in the "if" direction. A *Boolean valuation* is a saturated valuation that satisfies the additional condition of being *total*, i.e. defined for all sentences. Observe that, for total valuations, conditions C5–C8 are redundant, in that they can be derived from conditions C1–C4.

According to the standard view, the intended meaning of the classical logical operators is fixed by accepting only Boolean valuations as admissible. Moving from their classical to their informational meaning, not only must the requirement of total valuations be dropped, but also some of the saturation conditions become obviously unsound. In particular, when represented as valuations, information states *are not downward saturated*. Indeed, as explained above, it may well be the case, under the ordinary notion of information, that we hold the information that a disjunction is true or the information that a conjunction is false without holding any information about the component sentences. Therefore, in general, when both A and B are informationally indeterminate, the value of their conjunction $A \wedge B$ may be either informational falsity 0, or informational indeterminacy $\perp$, depending on whether or not we hold the information that A and B cannot be simultaneously true. And the value of their disjunction $A \vee B$ may be either informational truth 1 or informational indeterminacy $\perp$, depending on whether or not we hold the information that at least one of A and B must be true.[21] In this context, admissible valuations *cannot be specified by means of necessary and sufficient conditions* such as C1–C8. All we can do is specify a set of *negative constraints*, which restrict the domain of all possible valuations to those which are compliant with the intended (informational) meaning of the logical operators. In (D'Agostino & Floridi, 2009) the meaning of the logical operators is characterized in terms of such negative constraints on the admissible partial valuations for $\mathscr{L}$.

Let us call *$\mathscr{L}$-module* any set consisting of a non-atomic $\mathscr{L}$-formula, called the *top formula* of the module, and of its *immediate* subformulae. We shall denote by $\mathrm{Mod}(A)$ the unique $\mathscr{L}$-module whose top formula is A. The informational meaning of a logical operator $\sharp$ can be fixed by determining which subvaluations of $\mathrm{Mod}(A)$ are *not admissible* for a formula A containing $\sharp$ as the main operator. This is a negative way of defining this meaning. It allows us to detect valuations that are *immediately* forbidden to any agent who "understands" it. For example, a valuation such that $v(A \vee B) = 1, v(A) = 0$ and $v(B) = 0$ would clearly be inadmissible and therefore provides a negative constraint that can be taken as part of the definition of "$\vee$".

These negative constraints are summarized in Table 1, where each line represents a *minimal* non-admissible valuation (the asterisk means that the corresponding informational value of the sentence may indifferently be true, false or indeterminate). A valuation v is *admissible* if, for every formula A, v does not contain any subvaluation of $\mathrm{Mod}(A)$ that is ruled out by the accepted constraints expressing the informational meaning of the main operator of A. We shall denote

[21] As far as the operator $\vee$ is concerned, its informational meaning we are trying to characterize clearly departs from its intuitionistic meaning, according to which a disjunction $A \vee B$ is intuitionistically true (roughly speaking, provable[22]) if and only if either A is intuitionistically true or B is intuitionistically true. This is the so-called disjunction property of intuitionistic logic. While this property is appropriate for (constructive) mathematics, it is quite at odds with ordinary usage outside mathematics.

$\neg A$	A
1	1
0	0

$A \vee B$	A	B
1	0	0
0	1	$*$
0	$*$	1

$A \wedge B$	A	B
1	0	$*$
1	$*$	0
0	1	1

$A \to B$	A	B
1	1	0
0	$*$	1
0	0	$*$

TABLE 1: *The informational meaning of the logical operators. Each line represents a minimal non-admissible valuation.*

by $\mathscr{A}$ the domain of all admissible valuations. Admissible valuations are partially ordered by the usual approximation relation $\sqsubseteq$ defined as follows: $v \sqsubseteq w$ (read "w is a *refinement* of v" or "v is an *approximation* of w") if and only if w agrees with v on all the formulae A for which $v(A) \neq \bot$. We can identify a partial valuation v with a set of pairs of the form $\langle A, i\rangle$, where A is a sentence in the given language and i is equal to 0 or 1, such that for no A, $\langle A, 1\rangle$ and $\langle A, 0\rangle$ are both in v. Each of these pairs can be thought of as a "piece of information" and the partial valuation itself as an attempt to put together such pieces of information in a way which is consistent with the intended meaning of the logical operators. Under this interpretation $v(A) = 1$ means that $\langle A, 1\rangle \in v$, $v(A) = 0$ that $\langle A, 0\rangle \in v$ and $v(A) = \bot$ that neither of these two pieces of information is in v. The partial ordering $\sqsubseteq$ is a meet-semilattice with a bottom element equal to $\emptyset$, the valuation which is undefined for all formulae of the language. It fails to be a lattice because the join of two admissible valuations may be inadmissible.

Now, consider a valuation v such that $v(A \vee B) = 1$ and $v(A) = 0$, while B is undefined. We can legitimately say that the value of B in this valuation is implicitly determined by the values of $A \vee B$ and A and by our understanding of the meaning of $\vee$ based on the constraints specified in Table 1. For, there is *no admissible refinement* of v such that $v(B) = 0$, since such a refinement would fail to satisfy one of the constraints that define the meaning of "$\vee$". In other words, any assignment other than 1 would be *immediately* recognized as inconsistent by any agent that understands $\vee$ via the specified constraints. Notice that checking whether a possible refinement of a valuation regarding a specific formula B is admissible is a task that can be performed in linear time on the basis of local information. It involves only checking all the modules that contain B. If assigning a certain value to B violates one of the constraints, which are specified in terms of the modules containing B, then the valuation is not admissible.

Hence, if we understand the meaning of $\vee$, we are able to detect immediately that B cannot be assigned, consistently with this meaning, the value 0, given that $A \vee B$ has been assigned the value 1 and A the value 0. In such a situation, we can say that the information concerning the value of B is *implicitly contained* in our information state — in particular, in that portion of it concerning $\mathrm{Mod}(A \vee B)$, i.e. $A \vee B$ and its immediate subformulae — because the intended meaning of $\vee$ leaves us no option about this value. This is comparable, interestingly, to what happens in those easy steps of the *sudoku* game where the digit to be inserted in a given empty cell is *dictated* by the digits already inserted in the cells belonging to the regions into which the empty cell is contained. (Recall that a "region" in the classic version of the *sudoku* game is either a column, or a row or one of the four sub-squares into which the main square is divided. In our context, a "region" is simply the locale of a formula.) This is the most basic consistency principle by means of which logical inference can be justified analytically, that is, by virtue of the informational meaning of the logical operators as specified by the constraints in Table 1.

Given a valuation v, let us say that a piece of information $\langle A, i\rangle$, with $i \in \{0,1\}$, is *implicitly contained in v at depth* 0, and write $v \Vdash_0 \langle A, i\rangle$, if the complementary piece of information $\langle A, |1-i|\rangle$ is *immediately ruled out* solely by virtue of the meaning of the logical operators, that is, if extending v with $\langle A, |i-1|\rangle$ makes it non-admissible by violating one of the constraints specifying the intended meaning of the main logical operator of A. In symbols (recalling that $\mathscr{A}$ is the set of admissible valuations):

$$v \Vdash_0 \langle A, i\rangle \Longleftrightarrow v \cup \{\langle A, i\rangle\} \notin \mathscr{A}. \tag{4}$$

Notice that, according to the above definition, if $\langle A, i\rangle \in v$, then $v \Vdash_0 \langle A, i\rangle$, since $v \cup \langle A, |i-1|\rangle$ is trivially inadmissible in that it is not a partial valuation. A *minimal* requirement on an information state is that it is closed under the implicit information that immediately stems from the meaning constraints. So, an admissible valuation v is a *shallow information state* if it satisfies:

$$v \Vdash_0 \langle A, i\rangle \Longrightarrow \langle A, i\rangle \in v \tag{5}$$

A shallow information state is a Boolean valuation if and only if it contains a *total* valuation of all the atomic formulae of the language. So, Boolean valuations can be seen as shallow information states that are closed under a *Principle of Omniscience*, the informational counterpart of the classical Principle of Bivalence:

> For every information state v and every atomic sentence p, either p is true in v or p is false in v.

Example 3.1. Consider an admissible valuation v such that

1. $v(p \vee q) = 1$
2. $v(p) = 0$
3. $v(q \to r) = 1$
4. $v(q \to s) = 1$
5. $v(\neg t \to \neg(r \wedge s)) = 1$
6. $v(t \wedge u) = 0$.

We show that $v'(u) = 0$ for every shallow information state that contains v, i.e. for every admissible valuation closed under D_0 which contains v. From (i) and (ii), by the constraints on $\vee$, it follows that $v' \cup \{\langle q, 0\rangle\}$ would be non-admissible. So, $v' \Vdash_0 \langle q, 1\rangle$ and, by D_0:

7. $v'(q) = 1$.

Then, from (vii) and (iii) it follows that $v' \cup \{\langle r, 0\rangle\}$ would be non-admissible by the meaning constraints on $\to$. Hence, $v' \Vdash_0 \langle r, 1\rangle$, and by D_0 again:

8. $v'(r) = 1$.

The remaining steps of the argument are similar and can be summarized as follows:

9. $v'(s) = 1$, by (iv), (vii), the meaning constraints on $\to$ and D_0.
10. $v'(r \wedge s) = 1$, by (viii), (ix), the meaning constraints on $\wedge$ and D_0.

11. $v'(\neg(r \wedge s)) = 0$, by (x), the meaning constraints on $\neg$ and D_0.

12. $v'(\neg t) = 0$, by (v), (xi), the meaning constraints on $\rightarrow$ and D_0.

13. $v'(t) = 1$, by (xii), the meaning constraints on $\neg$ and D_0.

14. $v'(u) = 0$, by (vi), (xiii), the meaning constraints on $\wedge$ and D_0.

Intuitively, a shallow information state represents the overall information that a reasoner holds, either explicitly or implicitly, on the sole basis of the intended (informational) meaning of the logical operators and of the basic consistency principle expressed by the closure condition D_0.

In what follows we shall make use of *signed formulae* (*S-formulae* for short), namely expressions of the form TA or FA with the intended meaning of "A is informationally true" and "A is informationally false". This choice allows us to express an information state V as a set of S-formulae, namely the set $\{TA \mid v(A) = 1\} \cup \{FA \mid v(A) = 0\}$. We shall use "$\varphi, \psi, \theta, \ldots$", as variables ranging over S-formulae. We shall also use "$X, Y, Z, \ldots$", as variables ranging over sets of S-formulae and "$\Gamma, \Delta, \Lambda, \ldots$", as variables ranging over sets of unsigned formulae.

Let us say that an information state v *satisfies* an S-formula TA if $V(A) = 1$ and an S-formula FA if $v(A) = 0$. For every set X of S-formulae and every S-formula φ, we say that:

- φ is a 0-*depth consequence* of X if v satisfies φ for every information state v such that v satisfies all the S-formulae in X.

- X is 0-*depth inconsistent* if there is no information state v such that v satisfies all the S-formulae in X.

We use the symbol "$\vDash_0$" for the 0-depth consequence relation and write "$X \vDash_0 \varphi$" for "φ is a 0-depth consequence of X". The notions of 0-depth consequence and 0-depth inconsistency can be extended to unsigned formulae by stipulating that an unsigned formula A is a 0-depth consequence of a set Γ of unsigned formulae if and only if $T\Gamma \vDash_0 TA$ and that Γ is 0-depth inconsistent if and only if $T\Gamma$ is 0-depth inconsistent. In (D'Agostino *et al.*, 2013) (Proposition 2.49) it is shown that 0-depth consequence is not a (finite) many-valued logic, that is, it cannot be characterized by any set of finitely valued matrices.

$$\frac{F\,A}{T\,\neg A}\;T\neg\text{-}\mathscr{I} \qquad \frac{T\,A}{F\,\neg A}\;F\neg\text{-}\mathscr{I}$$

$$\frac{T\,A}{T\,A\vee B}\;T\vee\text{-}\mathscr{I}1 \qquad \frac{T\,B}{T\,A\vee B}\;T\vee\text{-}\mathscr{I}2 \qquad \frac{\begin{array}{c}F\,A\\ F\,B\end{array}}{F\,A\vee B}\;F\vee\text{-}\mathscr{I}$$

$$\frac{F\,A}{F\,A\wedge B}\;F\wedge\text{-}\mathscr{I}1 \qquad \frac{F\,B}{F\,A\wedge B}\;F\wedge\text{-}\mathscr{I}2 \qquad \frac{\begin{array}{c}T\,A\\ T\,B\end{array}}{T\,A\wedge B}\;T\wedge\text{-}\mathscr{I}$$

$$\frac{F\,A}{T\,A\to B}\;T\to\text{-}\mathscr{I}1 \qquad \frac{T\,B}{T\,A\to B}\;T\to\text{-}\mathscr{I}2 \qquad \frac{\begin{array}{c}T\,A\\ F\,B\end{array}}{F\,A\to B}\;F\to\text{-}\mathscr{I}$$

TABLE 2: *Introduction rules for the standard Boolean operators.*

It is not difficult to show that he relation $\vDash_0$ is a *logic in Tarski's sense*, that is, it satisfies the following conditions:

(Reflexivity)	$A \vDash_0 A$
(Monotonicity)	$\Gamma \vDash_0 A \Longrightarrow \Gamma \cup \Delta \vDash_0 A$
(Cut)	$\Gamma \vDash_0 A$ and $\Gamma \cup \{A\} \vDash_0 B \Longrightarrow \Gamma \vDash_0 B$.
(Substitution Invariance)	$\Gamma \vDash_0 A \Longrightarrow \sigma\Gamma \vDash_0 \sigma A$ for all substitutions σ

where $\sigma\Delta$ is short for $\{\sigma A | A \in \Delta\}$.

3.3. Intelim sequences

A natural proof-theoretical characterization of the 0-depth consequence relation $\vDash_0$ is obtained by means of a set of introduction and elimination rules (*intelim* rules) for the logical operators. These rules are shown in Tables 2 and 3 and are expressed in terms of S-formulae. A version of these rules for unsigned formulae is obtained by removing all the occurrences of the sign T and replacing all the occurrences of the sign F with the negation sign $\neg$.

Given a set X of S-formulae:

- An *intelim sequence* for X is a sequence $\varphi_1, \ldots \varphi_n$ of S-formulae such that, for every $i = 0, \ldots, n$, either $\varphi_i \in X$ or is the conclusion of the application of an intelim rule to preceding formulae.

$$\frac{T\neg A}{FA}\;T\neg\text{-}\mathscr{E} \qquad \frac{F\neg A}{TA}\;F\neg\text{-}\mathscr{E}$$

$$\frac{\begin{matrix}TA\vee B\\ FA\end{matrix}}{TB}\;T\vee\text{-}\mathscr{E}1 \qquad \frac{\begin{matrix}TA\vee B\\ FB\end{matrix}}{TA}\;T\vee\text{-}\mathscr{E}2 \qquad \frac{FA\vee B}{FA}\;F\vee\text{-}\mathscr{E}1 \qquad \frac{FA\vee B}{FB}\;F\vee\text{-}\mathscr{E}2$$

$$\frac{\begin{matrix}FA\wedge B\\ TA\end{matrix}}{FB}\;F\wedge\text{-}\mathscr{E}1 \qquad \frac{\begin{matrix}FA\wedge B\\ TB\end{matrix}}{FA}\;F\wedge\text{-}\mathscr{E}2 \qquad \frac{TA\wedge B}{TA}\;T\wedge\text{-}\mathscr{E}1 \qquad \frac{TA\wedge B}{TB}\;T\wedge\text{-}\mathscr{E}2$$

$$\frac{\begin{matrix}TA\rightarrow B\\ TA\end{matrix}}{TB}\;T\!\rightarrow\!\text{-}\mathscr{E}1 \qquad \frac{\begin{matrix}TA\rightarrow B\\ FB\end{matrix}}{FA}\;T\!\rightarrow\!\text{-}\mathscr{E}2 \qquad \frac{FA\rightarrow B}{TA}\;F\!\rightarrow\!\text{-}\mathscr{E}1 \qquad \frac{FA\rightarrow B}{FB}\;F\!\rightarrow\!\text{-}\mathscr{E}2$$

TABLE 3: *Elimination rules for the four standard Boolean operators*

- An intelim sequence is *closed* when it contains both TA and FA for some A.
- An *intelim refutation* of X is a closed intelim sequence for X.
- A *intelim proof* of φ from X is an intelim sequence for X such that φ is the last S-formula in the sequence.
- X is *intelim-refutable* if there is a closed intelim sequence for X.
- An *S-formula* φ is *intelim deducible* from X if there is an intelim proof of φ from X.

In Figure 1 we show simple examples of intelim sequences using, respectively, the intelim rules for signed formulae and their version for unsigned formulae. We use the symbol "$\vdash_0$" to denote the relation of intelim-deducibility and write "$X \vdash_0 \varphi$" for "φ is a 0-depth deducible from X".

Again, we can extend the notions of intelim deducibility and refutability to unsigned formulae by stipulating that an unsigned formula A is intelim deducible from a set Γ of unsigned formulae if TA is intelim deducible from $T\Gamma$ and that a set Γ of unsigned formulae is intelim refutable if $T\Gamma$ is intelim refutable.

1	$T\,(p \vee q) \to \neg r$	Assumption	1	$(p \vee q) \to \neg r$	Assumption
2	$T\,p$	Assumption	2	p	Assumption
3	$T\,(p \wedge t) \to r$	Assumption	3	$(p \wedge t) \to r$	Assumption
4	$T\,p \vee q$	$T\vee$-$\mathscr{I}1$ (2)	4	$p \vee q$	$T\vee$-$\mathscr{I}1$ (2)
5	$T\,\neg r$	$T\to$-$\mathscr{E}1$ (1,4)	5	$\neg r$	$T\to$-$\mathscr{E}1$ (1,4)
6	$F\,r$	$T\neg$-$\mathscr{E}$ (5)	6	$\neg(p \wedge t)$	$T\to$-$\mathscr{E}2$ (3,5)
7	$F\,(p \wedge t)$	$T\to$-$\mathscr{E}2$ (3,6)	7	$\neg t$	$F\wedge$-$\mathscr{E}1$ (6,2).
8	$F\,t$	$F\wedge$-$\mathscr{E}1$ (7,2).			

FIG. 1: *On the left, an intelim sequence using the rules for signed formulae. On the right the corresponding sequence using the rules for unsigned formulae.*

The *unsigned part* of an S-formula is the unsigned formula that results from it by removing the sign T or F. Given an S-formula φ, we denote by φ^u the unsigned part of φ. We say that an intelim proof of φ from X (an intelim refutation of X) has the *subformula property* (SFP) if, for every S-formula ψ occurring in it, ψ^u is a subformula of θ^u for some θ in $X \cup \{\varphi\}$ (in X).

For an intelim sequence π, let $|\pi|$ denote the length of π.

Proposition 3.2 (Subformula Property)**.**

1. *For every intelim proof π of φ from X:*
   - *if π is an open intelim sequence, π can be transformed into an intelim proof π' of φ from X such that π' has the SFP and $|\pi'| \leq |\pi|$;*
   - *if π is a closed intelim sequence, there is an intelim refutation π' of X such that π' has the SFP and $|\pi'| \leq |\pi|$.*
2. *For every intelim refutation π of X, π can be transformed into an intelim refutation π' of X such that π' has the SFP and $|\pi'| \leq |\pi|$.*

A proof of the subformula property for generalized intelim systems with arbitrary Boolean operators can be found in (D'Agostino *et al.*, 2013). The subformula property of intelim proofs and refutations paves the way for a feasible decision procedure for intelim deducibility and refutability. A generalization of the following proposition is also proven in (D'Agostino *et al.*, 2013).

Proposition 3.3. *Whether or not $X \vdash_0 \varphi$ (X is 0-depth refutable) can be decided in time $O(n^2)$ where n is the total number of occurrences of symbols in $X \cup \{\varphi\}$ (in X).*

$$\frac{T\,A \vee A}{T\,A}\; T\vee\text{-mingle} \qquad\qquad \frac{F\,A \wedge A}{F\,A}\; F\wedge\text{-mingle}$$

TABLE 4: *Mingle rules for* $\vee$ *and* $\wedge$.

The reader can verify that all the intelim rules are sound, namely that the conclusion is a 0-depth logical consequence of the premises. As for completeness, there is a technical point which must be taken care of. It follows from the principle D_0 above that in every shallow information state such that $v(A \vee A) = 1$, it must hold true that $v(A) = 1$. Similarly, in every shallow information state such that $v(A \wedge A) = 0$ it must hold true that $v(A) = 0$. So, the value of A is by all means dictated, in both cases, by the value of $A \vee A$ ($A \wedge A$) and by the intended meaning of $\vee$ ($\wedge$) as specified by the meaning constraints. However, our intelim rules do not allow us, as they stand, to infer $T\,A$ from $T\,A \vee A$, or $F\,A$ from $F\,A \wedge A$, This technical problem can be addressed in two different ways:

1. we pre-process all formulae and replace every occurrence of $A \vee A$, respectively $A \wedge A$, with A;

2. we introduce the two sound "mingle" rules in Table 4 in addition to the intelim rules in Tables 2 and 3; this solution is the same as the one adopted in (Finger & Gabbay, 2006) in response to a similar problem arising in their investigations into tractable subsystems of classical propositional logic.

Both solutions appear quite reasonable, although not terribly elegant. The first one requires restrictions on the language that do not undermine its expressive power. The second solution requires the addition of *ad hoc* rules that somehow spoil the harmony of the intelim approach, but are, on the other hand, perfectly justified by the aim of extracting all the information that can possibly be obtained on the sole basis of the meaning of the logical operators. A more elegant solution would consist in recognizing that the traditional approach, based on standard inference rules, is perhaps not ideal to represent the flow of logical information, and revert to a less traditional approach based on logical networks, along the lines suggested in the Appendix of (D'Agostino & Floridi, 2009).

Let us call intelim* the set consisting of the intelim rules in Tables 2-3 and of the mingle rules in Table 4 and assume the notions of intelim$^+$-sequence, intelim* deduction and intelim* refutation are defined as above replacing "intelim" with "intelim*". Then we have the following:

Proposition 3.4. *For every finite set* Γ *of sentences and every sentence* A, $\Gamma \vDash_0 A$ *if and only if there is an intelim* deduction of* A *from* Γ.

Clearly the tractability of intelim-deducibility extends to intelim*-deducibility. Hence the 0-depth logic is truly informationally closed in that there exists a simple feasible procedure for deciding whether the information that the conclusion is true is implicitly contained, by virtue of the informational meaning of the logical operators, in the information that the premises are true.

3.4. Modular semantics

In this section we present a variant of the constraint-based semantics first introduced in (D'Agostino *et al.*, 2013). Under this "modular semantics" the mingle rules of Table 4 are no longer sound and so a completeness theorem can be shown with respect to the pure intelim rules of Table 2 and 3. Here we consider any propositional language $\mathscr{L}$ with *arbitrary* Boolean operators and regard partial valuations as total mappings from the well-formed formulae of $\mathscr{L}$ to $\{0,1,\bot\}$. So"$\bot$" is no longer just a notational device for expressing that a valuation v is undefined for a given argument, but denotes a third "indeterminate" value. We call these mappings **3**-valuations.

We take the three values as partially ordered by the relation $\preceq$ such that $x \preceq y$ ("x is less defined than, or equal to, y") if, and only if, $x = \bot$ or $x = y$ for $x, y \in \{0,1,\bot\}$. As before, **3**-valuations are partially ordered by the usual approximation relation $\sqsubseteq$ defined as follows: $v \sqsubseteq w$ if and only if $v(A) \preceq w(A)$ for all A. The set of all **3**-valuations partially ordered by $\sqsubseteq$ forms a meet-semilattice with a bottom element consisting of the **3**-valuation that takes the undefined value for all formulae. We denote by $v \sqcap w$ the meet of the **3**-valuations v and w. In this setting an "information state" will be a special kind of **3**-valuation that (i) is "locally" compatible with the classical truth-table and (ii) is closed under a most basic kind of implicit information.

First, we need to pick out, from the set of all **3**-valuations, those that agree with the classical truth-tables. In the presence of the "indeterminate" value $\bot$, this can be done in a variety of ways. Here we describe a *modular* approach, in which agreement is checked only "locally" with respect to the main logical operator of a formula, ignoring the other operators that may occur in it, that is, treating the operands as if they were (distinct) atomic formulae.

Recall that an $\mathscr{L}$-module is a set consisting of a non-atomic $\mathscr{L}$-formula, called the *top formula* of the module, and of its *immediate* subformulae that we call *secondary formulae*. As in the previous section we shall denote by $\mathrm{Mod}(A)$ the unique $\mathscr{L}$-module whose top formula is A. A *valuation module* is any mapping $\alpha : \mathrm{Mod}(A) \to \{0,1,\perp\}$, where A is some non-atomic formula. Given a **3**-valuation v, we say that α *is a module of* v if $\alpha = v \mid \mathrm{dom}(\alpha)$.

By extension, we shall call *top formula of* α the top formula of $\mathrm{dom}(\alpha)$ (the domain of α), and *main operator of* α the main operator of the top formula of α.

For each non-atomic formula $B = \star(A_1, \ldots, A_n)$ and each $i = 1, \ldots, n+1$, let

$$\pi_i(B) =_{\mathrm{def}} \begin{cases} A_i & \text{if } 1 \leq i \leq n \\ B & \text{if } i = n+1 \end{cases}$$

In the sequel we shall use the boldface letters $\mathbf{x}, \mathbf{y}, \mathbf{z}$, etc. to denote finite vectors with values in $\{0,1,\perp\}$. For each valuation module $\alpha : \mathsf{Mod}(A) \to \{0,1,\perp\}$, the *value vector* of α, denoted by $\vec{\alpha}$, is the vector $\mathbf{x} \in \{0,1,\perp\}^{a+1}$, where a is the arity of the main logical operator of A, such that $x_i = \alpha(\pi_i(A))$ for all $i = 1, \ldots, a+1$.

We say that a vector $\mathbf{x} \in \{0,1,\perp\}^n$ *approximates* a vector $\mathbf{y} \in \{0,1,\perp\}^n$, and write $\mathbf{x} \preceq \mathbf{y}$ if, and only if, $x_i \preceq y_i$ for all $i = 1, \ldots, n$. Notice that, for every fixed n, the set of all vectors in $\{0,1,\perp\}^n$ partially ordered by $\preceq$ is a meet semilattice. We denote by $\mathbf{x} \wedge \mathbf{y}$ the meet of $\mathbf{x}$ and $\mathbf{y}$. For each n-ary operator $\star$, let us denote by $\mathscr{A}_\star$ the set of all vectors in $\{0,1,\perp\}^{n+1}$ that approximate some vector in the graph of $f_\star$ (i.e., the Boolean function associated with $\star$), namely the set

$$\{\mathbf{x} \mid \exists \mathbf{y}(\mathbf{x} \preceq \mathbf{y} \text{ and } \mathbf{y} \in f_\star)\},$$

where by "$\mathbf{y} \in f_\star$" we mean that $\mathbf{y}$ belongs to the graph of $f_\star$. Let $\mathbf{M}_\star$ denote the set of all valuation modules α such that the top formula of α is a $\star$-formula. A *module* $\alpha \in \mathbf{M}_\star$ is *admissible* iff its value vector $\vec{\alpha}$ is in $\mathscr{A}_\star$. A ***3**-valuation* v is *admissible* iff all its modules are admissible. Thus an admissible **3**-valuation is one that conveys *partial* information about the sentences of the given language in a way that agrees with the Boolean truth-tables for the logical operators in the modular sense specified above: the value assigned to each formula A is compatible with those assigned to its immediate subformulae according to the truth-table for the *main* logical operator of A.

Now, we explain what it means, in this setting, that a certain piece of information, specifying that a given sentence is true or false, is "implicitly contained" in the information that is explicitly conveyed by a partial valuation v. Let us denote by $\mathbf{x}[i := z]$, for $z \in \{0, 1, \bot\}$ the vector $\mathbf{y}$ such that $y_i = z$ and $y_j = x_j$ for all $j \neq i$. For every n-ary operator $\star$, a vector $\mathbf{x} \in \mathscr{A}_\star$ is *stable in* $\mathscr{A}_\star$ if, and only if, for every i, $1 \leq i \leq n+1$,

$$x_i = \bot \Longrightarrow \mathbf{x}[i := 0] \in \mathscr{A}_\star \text{ and } \mathbf{x}[i := 1] \in \mathscr{A}_\star.$$

In other words, a stable approximation of a vector in the graph of $f_\star$ is a vector such that each of its "undefined" components may indifferently take, in some better approximation, either of the two defined values 0 and 1. On the other hand, a vector in $\mathscr{A}_\star$ is unstable whenever, for some of its undefined components, one of the two possible defined values is ruled out by the truth-table for $\star$, and so the other one is uniquely determined as the only possible defined value that this component can take. A *valuation module* $\alpha \in \mathbf{M}_\star$ is *informationally closed* iff $\vec{\alpha}$ is a stable element of $\mathscr{A}_\star$. A *valuation* v is *informationally closed* if all its modules are informationally closed. If α is not informationally closed, that is, $\vec{\alpha}$ is *un*stable, then there is an $A \in \mathrm{dom}(\alpha)$ such that $\alpha(A) = \bot$, but α carries implicit information about A, because there is a defined value for A that is *deterministically dictated* by the truth-table for the main operator of α. An informationally closed module is one in which all such implicit information has been made explicit. It is not difficult to verify that:

Proposition 3.5. *For every n-ary operator $\star$, if $\mathbf{x}$ and $\mathbf{y}$ are both stable elements of $\mathscr{A}_\star$, their meet $\mathbf{x} \wedge \mathbf{y}$ is also a stable element of $\mathscr{A}_\star$.*

This is sufficient to ensure that, for every α, the set of all informationally closed β such that $\alpha \sqsubseteq \beta$ has a minimum element. Therefore, for every valuation v, the set of informationally closed **3**-valuations w such that $v \sqsubseteq w$ also has a minimum element. Now, we can say that a **3**-valuation is an *information state* if it is informationally closed.

Examples 3.6. According to the above definition, the module α described by the following table:

(6)

$p \vee q$	p	$(p \vee q) \wedge p$
1	$\bot$	0

is admissible, because its value vector $\vec{\alpha} = (1,\bot,0)$ approximates the vector $(1,0,0)$ which is in the graph of $f_\wedge$ (since $f_\wedge(1,0) = 0$). On the other hand, the following modules:

$$(7) \qquad \frac{p\vee r \quad q\to s \quad (p\vee r)\wedge(q\to s)}{0 \qquad \bot \qquad\qquad 1} \qquad \frac{p\wedge r \quad q \quad (p\wedge r)\vee q}{1 \qquad \bot \qquad 0},$$

are *not* admissible, because $(0,\bot,1)$ does not approximate any vector in the graph of $f_\wedge$ and $(1,\bot,0)$ does not approximate any vector in the graph of $f_\vee$.

The following modules

$$(8) \qquad \frac{p\vee r \quad q\to s \quad (p\vee r)\wedge(q\to s)}{\bot \qquad \bot \qquad\qquad 0} \qquad \frac{p\wedge r \quad q \quad (p\wedge r)\vee q}{\bot \qquad \bot \qquad 1},$$

are informationally closed. The first module, is informationally closed because its value vector $(\bot,\bot,0)$ is a stable element of $\mathscr{A}_\wedge$, that is, the vectors $(0,\bot,0)$, $(1,\bot,0)$, $(\bot,0,0)$ and $(\bot,1,0)$ are all in $\mathscr{A}_\wedge$. The second one is informationally closed because its value vector $(\bot,\bot,1)$ is a stable element of $\mathscr{A}_\vee$, that is, $(0,\bot,1)$, $(1,\bot,1)$, $(\bot,0,1)$ and $(\bot,1,1)$ are all in $\mathscr{A}_\vee$. On the other hand, the module (6) is not *informationally* closed, because $(1,1,0)$ is not in $\mathscr{A}_\wedge$.

Observe that under this definition of information state: (i) every Boolean valuation is an information state; (ii) for every information state v, if $v(A) \neq \bot$ for all well-formed formulae A of $\mathscr{L}$, then v is a Boolean valuation. The consequence relation of depth 0, $\models_0$ can now be defined as in the previous section: for every set X of S-formulae and every S-formula φ, $X \models_0 \varphi$ if and only if φ is satisfied by every information state that satisfies all the S-formulae in X. (Notice the the symbol $\models_0$ is slightly different from the symbol $\vDash_0$ used in the constraint-based semantics of the previous section.) Again $\models_0$ is a logic in Tarski's sense, that is, it is closed under Reflexivity, Monotonicity, Cut and Substitution Invariance. However, the notion of information state defined in terms of this modular semantics is different (slightly weaker) than the notion of information state defined in terms of the constraint-based semantics of the previous section. Indeed, as anticipated, the consequence relation $\models_0$ defined in terms of the modular semantics is faithful to the deducibility relation based on the intelim rules only, since the mingle rules turn out to be unsound. So we have the following:

Proposition 3.7. *For every finite set Γ of sentences and every sentence A, $\Gamma \models_0 A$ if and only if there is an intelim deduction of A from Γ.*

3.5. Virtual information and depth-bounded deduction

In Sections 3 and 3 we have introduced two related basic mechanisms for extracting implicit information from the data (expressed by a partial valuation) by virtue of the intended (informational) meaning of the logical operators. Any "deeper" processing of the data must introduce and use information that is not uniquely determined in this way. For example, to establish that r is a Boolean consequence of $p \vee q$, $p \to r$ and $q \to r$, we start with the partial valuation v such that

$$v(A) = \begin{cases} 1 & \text{if } A \in \{p \vee q, p \to r, q \to r\} \\ \perp & \text{otherwise} \end{cases}$$

Given the subformula property of intelim deducibility, the only valuation modules of v that are significant for the deduction problem under consideration are those extensionally described by the following diagrams (the links represent the subformula relation):

$$\begin{tikzcd}[column sep=small]
& \langle p \vee q, 1\rangle \arrow[dl, dash] \arrow[dr, dash] & & & \langle p \to r, 1\rangle \arrow[dl, dash] \arrow[dr, dash] & & & \langle q \to r, 1\rangle \arrow[dl, dash] \arrow[dr, dash] & \\
\langle p, \perp\rangle & & \langle q, \perp\rangle & \langle p, \perp\rangle & & \langle r, \perp\rangle & \langle q, \perp\rangle & & \langle r, \perp\rangle
\end{tikzcd}$$

These valuation modules are all informationally closed: in each of them there is no way of determining the value of some undefined formula from the information explicitly stored in the module. To extract the information that r must be true, we necessarily have to consider possible refinements of v containing information that is not implicitly determined by it. For example, we can consider its possible alternative refinements v_1 and v_2 such that $v_1(p) = 1$, $v_2(p) = 0$ and $v_1(A) = v_2(A) = v(A)$ for all $A \neq p$. Both such refinements of v contain information concerning p that is not contained in v. This is what we call *virtual information*. We may assume that either v_1 or v_2 must "eventually" obtain, that is, one of the two valuations will eventually express the information that is explicitly held by the given agent. Or, alternatively, argue that either v_1 or v_2 must agree with all "possible worlds" that agree with v. In either case, the relevant modules of v_1 and v_2 are, respectively,

$$
\begin{tikzcd}[arrows=dash]
 & \langle p \vee q, 1\rangle \arrow[dl] \arrow[dr] & & & \langle p \to r, 1\rangle \arrow[dl] \arrow[dr] & & & \langle q \to r, 1\rangle \arrow[dl] \arrow[dr] & \\
\langle p, 1\rangle & & \langle q, \bot\rangle & \langle p, 1\rangle & & \langle r, \bot\rangle & \langle q, \bot\rangle & & \langle r, \bot\rangle
\end{tikzcd}
$$

and

$$
\begin{tikzcd}[arrows=dash]
 & \langle p \vee q, 1\rangle \arrow[dl] \arrow[dr] & & & \langle p \to r, 1\rangle \arrow[dl] \arrow[dr] & & & \langle q \to r, 1\rangle \arrow[dl] \arrow[dr] & \\
\langle p, 0\rangle & & \langle q, \bot\rangle & \langle p, 0\rangle & & \langle r, \bot\rangle & \langle q, \bot\rangle & & \langle r, \bot\rangle
\end{tikzcd}
$$

It is easy to check that, in both cases, the least valuation (with respect to the $\sqsubseteq$ relation) in which all the relevant modules are informationally closed must verify r. So, we may wish to conclude that the piece of information that r is true is also implicitly contained in v. However, this essentially requires the consideration of refinements of v that contain what we have called virtual information, in the above example the information that p is true in v_1 and the information that p is false in v_2, namely information that is not implicitly contained in the original valuation v.

The unbounded use of virtual information, in the way just explained, turns an information state into a Boolean valuation. A natural way of approximating Boolean valuations, starting from the information states defined in Section 3 or in Section 3 — which correspond to the basis case in which *no* virtual information is allowed — consists in imposing an upper bound on the depth at which the nested use of virtual information is allowed. In (D'Agostino *et al.*, 2013) the reader can find a discussion of a variety of ways in which such approximations can be defined. Here we focus on a particularly simple approach based on the notion of *intelim tree* introduced in (Mondadori, 1989) and further developed in (D'Agostino, 2005).

Given a set of formulae Γ, we use the notation $\mathrm{Sub}(\Gamma)$ to denote the set of all subformulae of the formulae in Γ. For all $k \in \mathbb{N}_+$, the relation $\models_k$ is defined as follows:

1. $X \models_k \varphi$ if and only if $X \cup \{TA\} \models_{k-1} \varphi$ and $X \cup \{FA\} \models_{k-1} \varphi$ for some $A \in \mathrm{Sub}(X^u \cup \{\varphi^u\})$;

2. X is k-depth inconsistent if and only if $X \cup \{TA\}$ and $X \cup \{FA\}$ are both $k-1$-depth inconsistent for some $A \in \mathrm{Sub}(X^u)$.

Observe that, since $\models_0$ is monotonic, $\models_j \subseteq \models_k$ whenever $j \leq k$. The transition from $\models_k$ to $\models_{k+1}$ corresponds to an increase in the depth at which the nested use of virtual information is allowed. It is not difficult to show that:

Proposition 3.8. *The relation $\models_\infty = \bigcup_{k \in \mathbb{N}} \models_k$ is the consequence relation of classical propositional logic.*

Given Proposition 3.4 and the definitions of k-depth consequence $\models_k$ and k-depth inconsistency given above, the corresponding notions of k-depth deducibility and k-depth refutability for $k \in \mathbb{N}_+$, are trivially defined as follows: for all $k \in \mathbb{N}_+$,

1. $X \vdash_k \varphi$ if and only if $X \cup \{TA\} \vdash_{k-1} \varphi$ and $X \cup \{FA\} \vdash_{k-1} \varphi$ for some $A \in \mathrm{Sub}(X^u \cup \{\varphi^u\})$;
2. X is k-depth refutable if and only if $X \cup \{TA\}$ and $X \cup \{FA\}$ are both $k-1$-depth refutable for some $A \in \mathrm{Sub}(X^u)$.

As before, we extend the relations $\models_k$ and $\vdash_k$ to unsigned formulae by stipulating that $\Gamma \models_k A$ ($\Gamma \vdash_k A$) if and only if $T\Gamma \models_k TA$ ($T\Gamma \vdash_k TA$). Similarly, we stipulate that Γ is k-depth inconsistent (k-depth refutable) if and only if $T\Gamma$ is k-depth inconsistent (k-depth refutable).

While deductions of depth 0 are represented by *intelim sequences*, deductions of depth $k > 0$ may be aptly represented in the format of *intelim trees*. For this purpose it is sufficient to add to the intelim rules the following branching rule:[23]

$$\begin{array}{ccc} & & \text{PB} \\ / & & \backslash \\ TA & & FA \end{array}$$

Each application of this rule allows us to introduce virtual information concerning an arbitrary formula A by appending both TA and FA as sibling nodes at the end of any branch of the tree, generating two new branches. The formula A involved in a specific application of the rule is called *PB-formula*. The S-formulae

[23] "PB" stands for "Principle of Bivalence".

TA and FA are called *virtual assumptions*. Such a step invites us to consider information states that, besides containing all the information expressed by the preceding S-formulae in the branch, also contain definite information about the truth or falsity of the PB-formula A. An *intelim tree* for X is a tree $\mathscr{T}$ of S-formulae such that, for every S-formula φ in a branch of $\mathscr{T}$, either

1. $\varphi \in X$, or
2. φ is obtained from preceding S-formulae in the same branch by an application of an intelim rule, or
3. φ is a virtual assumption introduced by an application of the branching rule PB.

We say that a branch of an intelim tree is *closed* if it contains both TA and FA for some formula A, otherwise it is *open*. The *depth* of an intelim tree $\mathscr{T}$ is the maximum number of virtual assumptions occurring in a branch of $\mathscr{T}$. For all $k \in \mathbb{N}$,

1. A k-depth *intelim proof* of φ from X is an intelim tree $\mathscr{T}$ for X of depth k such that φ occurs in all open branches of $\mathscr{T}$;
2. A k-depth refutation of X is an intelim tree $\mathscr{T}$ for X of depth k such that every branch of $\mathscr{T}$ is closed.

For each intelim tree $\mathscr{T}$, let us denote by $\mathrm{PB}(\mathscr{T})$ the set of all PB-formulae occurring in $\mathscr{T}$.

Proposition 3.9. *For all $k \in \mathbb{N}$,*

1. *Every k-depth intelim proof $\mathscr{T}$ of φ from X can be transformed into a $k+j$-depth (with $j \geq 0$) intelim proof $\mathscr{T}'$ of φ from X such that* $\mathrm{PB}(\mathscr{T}') \subseteq \mathrm{Sub}(X^u \cup \varphi^u)$.
2. *every k-depth refutation $\mathscr{T}$ of X can be transformed into a $k+j$-depth (with $j \geq 0$) $\mathscr{T}'$ of X such that* $\mathrm{PB}(\mathscr{T}') \subseteq \mathrm{Sub}(X^u)$.

A proof of the above proposition can be adapted from (D'Agostino, 2005). Let us say that an intelim proof $\mathscr{T}$ of φ from X (an intelim refutation $\mathscr{T}$ of X) has the *subformula property* (SFP) if $\psi^u \in \mathrm{Sub}(X^u \cup \varphi^u)$ ($\psi^u \in \mathrm{Sub}(X^u)$) for every S-formula ψ occurring in $\mathscr{T}$.

Proposition 3.10. *For all $k \in \mathbb{N}$,*

1. *Every k-depth intelim proof $\mathscr{T}$ of φ from X such that $\mathrm{PB}(\mathscr{T}) \subseteq \mathrm{Sub}(X^u \cup \varphi^u)$ can be transformed into a k-depth intelim proof $\mathscr{T}'$ of φ from X such that $\mathscr{T}'$ has the SFP.*

2. *Every k-depth refutation $\mathscr{T}$ of X such that $\mathrm{PB}(\mathscr{T}) \subseteq \mathrm{Sub}(X^u)$ can be transformed into a k-depth refutation $\mathscr{T}'$ of X such that $\mathscr{T}'$ has the SFP.*

Then, it is easy to show that:

Proposition 3.11. *For all $k \in \mathbb{N}$,*

1. *$X \vdash_k \varphi$ if and only if, for some $j \leq k$, there is a j-depth intelim proof $\mathscr{T}$ of φ from X such that $\mathscr{T}$ has the SFP;*

2. *X is k-depth refutable if and only if for some $j \leq k$, there is a j-depth intelim refutation $\mathscr{T}$ of X such that $\mathscr{T}$ has the SFP.*

An example of an intelim proof of depth 2with the SFP is given in Figure 2. This is an intelim proof of TH from premises 1–6. The reader can check that each S-formula that is not a premise either is obtained from previous S-formulae on the same branch by an application of one of the intelim rules in Tables 2 and 3, or is one of the virtual assumptions introduced by the branching rule PB. All the open branches end with the S-formula TH. The rightmost branch is closed since it contains both TC and FC. Each open branch is a 0-depth intelim proof of TH from the union of the initial premises 1–6 plus the virtual assumptions introduced by the rule PB on that branch. Given Propositions 3.3 and 3.11 it is not difficult to show that, for each fixed k, $\vdash_k$ admits of a feasible decision procedure:

Proposition 3.12. *For each $k \in \mathbb{N}$, whether or not $X \vdash_k \varphi$ (X is k-depth refutable), can be decided in time $O(n^{k+2})$, where n is the total number of occurrences of symbols in $X \cup \{\varphi\}$ (in X).*

As for the intelim rules, a version of PB for unsigned formulae is obtained simply by removing the sign T and replacing the sign F with the negation sign. Examples of intelim trees using the rules for unsigned formulae are given in Figure 3. In each tree the premises are the formulae labelled with a circled "p". The first tree is a 1-depth intelim deduction of v from the premises. The second is a 2-depth intelim deduction of v from the premises. The third is a 1-depth intelim refutation of the premises.

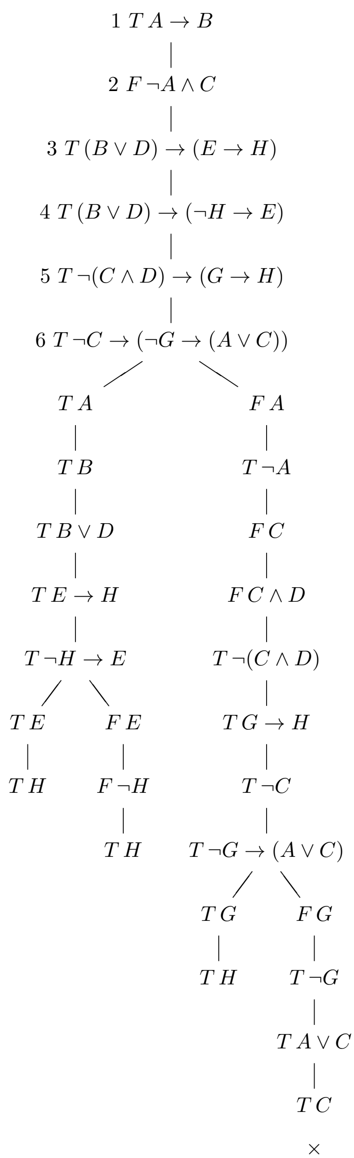

FIG. 2: *An intelim proof of depth 2.*

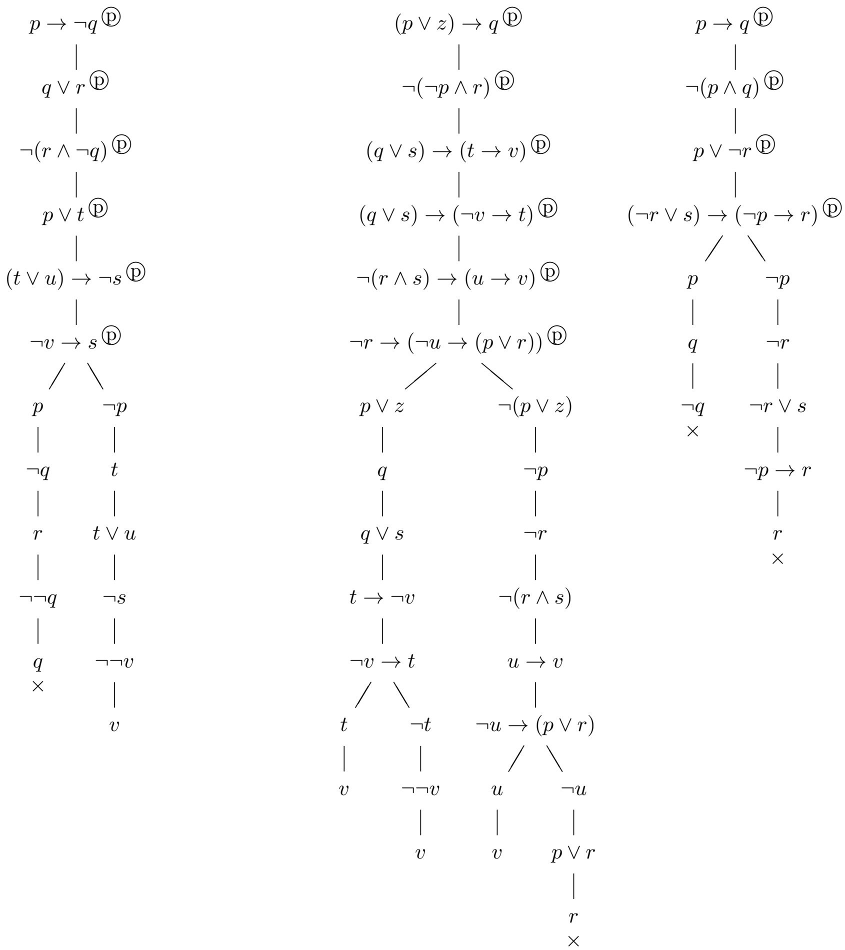

FIG. 3: *Intelim trees for unsigned formulae*

3.6. Solving the anomalies of the received view

We now show how the informational view of classical logic outlined in the previous sections may help to solve the anomalies of the received view discussed in Section 2. In what follows, we use "information state" to refer ambiguously to a partial valuation closed under implicit information either in the sense of the constraint-based semantics of Section 3 or in the sense of the modular semantics of Section 3. Clearly the proposed notions of information content will reflect the differences in these two notions of information state. In either case in this section we think of an information state as a set S of pieces of information of the form $\langle A,1\rangle$ or $\langle A,0\rangle$ such that for no A, $\langle A,1\rangle$ and $\langle A,0\rangle$ are both in S.

The set $\mathscr{S}$ of all information states is naturally ordered by set inclusion. Every non-empty subset $\mathscr{P}$ of $\mathscr{S}$ has a meet in $\mathscr{S}$ given by $\bigcap\mathscr{P}$. On the other hand, two information states might not have a join in $\mathscr{S}$. Indeed, if two information states are mutually inconsistent they have no upper bounds in $\mathscr{S}$. Observe also that, even when two information states have a join in $\mathscr{S}$ this is not, in general, their set union. For example, the join in $\mathscr{S}$ of two information states containing, respectively, $\langle p\vee q,1\rangle$ and $\langle p,0\rangle$, must contain also the signed sentence $\langle q,1\rangle$ that may not be contained in either of them. Given a subset $\mathscr{P}$ of $\mathscr{S}$, let $\mathscr{P}^u$ be the set of all upper bounds of $\mathscr{P}$ in $\mathscr{S}$. Then, $\mathscr{P}$ has a join in $\mathscr{S}$ whenever $\mathscr{P}^u$ is non-empty, and this is given by $\bigcap\mathscr{P}^u$. Now, since $\mathscr{S}$ itself has no upper bounds in $\mathscr{S}$, this ordering is topless. Let $\top$ be the set of *all* pieces of information (which is *not* an information state) and let $\mathscr{S}^* = \mathscr{S}\cup\{\top\}$. Then $(\mathscr{S}^*,\subseteq)$ is a complete lattice, where the meet of an arbitrary subset $\mathscr{P}$ of $\mathscr{S}^*$ is given by $\bigcap\mathscr{P}$, while its join is equal either to the top element $\top$, if $\mathscr{P}^u$ is empty, or to $\bigcap\mathscr{P}^u$ otherwise.

Now, the *surface information* carried by a *sentence* A, $\mathrm{INF}(A)$ can be defined as

(9)
$$\mathrm{INF}(A) = \bigsqcap\{S\in\mathscr{S}\mid TA\in S\}.$$

More generally, the *surface information* carried by a *set* Γ of sentences can be defined as

(10)
$$\mathrm{INF}(\Gamma) = \bigsqcap\{Y\in\mathscr{S}\mid \Gamma\subseteq Y\}.$$

Observe that, since $\bigsqcap\emptyset = \top$, (10) yields $\mathrm{INF}(\Gamma)=\top$ whenever Γ is 0-depth inconsistent, for there is no $Y\in\mathscr{S}$ that may include Γ. Recall that $\top$ is not an information state, but only denotes a situation in which all information is "suspended" and can be rather interpreted as a call for revision. So $\top$ is conceptually distinct from the *empty information state*, that is, the partial valuation that is undefined for all formulae. However, an agent whose informational situation is described by $\top$ holds no genuine information just as any agent whose information state is empty. Then, in order to be informative for an agent a, a (set of) sentence(s) must be 0-depth consistent.

This requirement of 0-depth consistency (not classical consistency) can be seen as a substantial mitigation of the "veridicality thesis" put forward by Luciano Floridi as a solution to the BHC paradox. Even if one is not willing to endorse the somewhat controversial view that "information encapsulates truth" (Floridi 2004 and Floridi 2011, Chapters 4–5), one can still maintain that a minimal interpretation of "holding information" is one that satisfies the requirement that no agent may hold information that is explicitly inconsistent. And if a set of sentences Γ is 0-depth inconsistent, no agent a can "hold the information" that all the sentences in Γ are true, because adding Γ to a's current information state would destroy the latter as an information state.

The *informativeness of* Γ *for an agent* a, $\iota_a(\Gamma)$ can be characterized as follows:

(11) $$\iota_a(\Gamma) = \mathrm{INF}(\mathrm{INF}(S_a \cup \Gamma) \sim S_a),$$

where S_a is the current information state of a. Again, it follows from (11) that $\iota_a(\Gamma) = \top$ whenever Γ is 0-depth inconsistent.

On the basis of the above definitions, the 0-depth consequence relation can be equivalently defined as follows:

$\Gamma \vdash \varphi$ if and only if $\mathrm{INF}(\varphi) \subseteq \mathrm{INF}(\Gamma)$.

Hence, $\vdash$ is informationally trivial, in that every agent that actually holds the information that the premises are true must thereby hold the information that the conclusion is true, or equivalently, the surface semantic information carried by the conclusion is included in the surface semantic information carried by the premises. The latter wording covers the limiting case in which the surface information carried by the premises is $\top$ which do not qualify as genuine information ($\top$ is not an information state).

It may be objected that the consequence relation $\models$ is still "explosive" when Γ is 0-depth inconsistent, for there is no information state for a that contains

$\langle A, 1\rangle$ for all $A \in \Gamma$. So, if Γ is 0-depth inconsistent $\Gamma \models A$ for every sentence A. Similarly, if we follow the informational definition of 0-depth consequence just given, when Γ is 0-depth inconsistent, $\mathrm{INF}(\Gamma) = \top$ and so, for every A, $\mathrm{INF}(A) \subseteq \mathrm{INF}(\Gamma)$. However, the problem raised by this kind of explosivity is far less serious than the similar problem for the classical consequence relation. For, we can detect that the premises are 0-depth inconsistent in feasible time. Unlike hidden classical inconsistencies, that may be hard to discover even for agents equipped with powerful (but still bounded) computational resources, 0-depth inconsistency lies, as it were, on the surface. So, we always have a feasible means to ensure that our premises are 0-depth consistent, in which case the consequence relation $\models_0$ is *not* explosive, even if these premises are *classically inconsistent*.

We stress again that our definition of information state and surface information do not require that information "encapsulates truth", nor do they even require that it "encapsulates consistency", but only that information "encapsulates surface (0-*depth*) consistency". According to this characterization, $\models$ is informationally trivial by definition, and this is in accordance with the tenet that analytic inferences are utterly uninformative. The valid inferences of $\models$ are only a subclass of the classically valid inferences and their validity can be recognized in feasible time. These are the "easy" inferences that (nearly) everybody learns to make correctly in the very process of learning the meaning of the logical operators.

4. Conclusions

In (D'Agostino & Floridi, 2009) we described "virtual information" as information that is by no means contained in the information carried by the premises of an inference, but is still essentially, if only temporarily, involved in obtaining the conclusion. It is the kind of provisional assumptions that occur in the so-called "discharge rules" of Gentzen's natural deduction and, more generally, in any kind of "reasoning by cases". For example, the following inference:

(12)
$$\begin{array}{c} T\,A \vee B \\ T\,A \to C \\ T\,B \to C \\ \hline T\,C \end{array}$$

is classically valid, but cannot be immediately justified by means of the intelim rules that mirror what we have called the "informational semantics" of the logical operators.

An argument to show the validity of (12) based on these rules will have necessarily to resort to an intelim tree such as the following:

$$
\begin{array}{ccc}
 & T\,A \vee B & \\
 & | & \\
 & T\,A \to C & \\
 & | & \\
 & T\,B \to C & \\
 / & & \backslash \\
T\,A & & F\,A \\
| & & | \\
T\,C & & T\,B \\
 & & | \\
 & & T\,C
\end{array}
$$

Here, the information expressed by the signed sentences $T\,A$ and $F\,A$ is *not*, in general, information that is actually held by the agent who holds the information carried by the premises. It is *virtual information* that goes *beyond what is "given"* in the premises. This use of virtual information, that is not contained in the data and so may not be actually held by any agent who holds the information carried by the data, appears to qualify this kind of argument as "synthetic" in a sense close to Kant's sense, in that it forces the agent to consider potential information that is not included in the information "given" to him:

> Analytical judgements (affirmative) are therefore those in which the connection of the predicate with the subject is cogitated through identity; those in which this connection is cogitated without identity, are called synthetical judgements. The former may be called *explicative*, the latter *augmentative* judgements; because the former add in the predicate nothing to the conception of the subject, but only analyse it into its constituent conceptions, which were thought already in the subject, although in a confused manner; the latter add to our conceptions of the subject a predicate which was not contained in it, and which no analysis could ever have discovered therein. [...]

$$
\begin{array}{cc}
/ & \backslash \\
T\,A & F\,A \\
| & | \\
T\,A \vee \neg A & T\,\neg A \\
 & | \\
 & T\,A \vee \neg A
\end{array}
$$

FIG. 4: *An intelim tree of depth 1 proving the law of excluded middle.*

> In an analytical judgement I do not go beyond the given conception, in order to arrive at some decision respecting it. [...] But in synthetical judgements, I must go beyond the given conception, in order to cogitate, in relation with it, something quite different from what was cogitated in it [...][24]

One could say, by analogy, that analytical inferences are those which are recognized as sound via steps which are all "explicative", that is, descending immediately from the (informational) meaning of the logical operators, while synthetic ones are those that are "augmentative", involving some intuition that goes beyond this meaning, i.e., in our case involving the consideration of virtual information. So, we could paraphrase Kant and say that an inference is analytic only if it adds in the conclusion nothing to the information contained in the premises, but only analyses it in its constituent pieces of information, which were "thought already in the premises, although in a confused manner". The confusion vanishes once the meaning of the logical operators is properly explicated.

On the other hand, the synthetic inferences of classical propositional logic are precisely those that *essentially* require the introduction, via the branching rule PB, of virtual information. This is only the intuitive idea, the formal details are in (D'Agostino, 2005), (D'Agostino & Floridi, 2009) and (D'Agostino,

[24] I. Kant, *Critique of Pure Reason* [1781], Book II, Chapter II, Section II. Quoted from the english translation by J.M.D. Meiklejohn, ebooks Adelaide, 2009, `http://ebooks.adelaide.edu.au/k/kant/immanuel/k16p/index.html`.

2013). For each $n \in \mathbb{N}$, $\vdash_k$ is the consequence relation that allows for bounded applications of this branching rule up to a given fixed depth k. A classically valid inference can be said *synthetic at degree k* when k is the smallest natural number such that the inference in question is provable in $\vdash_k$ but not in $\vdash_{k-1}$. All classical tautologies are synthetic at some degree greater than 0. For example, the law of excluded middle is synthetic at degree 1, as shown by the argument in Figure 4, proving the law of excluded middle from the empty set of premises.

References

ANDERSON, A. R., & BELNAP JR, N.D. 1975. *Entailment: the Logic of Relevance and Necessity*. Vol. 1. Princeton: Princeton University Press.

ARTEMOV, S., & KUZNETS, R. 2009. Logical omniscience as a computational complexity problem. *Pages 14–23 of: TARK '09: Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*. New York, NY, USA: ACM.

BAR-HILLEL, Y., & CARNAP, R. 1953. Semantic Information. *British Journal for the Philosophy of Science*, **4**(14), 147–157.

BELNAP JR., N.D. 1976. How a computer should think. *Pages 30–55 of:* RYLE, G. (ed), *Contemporary Aspects of Philosophy*. Oriel Press.

BELNAP JR., N.D. 1977. A useful four-valued logic. *Pages 8–37 of:* DUNN, J. M., & EPSTEIN, G. (eds), *Modern uses of multiple-valued logics*. Dordrecht: Reidel.

BERGSTRÖM, L., & FØLLESDAG, D. 1994. Interview with Willard van Orman Quine, November 1993. *Theoria*, **60**, 193–206.

BLAMEY, S. 1986. Partial Logic. *Pages 1–70 of:* GABBAY, D.M., & GUENTHNER, F. (eds), *Handbook of Philosophical Logic*, vol. 3. Kluwer. Republished in the 2nd edition, volume 5, Kluwer, Dordrecht, 2002.

CADOLI, M., & SCHAERF, M. 1992. Approximate reasoning and non-omniscient agents. *Pages 169–183 of: TARK '92: Proceedings of the 4th* *conference on Theoretical aspects of reasoning about knowledge*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

CARAPEZZA, M., & D'AGOSTINO, M. 2010. Logic and the myth of the perfect language. *Logic and Philosophy of Science*, **8**(1), 1–29.

CARNAP, R., & BAR-HILLEL, Y. 1953. An Outline of a Theory of Semantic Information. *Pages 221–74 of:* BAR-HILLEL, Y. (ed), *Language and information*. Reading, Mass. and London: Addison-Wesley.

CHERNIAK, C. 1986. *Minimal Rationality*. MIT Press.

COHEN, M.R., & NAGEL, E. 1934. *An Introduction to Logic and Scientific Method*. Routledge and Kegan Paul.

COOK, S.A. 1971. The complexity of theorem-proving procedures. *Pages 151–158 of: STOC '71: Proceedings of the third annual ACM symposium on Theory of computing*. New York, NY, USA: ACM Press.

CRAWFORD, J.M., & ETHERINGTON, D.W. 1998. A Non-Deterministic Semantics for Tractable Inference. *Pages 286–291 of: AAAI/IAAI*.

D'AGOSTINO, M. 2005. Classical Natural Deduction. *Pages 429–468 of:* ARTËMOV, S.N., BARRINGER, H., D'AVILA GARCEZ, A.S., LAMB, L.C., & WOODS, J. (eds), *We Will Show Them! (1)*. College Publications.

D'AGOSTINO, M. 2010. Tractable Depth-Bounded Logics and the Problem of Logical Omniscience. *Pages 245–275 of:* MONTAGNA, F., & HOSNI, H. (eds), *Probability, Uncertainty and Rationality*. CRM series. Springer.

D'AGOSTINO, M. 2013. Semantic Information and the Trivialization of Logic: Floridi on the Scandal of Deduction. *Information*, **4**(1), 33–59.

D'AGOSTINO, M., & FLORIDI, L. 2009. The enduring scandal of deduction. Is propositionally logic really uninformative? *Synthese*, **167**, 271–315.

D'AGOSTINO, M., FINGER, M., & GABBAY, D.M. 2013. Semantics and proof-theory of Depth-Bounded Boolean Logics. *Theoretical Computer Science*, **480**, 43–68.

DALAL, M. 1996. Anytime families of tractable propositional reasoners. *In: Proceedings of the Fourth International Symposium on AI and Mathematics (AI/MATH-96), 42–45.*

DALAL, M. 1998. Anytime Families of Tractable Propositional Reasoners. *Annals of Mathematics and Artificial Intelligence*, **22**, 297–318.

DECOCK, L. 2006. *True by virtue of meaning. Carnap and Quine on the analytic-synthetic distinction.* `http://vu-nl.academia.edu/LievenDecock/Papers/866728/Carnap_and_Quine_on_some_analytic-synthetic_distinctions`.

DUMMETT, M. 1991. *The logical basis of metaphysics.* London: Duckworth.

DUŽÍ, M. 2010. The Paradox of Inference and the Non-Triviality of Analytic Information. *Journal of Philosophical Logic*, **39**, 473–510.

FAGIN, R., HALPERN, J.Y., & VARDI, M.Y. 1995. A nonstandard approach to the logical omniscience problem. *Artificial Intelligence*, **79**, 203–240.

FINGER, M. 2004a. Polynomial Approximations of Full Propositional Logic via Limited Bivalence. *Pages 526–538 of: 9th European Conference on Logics in Artificial Intelligence (JELIA 2004).* Lecture Notes in Artificial Intelligence, vol. 3229. Springer.

FINGER, M. 2004b. Towards polynomial approximations of full propositional logic. *Pages 11–20 of:* BAZZAN, A.L.C., & LABIDI, S. (eds), *XVII Brazilian Symposium on Artificial Intel ligence (SBIA 2004).* Lecture Notes in Artificial Intel lingence, vol. 3171. Springer.

FINGER, M., & GABBAY, D.M. 2006. Cut and Pay. *Journal of Logic, Language and Information*, **15**(3), 195–218.

FINGER, M., & WASSERMANN, R. 2004. Approximate and limited reasoning: Semantics, proof theory, expressivity and control. *Journal of Logic and Computation*, **14**(2), 179–204.

FINGER, M., & WASSERMANN, R. 2006. The universe of propositional approximations. *Theoretical Computer Science*, **355**(2), 153–166.

FLORIDI, L. 2004. Outline of a Theory of Strongly Semantic Information. *Minds and Machines*, **14**(2), 197–222.

FLORIDI, L. 2006. The Logic of Being Informed. *Logique et Analyse*, **49**(196), 433–460.

FLORIDI, L. 2011. *The Philosophy of Information*. Oxford: Oxford University Press.

FRISCH, A.M. 1987. Inference without chaining. *Pages 515–519 of: IJCAI'87: Proceedings of the 10th international joint conference on Artificial intelligence*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

HAHN, H., NEURATH, O., & CARNAP, R. 1973. The scientific conception of the world [1929]. *In:* NEURATH, M., & COHEN, R.S. (eds), *Empiricism and Sociology*. Dordrecht: Reidel.

HALPERN, J.Y. 1995. Reasoning About Knowledge: a Survey. *Pages 1–34 of:* GABBAY, D.M., HOGGER, C.J., & ROBINSON, J.A. (eds), *Handbook of Logic in Artificial Intelligence and Logic Programming*, vol. 4. Oxford: Clarendon Press.

HEMPEL, C.G. 1945. Geometry and Empirical Science. *American Mathematical Monthly*, **52**, 7–17.

HINTIKKA, J. 1973. *Logic, language games and information. Kantian themes in the philosophy of logic*. Oxford: Clarendon Press.

JAGO, M. 2012. The content of deduction. *The Journal of Philosophical Logic*. Forthcoming. Published online February 2012.

KANOVICH, M.I. 1992. Horn programming in linear logic is NP-complete. *Pages 200–210 of: Proceedings of the Seventh Annual Symposium on Logic in Computer Science*. IEEE Computer Science Press.

KRIPKE, S.A. 1965. Semantical analysis of intuitionistic logic I. *Pages 92–130 of:* CROSSLEY, J., & DUMMETT, M. A. E. (eds), *Formal Systems and Recursive Functions*. Amsterdam: North-Holland Publishing.

KWANG, M. S. 1997. Epistemic Logic and Logical Omniscience: a Survey. *International Journal of Intelligent Systems*, **12**(1), 57–81.

LADNER, R.E. 1977. The computational complexity of provability in systems of model propositional logic. *SIAM Journal of Computing*, **6**(3), 467–480.

LEVESQUE, H.J. 1988. Logic and the complexity of reasoning. *Journal of Philosophical Logic*, **17**(4), 355–389.

LINCOLN, P. 1995. Deciding probability of linear logic formulas. *Pages 197–210 of:* GIRARD, J.-Y, LAFONT, Y., & REGNIER, L. (eds), *Proceedings of the Workshop on Advances in Linear Logic*. Cambridge University Press.

LINCOLN, P., & WINKLER, T. 1994. Constant-Only Multiplicative Linear Logic is NP-Complete. *Theoretical Computer Science*, **135**, 155–169.

LINCOLN, P., MITCHELL, J., SCEDROV, A., & SHANKAR, N. 1992. Decision problems for propositional linear logic. *Annals of pure and applied logic*, **56**, 239–311.

MASSACCI, F. 1998. *Efficient Approximate Deduction and an Application to Computer Security*. Ph.D. thesis, Università degli Studi di Roma "La Sapienza".

MEYER, J.C. 2003. Modal Epistemic and Doxastic Logic. *Pages 1–38 of:* GABBAY, DOV M., & GUENTHNER, FRANZ (eds), *Handbook of Philosophical Logic*, 2nd edn., vol. 10. Kluwer Academic Publishers.

MONDADORI, M. 1989. *An Improvement of Jeffrey's Deductive Trees*. Annali dell'Università di Ferrara; Sez. III; Discussion paper 7. Università di Ferrara.

PARIKH, R. 2008. Sentences, belief and logical omniscience, or what does deduction tell us? *The Review of Symbolic Logic*, **1**(4).

POPPER, K.R. 1934. *Logik Der Forschung : Zur Erkenntnistheorie Der Modernen Naturwissenschaft*. Wien: J. Springer. Eng. tr. The Logic of Scientific Discovery (London: Hutchinson, 1959).

POPPER, K.R. 1959. *The Logic of Scientific Discovery*. London: Hutchinson. Eng. tr. of Popper (1934).

PRAWITZ, D. 2006. Meaning approached via proofs. *Synthese*, **148**(3), 507–524.

PRIMIERO, G. 2008. *Information and Knowledge. A Constructive Type-Theoretical Approach*. Springer.

PRIMIERO, G. 2009. An epistemic logic for being informed. *Synthese*, **167**(2), 363–389.

QUINE, W.V.O. 1951. Two dogmas of empiricism. *The Philosophical Review*, **60**, 20–43. Reprinted in W.V.O. Quine, *From a Logical Point of View*, Harvard University Press, Cambridge 1953, 2nd 1961.

QUINE, W.V.O. 1961. Two dogmas of empiricism. *Pages 20–46 of: From a Logical Point of View*. Cambridge, MA: Harvard University Press. First published in *Philosophical Review*, 60:20–43, 1951.

QUINE, W.V.O. 1973. *The Roots of Reference*. Open Court.

RAATIKAINEN, P. 2004. Conceptions of truth in intuitionism. *History and Philosophy of Logic*, **25**, 131–145.

SAVAGE, L. 1967. Difficulties in the theory of personal probability. *Philosophy of Science*, **44**, 305–310.

SEQUOIAH-GRAYSON, S. 2008. The Scandal of Deduction. Hintikka on the Information Yield of Deductive Inferences. *The Journal of Philosophical Logic*, **37**(1), 67–94.

SHANNON, C.E., & WEAVER, W. 1949. *The Mathematical Theory of Communication*. Urbana: University of Illinois Press. Foreword by R.E. Blahut and B. Hajek.

SHEERAN, M., & STÅLMARCK, G. 2000. A Tutorial on Stalmarck's Proof Procedure for Propositional Logic. *Formal Methods in System Design*, **16**, 23–58.

SILLARI, G. 2008a. Models of Awareness. *Pages 209–240 of:* BONANNO, G., VAN DER HOEK, W., & WOOLDRIDGE, M. (eds), *Logic and the Foundations of Games and Decisions*. University of Amsterdam.

SILLARI, G. 2008b. Quantified Logic of Awareness and Impossible Possible Worlds. *Review of Symbolic Logic*, **1**(4), 1–16.

STATMAN, R. 1979. Intuitionistic Propositional Logic is Polynomial-Space Complete. *Theoretical Computer Science*, **9**, 67–72.

STOCKMEYER, L. 1987. Classifying the computational complexity of problems. *Journal of Symbolic Logic*, **52**(1), 1–43.

SVEJDAR, V. 2003. On the polynomial-space completeness of intuitionistic propositional logic. *Archive for Mathematical Logic*, **42**(7), 711–716.

URQUHART, A. 1984. The undecidability of entailment and relevant implication. *Journal of Symbolic Logic*, **49**(4), 1059–1073.

URQUHART, A. 1990. The complexity of decision procedures in relevance logic. *Pages 61–76 of:* E A. GUPTA, J.M. DUNN (ed), *Truth or Consequences*. Kluwer Academic Publishers.

WEIR, A. 1986. Classical Harmony. *Notre Dame Journal of Formal Logic*, **27**(4), 459–482.