

Article

Semantic Information and the Trivialization of Logic: Floridi on the Scandal of Deduction

Marcello D’Agostino

Department of Economics and Management, University of Ferrara, Ferrara, 44121 Italy;

E-Mail: dgm@unife.it

Received: 19 November 2012 / Accepted: 27 December 2012 / Published: 11 January 2013

Abstract: In this paper we discuss Floridi’s views concerning semantic information in the light of a recent contribution (in collaboration with the present author) [1] that defies the traditional view of deductive reasoning as “analytic” or “tautological” and construes it as an informative, albeit non-empirical, activity. We argue that this conception paves the way for a more realistic notion of semantic information where the “ideal agents” that are assumed by the standard view can be indefinitely approximated by real ones equipped with growing computational resources.

Keywords: semantic information; Cohen–Nagel paradox; tractability; analytic-synthetic distinction

1. Introduction

“Philosophical work on the concept of [...] information is still at that lamentable stage when disagreement affects even the way in which the problems themselves are provisionally phrased and framed” [2]. Thus Luciano Floridi, in his entry on “Semantic conceptions of information” in the *Stanford Encyclopedia of Philosophy*, takes a snapshot of the state of the art in the philosophy of information A.D. 2005. As early as 1953, a few years after the appearance of his “Mathematical Theory of Communication”, Claude Shannon had warned that the new area of Information Theory was surrounded by some conceptual and terminological confusion that prompted for further clarification and differentiation:

The word “information” has been given many different meanings by various writers in the field of information theory. It is likely that at least a number of these will prove sufficiently useful in certain applications to deserve further study and permanent recognition. It is hardly

to be expected that a single concept of information would satisfactorily account for the numerous possible applications of this general field [3].

Half a century later, Floridi's main contribution has probably been that of taking up Shannon's challenge from a philosophical perspective, by engaging in the bold project of cleaning up the area and taming the "notoriously polymorphic and polysemantic" [2], although increasingly pervasive, notion of information. This ambitious program has been carried out to a considerable extent in a series of papers and, above all, in his main work on this subject, *The Philosophy of Information* [4].

In this paper we discuss Floridi's views concerning *semantic* information in the light of a recent contribution (in collaboration with the present author) [1] that defies the traditional view of deductive reasoning as "analytic" or "tautological" and construes it as an informative, albeit non-empirical, activity. We argue that this conception paves the way for a more realistic notion of semantic information where the "ideal agents" that need to be assumed to defend the standard view can be indefinitely approximated by real ones equipped with growing computational resources. We also argue that this conception offers a partial vindication of the Kantian notion of "synthetic a priori" even in the allegedly trivial domain of propositional logic. (A similar vindication of Kantian ideas, but restricted to the notoriously harder domain of quantification logic, is the leitmotiv of Hintikka's well-known book *Logic, Language Games and Information* [5].)

We start, in Section 2, by presenting the "received view" that logical inference is tautological, in the sense that it does not increase information. In this view, the conclusion can be seen to be true, given the truth of the premises, by the very meaning of the logical operators and the corresponding judgment that the conclusion follows from the premises is therefore "analytic", a purely linguistic truth that we learn by learning the language. Accordingly, the "semantic information" carried by the conclusion—that is, the information that it carries by virtue of the meaning of the logical operators—must be contained in the information carried by the premises. In Section 3, we discuss various anomalies of the received view—the Bar-Hillel–Carnap paradox, the "scandal of deduction", the problem of logical omniscience—and, in Section 4, Floridi's approach to them as it emerges from [6–8]. In Section 5 we present our new approach based on an informational semantics for the logical operators. This consists in fixing their meaning in terms of the information that is *actually possessed* by an agent. The main idea is that their classical meaning, being based on information-transcending notions, such as classical truth and falsity, is not apt to justify the claim that inferences that are "analytic"—*i.e.*, licensed by this very meaning—are also "tautological", *i.e.*, informationally trivial. By contrast, the central semantic notions of informational semantics are those of "actually possessing the information" that a given sentence is true, respectively false, and the meaning of the logical operators is defined exclusively in terms of these notions. As a consequence, inferences that are "analytic" according to the informational meaning of the operators turn out to be also "tautological" in a particularly strict sense—that an agent that actually possesses the information carried by the premises actually possesses also the information carried by the conclusion—and this is confirmed by the fact that the corresponding consequence relation is computationally feasible. Non-analytic (or "synthetic") inferences, on the other hand, are characterized by the fact that they essentially require the use of "virtual information", in the form of provisional hypothetical information that is not actually possessed by the agent who makes the inference, such as the one that plays a crucial role in the "discharge rules" of Gentzen-style Natural

Deduction (see [9] for an excellent introduction). Gradually allowing for the nested use of such virtual information naturally leads to define a hierarchy of increasingly informative inferential systems that indefinitely approximate classical propositional logic.

2. The Received View

According to the received view, logical deduction never increases (semantic) information. This tenet clashes with the intuitive idea that deductive arguments are useful just because, by their means, we obtain information that we did not possess before. However, it allowed philosophers and mathematicians to justify their view of logic and mathematics as infallible activities that are not subject to the tribunal of experience.

2.1. Logical Empiricism and the Trivialization of Logic

One of the trademarks of modern, or “logical”, empiricism was the rejection of the possibility of “synthetic a priori” judgements. (We briefly recall some philosophical terminology. A judgment is *analytic* when it does not extend our knowledge, but asserts something that is trivially true solely by the meaning of the words, e.g., “all bachelors are unmarried”. A *synthetic* judgment is one that does extend our knowledge and cannot be established by mere semantic analysis, such as “all bachelors eat in front of tv”. A judgment is *a priori* when its truth does not depend on experience, and *a posteriori* when it does. Analytic judgments are always a priori. Synthetic judgments are normally a posteriori, but Kant argued that some mathematical judgments are synthetic a priori: they do extend our knowledge but do not depend on experience.) This position was clearly stated in the *manifesto* of the Vienna Circle as “the basic thesis of modern empiricism”:

In such a way logical analysis overcomes not only metaphysics in the proper, classical sense of the word, especially scholastic metaphysics and that of the systems of German idealism, but also the hidden metaphysics of Kantian and modern apriorism. The scientific world-conception knows no unconditionally valid knowledge derived from pure reason, no “synthetic judgments a priori” of the kind that lie at the basis of Kantian epistemology and even more of all pre- and post-Kantian ontology and metaphysics. [...] 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 [10] (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”. More explicitly, they thought 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 [10] (p. 311).

This view of deductive reasoning as informationally void is usually supported by resorting to elementary examples, as in this well-known passage by 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 [11] (p. 9).

Despite its highly counterintuitive implications—at least in the ordinary sense of the word “information”, it is hard to accept that all the mathematician’s efforts never go “one iota” beyond the information that was already contained in the axioms of a mathematical theory—this view of deductive reasoning caught on and became part of the logical folklore. Most of its philosophical appeal probably lies in the fact that it appears to offer 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. However, 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 [12] (p. 195).

Indeed—as we shall argue in Section 3—this trivialization of logic is a philosophical overkill: a definitive foundation for deductive practice is obtained at the price of its informativeness. Logic lies on a bedrock of platitude.

2.2. Quine on Logical Truth

The conception of logical deduction as “analytic”, and therefore “tautological”, is a persistent dogma of (logical) empiricism that seems to be somewhat independent of Quine’s two dogmas [13] as well as from Davidson’s “third dogma” [14]. 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 [13], 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 was 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” ([13], pp. 22–32 of the 1961 edition). Four decades later, while his reservations over the notion of analyticity remained the “the same as ever”, Quine clarified that they concerned only “the tracing of any demarcation, even a vague and approximate one, across the domain of sentences in general” [15] (p. 270). But the impossibility of tracing a sharp demarcation does not exclude that there may be undebatable cases of analytic sentences. Indeed, “It is intelligible and often useful in discussion to point out that some disagreement is purely a matter of words rather than of fact” [15] (p. 270). The so-called “logical laws” are the most natural candidates for such paradigmatic examples of analytic sentences: it seems almost uncontroversial that a disagreement about a logical truth can always be reduced to a disagreement about the meaning of some logical word that occurs in it.

In fact, in *The Roots of Reference* Quine had already suggested that, in order to fit the undisputed cases of analytic sentences, one may provide a rough theoretical definition of analyticity by saying that (i) a sentence is analytic for the native speakers of a language if they learn its truth in the very process of learning how to use the words that occur in it; and (ii) “recondite” sentences should still count as analytic if they can be obtained by “a chain of inferences each of which individually is assured by the learning of the words” [16] (pp. 79–80). In this perspective, logical truths may qualify as analytic in the traditional sense, although the very existence of enduring disagreement on some logical laws—e.g., on the law of excluded middle on the part of intuitionists—may suggest that such laws are not similarly bound up with the learning of the logical words and “should perhaps be seen as synthetic” [16] (p. 80).

In his latest work Quine appears to leave aside this idea that some logical laws may be synthetic. For example, in his *Two dogmas in retrospect*, he argues that by the above criterion “all logical truths [...]—that is, the logic of truth functions, quantification, and identity—would then perhaps qualify as analytic, in view of Gödel’s completeness proof” [15] (p. 270) and later on, in a 1993 interview, seems to abandon any hesitation and make 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 [17] (p. 199). (Quoted in [18].)

2.3. Semantic Information

At the half of the 20th century, Bar-Hillel and Carnap’s theory of “semantic information” provided what is, to date, the strongest theoretical justification for the thesis that deductive reasoning is “tautological”. Although their effort was clearly inspired by the rising enthusiasm for Shannon and Weaver’s new Theory of Information [19], 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 role 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 symbolise, 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 [20].

By way of 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” [20] (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:

$$R \wedge W \quad R \wedge \neg W \quad \neg R \wedge W \quad \neg R \wedge \neg W$$

Now, the sentence “tomorrow will rain and will be windy” is intuitively more informative than the sentence “tomorrow 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 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 role 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.) [21] (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 [21] (p. 96).

3. The Anomalies of the Received View

A straightforward consequence of Bar-Hillel and Carnap’s notion of “semantic information” is that contradictions, like “tomorrow will rain and will not rain”, carry the maximum amount of information, since they exclude all possible states. Another inevitable consequence of the theory is that *all* logical truths are equally uninformative (they exclude no possible world), which justifies their being labelled as “tautologies”. But in classical logic a sentence φ is deducible from a finite set of premises $\psi_1, \dots, \psi_n$ if and only if the conditional $(\psi_1 \wedge \dots \wedge \psi_n) \rightarrow \varphi$ 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 (fourth?) empiricist dogma discussed in Section 2.1, both these consequences appear to be at odds with our intuitions and clash with the commonsense notion of information, to the extent that some authors have described them as true “paradoxes”.

3.1. The Bar-Hillel–Carnap Paradox

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 [22]:

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 [22] (p. 229).

Popper had also realized that his closely related notion of empirical content worked reasonably well only for *consistent* theories, since *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 role 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” [21] (p. 72). So, “whilst tautologies, purely existential statements and other nonfalsifiable 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” [21] (p. 71). In fact, what Popper claimed was that the information content of inconsistent theories is null, and so his definition of empirical 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 [21] (p. 72).

3.2. *The 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 [23] (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? [5] (p. 222).

The standard answer to this question has a strong psychologistic flavour. According to Hempel: “a mathematical theorem, such as the Pythagorean theorem in geometry, asserts nothing that is objectively or theoretically new as compared with the postulates from which it is derived, although its content may well be *psychologically new* in the sense that we were not aware of its being implicitly contained in the postulates” ([11] (p. 9), Hempel’s emphasis.) This implies that there is no objective (non-psychological) sense in which deductive inference yield new information. Hintikka’s reaction to this typical neopositivistic way out of the paradox is worth quoting in full:

If no objective, non-psychological increase of information takes place in deduction, all that is involved is merely psychological conditioning, some sort of intellectual psychoanalysis,

calculated to bring us to see better and without inhibitions what objectively speaking is already before your eyes. Now most philosophers have not taken to the idea that philosophical activity is a species of brainwashing. They are scarcely any more favourably disposed towards the much more far-fetched idea that all the multifarious activities of a contemporary logician or mathematician that hinge on deductive inference are as many therapeutic exercises calculated to ease the psychological blocks and mental cramps that initially prevented us from being, in the words of one of these candid positivists, “aware of all that we implicitly asserted” already in the premises of the deductive inference in question [5] (pp. 222–223).

3.3. Wittgenstein and the “Perfect Notation”

A non-psychologistic attempt to avoid the paradox consists in blaming it on the imperfection of our logical language. In his *Tractatus*, 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 proposition 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, 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 that 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. (On the theme of a “logically perfect language” see also [24].)

3.4. Hintikka on the Scandal of Deduction

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-Hillel 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 [5] (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 (for a criticism of Hintikka’s approach see [25]), we observe here that Hintikka’s proposal classifies as non-analytic only some inferences of the non-monadic predicate calculus and 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 premises. 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” ([5] (p. 154)).

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

$$\text{If } \Gamma \vdash \varphi \text{ the information carried by } \varphi \text{ is included in the information carried by } \Gamma \quad (1)$$

This is highly unsatisfactory, especially since the theory of computational complexity has revealed that the decision problem for Boolean logic is co-NP-complete [26], that is, among the hardest problems in co-NP. Although not a proved theorem, it is a widely accepted conjecture that Boolean logic is *practically undecidable*, i.e., admits of no feasible decision procedure. (This means that every decision procedure for Boolean logic is bound to be superpolynomial in the *worst case*. On the other hand, there are very efficient decision algorithms around that work quite efficiently on average. In [27] Finger and Reis present a very interesting empirical analysis of the runtime distribution of a variety of decision methods on randomly generated formulas.) To express the same idea in a different way, we could say that there cannot be any “perfect” propositional language, in Wittgenstein’s sense—one in which the logical relations between sentences can be recognized by mere inspection of the sentences themselves—into which a conventional logical language can be feasibly translated. (On this point see [24].)

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 reducing uncertainty, we should conclude that in some cases deductive reasoning *does* reduce our uncertainty, and therefore increases our information, even at the propositional level.

The scandal of deduction has recently received renewed attention leading to a number of original contributions (e.g., [28] (Chapter 2), [1,25,29–31] that do not appear, however, to be reducible to a single conceptual paradigm.

3.5. The Problem of Logical Omniscience

Another widely debated paradox connected with the received view on logical deduction arises in the context of modal characterizations of propositional attitudes and is nothing but a variant of the “scandal of deduction” described in the previous section. According to the standard logic 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), if an agent a knows (or believes, or is informed) that a sentence φ is true, and ψ is a logical consequence of φ , then a is supposed to know (or believe, or be informed) also that ψ is true. (For a survey on epistemic and doxastic logic see [32,33]; for information logic, or “the logic of being informed”, see [34,35].) 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,

$$\text{if } \Box_a \varphi \text{ for all } \varphi \in \Gamma \text{ and } \Gamma \vdash \psi, \text{ then } \Box_a \psi \quad (2)$$

where $\vdash$ stands for the relation of logical consequence. Observe that, letting $\Gamma = \emptyset$, it immediately follows from (2) 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 (2) emerges from the combined effect of the “distribution axiom”, namely,

$$(K) \quad \Box_a(\varphi \rightarrow \psi) \rightarrow (\Box_a\varphi \rightarrow \Box_a\psi)$$

and the “necessitation rule”:

$$(N) \quad \text{if } \vdash \varphi, \text{ then } \vdash \Box_a\varphi.$$

On the other hand, despite its paradoxical flavour, (2) 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 $(S, \tau, R_1, \dots, R_n)$, 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 $\models$. This incorporates the usual semantics of classical propositional logic and defines the truth of $\Box_a\varphi$ as “ φ 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, (2) 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, (2) may perhaps be satisfied by an “idealized reasoner”, in some sense to be made more precise, but is not satisfied, and is not likely to ever be satisfiable, in practice. (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.) As mentioned above, 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, and this means that any real agent, even if equipped with an up-to-date computer running a decision procedure for Boolean logic, will never be able to feasibly recognize that certain Boolean sentences logically follow from sentences that she regards as true. So, the clash between (2) 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 (2), 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 practical reasoner.

4. Floridi on the Received View

In this section we discuss Floridi’s ideas on the anomalies of the received view. The reader is warned that our exposition shows a strong bias for the ideas put forward in a joint paper by Floridi and the present author [1]. However, we hope that, as often is the case, this bias may have a heuristic value.

4.1. Floridi on the BCP

Floridi's key idea for dissolving the Bar-Hillel–Carnap paradox (BCP) is startlingly simple. Observe that the problem with the standard theory of semantic information is confined to inconsistent (sets of) sentences. So, if we endorse the view that “information encapsulates truth” [6], *i.e.*, that semantic information is not only “well-formed and meaningful data”, but must also be *truthful* [6,7], the paradox obviously vanishes. Since inconsistent sentences cannot be true, they cannot qualify as informative. To Floridi, “false information” is nearly an oxymoron: we should rather speak of *misinformation*. Once we realize that there is no such a thing as false information, clearly the BCP is still a problem for the standard theory of semantic information, renamed Theory of Weakly Semantic Information, where “truth values supervene on information”, but it is no longer a problem for Floridi's quantitative Theory of Strongly Semantic Information [6], [8] (Chapter 4), where information encapsulates truth.

We shall not discuss TSSI in detail, since we chose to concentrate on Floridi's approach to the scandal of deduction. We just remark, in passing, that despite all its merits, this theory may perhaps be criticized, very much like the neopositivistic tenet that logic is informationally trivial, as a philosophical overkill. If the problem lies with “inconsistent information”, why not simply say that “information encapsulates consistency” and try to define a notion of semantic information where inconsistent (sets of) sentences are qualified as uninformative, in line with Popper's informal view (Section 2.3 above)? Moreover, it can be argued that the “veridicality thesis” unduly imports a strong metaphysical commitment into Information Theory and that the latter would better remain as metaphysically neutral as possible, especially if we want to use it to solve philosophical controversies. (However, a similar objection could be raised as well against the Theory of Weakly Semantic Information where the classical notions of truth and falsity convey, via the tacit assumption of the principle of bivalence, a strong metaphysical commitment.) Finally, it seems plausible that no agent can, in general, practically distinguish a situation in which he holds genuine information from a situation in which he has been misinformed. We tend to agree, with Popper and the more sophisticated neopositivists such as Carnap and Neurath, that “statements can be logically justified only by other statements”[21] (p. 21), that perceptual statements recording experiences are not irrevocable, and that “experiences can *motivate a decision*, and hence an acceptance or a rejection of a statement, but a basic statements cannot be *justified* by them—no more than thumping the table” [21] (pp. 87–88). If we accept this view—interestingly enough, Popper describes it as “closer to the critical (Kantian) school of philosophy than to positivism”[21] (p. 88, footnote 3)—we cannot help recognizing that whether or not a set of statements provides, as a whole, genuine information ultimately depends on its consistency.

It must be observed however that even requiring that “information encapsulates consistency” appears too demanding. In many interesting contexts agents are not able to tell whether their data are (classically) consistent. So how can we practically distinguish a situation in which we hold genuine information (*i.e.*, well-formed, meaningful and consistent data) from a situation in which we do not because our data is inconsistent, but we have no practical means of detecting this inconsistency? It may be retorted that requiring that information encapsulates consistency leaves us at least with a *negative* criterion that is epistemologically robust: once a set Γ of statements is shown to be inconsistent, it definitively qualifies as misinformation and this judgement cannot be overthrown unless Γ is revised. However, if we want

also a positive criterion we have to weaken even the consistency requirement. We shall return to this point later on, in Section 5.2, where we shall make a positive proposal.

4.2. Floridi on the Scandal of Deduction

Floridi's Theory of Strongly Semantic Information implies that the degree of informativeness of any tautology is 0 (see [8] (Chapter 5, §6)), so leaving the paradox unsolved. On some occasions Floridi seems to endorse the idea that there is no scandal at all in claiming that logical truths are uninformative:

[...] indeed, according to MTI [Mathematical Theory of Information], TWSI [Theory of Weakly Semantic Information] and TSSI [Theory of Strongly Semantic Information], tautologies are not informative. This seems both reasonable and unquestionable. If you wish to know what the time is, and you are told that “it is either 5pm or it is not” then the message you have received provides you with no information [8] (p. 169).

Accordingly, in his *The logic of being informed* [34] Floridi chooses to downplay the related “problem of logical omniscience”, there renamed as the problem of “information overload”. He argues that the notion of “being informed” can be adequately captured by a normal modal logic, that he calls *Information Logic* (IL)—where $\Box_a$ is renamed I_a and the sentence $I_a\varphi$ is interpreted as “the agent a is informed (holds the information) that φ ”—and that **K** can be safely taken as an axiom of this logic. However, as already observed in Section 3.5, if combined with the “inevitable inclusion of the rule of necessitation” [34] (p. 18), that is **N**, the axiom **K** implies (2), namely, that “being informed” is preserved under logical entailment. So, endorsing **K** and **N**, under this interpretation of I_a , implies that logical truths and logical deductions are utterly uninformative: whenever $\psi \vdash \varphi$, any agent who is informed (holds the information) that ψ must also be informed (hold the information) that φ . This is tantamount to saying that logic is informationally trivial and that the information an agent obtains from the conclusion of a deductive argument, no matter how complex, is already included in the information she obtains from the premises. So the problem of information overload (aka. “logical omniscience”) and the scandal of deduction are two sides of the same coin.

In § 3.1 of [34], Floridi admits the problem and lists three possible strategies to mitigate it. First, one may claim that in Information Logic, as in any epistemic logic, the rule of necessitation describes only an *ideal* agent, equipped with unlimited computational resources. This is a typical and quite popular philosophical escape to which it may be retorted, with Gabbay and Woods, that

A logic is an idealization of certain sorts of real-life phenomena. By their very nature, idealizations misdescribe the behaviour of actual agents. This is to be tolerated when two conditions are met. One is that the actual behaviour of actual agents can defensibly be made out to approximate to the behaviour of the ideal agents of the logician's idealization. The other is the idealization's facilitation of the logician's discovery and demonstration of deep laws [36] (p. 158).

In the context of the problems discussed in this paper, the first condition is the crucial one: if the metaphor of the “ideal agent” is to be useful at all, we need to associate it with a theory of how the actual logical

behaviour of real agents approximates the theoretical behaviour of idealized agents. It does not seem, however, that this condition can be met by IL, at least in its current formulation.

Floridi's second strategy consists simply in observing that the problem is shared by all epistemic logics. While this is admittedly no solution, one can agree with him that a "a problem shared is a problem halved" [34] (p. 18) and, more importantly, that "any argument usable to limit the damage of cognitive overload in those logics [...] can be adapted to try IL as well". However, rather than just "limiting the damage", it would be interesting and useful to identify the conceptual glitch that is responsible for it. An effort in this direction is made in [1] of which we shall say more later on.

Interestingly enough, Floridi's third and last strategy seems to consist in removing the scandal of deduction altogether, by uncritically assuming that all tautologies are equally uninformative. Floridi's argument can be summarized as follows:

1. By the Inverse Relationship Principle, "information goes hand in hand with unpredictability" [34] (p. 19); since every tautology has probability 1, if φ is a tautology, then φ is completely uninformative.
2. Hence, an agent's information cannot be increased by receiving the (empty) information that φ .
3. This situation is indistinguishable from the one in which the agent actually holds the (empty) information that φ . In other words:

If you ask me when the train leaves and I tell you that either it does or it does not leave at 10:30 am, you have not been informed, although one may indifferently express this by saying that what I said was uninformative in itself or that (it was so because) you already were informed that the train did or did not leave at 10:30 am anyway [34] (p. 19).

So, we can consider " a holds the information that φ " as synonymous with " a 's information is not increased by receiving the information that φ ".

4. Hence, we can assume that, for every tautology φ , a holds the information that φ .

In conclusion:

It turns out that the apparent difficulty of information overload can be defused by interpreting $\vdash \varphi \implies \vdash I\varphi$ as an abbreviation for

$$\vdash \varphi \implies P(\varphi) = 1 \implies \text{Inf}(\varphi) = 0 \implies \vdash I\varphi$$

which does not mean that a is actually informed about all theorems provable in PC [Propositional Calculus] as well as in KTB-IL [Information Logic]—as if a contained a gigantic database with a lookup table of all such theorems—but that, much more intuitively, any theorem φ provable in PC or in KTB-IL (indeed, any φ that is true in all possible worlds) is uninformative for a [34] (p. 19).

It may be objected that, if $P \neq NP$, there will always be infinitely many tautologies that the agent a is not able to recognize within feasible time. More precisely, there will be infinite classes of tautologies that all the procedures available to a are not able to recognize in time bounded above by a polynomial in the size of the input. In practice, this means that, no matter how large a 's computational resources are, there will always be tautologies in one of these classes that go far beyond them. Floridi's argument seems to assume that even if φ is one of such tautologies that are "hard" for a , φ should be regarded as uninformative for a , although a has no feasible means of recognizing φ as a tautology. However, under these circumstances, a has no feasible means to distinguish this situation from an analogous situation in which φ is *not* a tautology and so a 's information has actually been increased from learning that φ . Following Floridi's line of argument, we should therefore say that, whenever φ is a tautology, a always holds, in some "objective" sense, the information that φ , although this situation is not practically distinguishable from a situation in which φ is not a tautology and a does not hold the information that φ . A similar problem arises, perhaps even more strikingly, when a valid inference of φ from Γ is "hard" for the agent a . Again, the conditional probability $P(\varphi \mid \Gamma)$ is equal to 1 and so φ adds nothing to the information carried by Γ . However, it may make a lot of *practical* difference for a to *actually* hold the information that φ indeed follows from Γ , and this may significantly affect his decisions and his overall behaviour.

Of course one might retort that, under these circumstances, a 's subjective judgement is wrong and that, objectively speaking, a is not really receiving any information at all. This may be made more palatable by resorting to the usual (and somewhat abused) philosophical trick of the "ideal agent" equipped with boundless resources, and by claiming that this is, after all, a "useful idealization" (p. 19). Such an ideal agent would always be able to recognize that φ is a tautology, whenever it is, and this may be taken to dissolve the difficulty outlined above. However, this would bring us back to where we had been left with the first strategy.

Perhaps times are ripe to ask the fundamental question: is this kind of metaphysical and ultimately unattainable "objective information"—as opposed to the information a (real) agent *actually* holds—the only possible subject matter of a theory of semantic information? Should the philosophical consolation of the "ideal agent" keep hiding the fact that the Theory of Semantic Information, in all its known variants, does not (yet) account for at least one important use of the word "information", perhaps the one that is most relevant in practice?

This is the sense in which, for example, an agent a may not hold the information that in a hard sudoku game a certain cell must be occupied by a certain digit, even if this follows from her initial information about the game by propositional logic only (as is indeed the case for all sudoku games). Would it be irrational for a to decline a bet that pays her 1 euro if she correctly identifies this digit and −1000000 euros if she does not? According to the current notion(s) of semantic information, it would. According to a more realistic notion, it clearly would not, unless a 's computational resources are sufficient to determine the digit that must fill the given cell within the time she is given to make her decision. It seems, though, that this commonsense notion of information still escapes any attempt to force it in the Procrustean bed of current theoretical accounts and that, for this purpose, the worn metaphor of the "ideal agent" is not that helpful, after all.

What we have in mind is a notion of information that satisfies the following commonsense requirement:

Strong Manifestability. If an agent a grasps the meaning of a sentence φ , then a should be able to tell, in practice and not only in principle, whether or not (s)he holds the information that φ is true, or the information that φ is false or neither of them.

Given the obvious interplay between meaning-theories—in our context theories on the meaning of the logical operators—and theories of semantic information, the above requirement is, at the same time, a requirement on the relevant notion of semantic information and on the meaning-theories that may be sensibly associated with this notion. One may well maintain, as we do, that it makes perfectly good sense to regard a sentence as uninformative for an agent a if it follows “analytically”, *i.e.*, by virtue only of the meaning of the logical operators, from the information that a actually holds. But, in the light of the Strong Manifestability Requirement (SMR), this meaning cannot be their classical meaning and the residual notion of semantic information cannot be any of the current ones, because both the standard way of fixing the meaning of the classical logical operators and the corresponding notion(s) of semantic information do not satisfy SMR. So, SMR constrains both the way in which the meaning of the logical operators is fixed and the way in which the notion of semantic information is characterized.

In our view, the scandal of deduction and the related problem of information overload are nothing but symptoms of a fundamental difficulty. This can be described as *the mismatch between the central semantic notions in terms of which the meaning of the logical operators is defined and the commonsense notion of information that underlies SMR*. The classical meaning of the logical operators, as defined by the standard truth-tables, is specified in terms of alethic notions of truth and falsity that are obviously information-transcendent. What we need is a meaning-theory whose central semantic notions are *themselves* of an informational nature. In this theory, the meaning of a complex sentence for an arbitrary agent a should not be specified—as in the truth-tables—in terms of the alethic notions of truth and falsity, but solely in terms of the information that the agent actually holds. More specifically, rather than the usual necessary and sufficient conditions for the truth/falsity of a complex sentence φ expressed in terms of the truth/falsity of its immediate components, we need necessary and sufficient conditions for an agent’s actually holding the information that φ is true/false in terms of the information actually held by the agent about its immediate components.

5. A Kantian View?

The problem raised at the end of the previous section is addressed in [1] and [38] via what we have called “the informational meaning of the logical operators”. In this section we explain how this alternative meaning-theory may provide an effective solution to the scandal of deduction and pave the way for an alternative notion of semantic information that satisfy SMR. According to this theory, only ideal agents are actually informed of the truth of all tautologies and of the validity of all classical inferences. Real agents are informed only of a well-defined subclass of tautologies and inferences that depend on their reasoning resources.

5.1. Informational Semantics

The informational semantics of the logical operators is based on the following principle:

Informational semantics. The meaning of an n -ary logical operator $\star$ is determined by specifying the necessary and sufficient conditions for an agent a to actually hold the information that a sentence of the form $\star(\varphi_1, \dots, \varphi_n)$ is true, respectively false, in terms of the information that a actually holds about the truth or falsity of $\varphi_1, \dots, \varphi_n$.

Here by saying that a *actually holds* the information that φ is true (respectively false) we mean that a possesses a feasible procedure to obtain this information.

Clearly, the classical truth-table semantics for the Boolean operators does not qualify as an informational semantics. Even if we interpret the truth values 1 and 0 that a sentence φ may assume in informational terms—that is, by stipulating that φ takes the value 1 when we actually hold the information that φ is true, and the value 0 when we actually hold the information that φ is false—this is by no means sufficient to turn the classical semantics into an informational one. A first obvious problem is related to the underlying *principle of bivalence*, according to which every sentence φ must take one of the two values 1 and 0. Under the informational re-interpretation of the truth-tables, this principle is turned into a *principle of omniscience*: for every sentence φ either we actually hold the information that φ is true or we actually hold the information that φ is false. This may be true for ideal Laplacian agents, but is obviously absurd for non-ideal ones. This problem might be circumvented by dropping the principle of bivalence and introducing a third “indeterminate” value, say $1/2$, but leaving the necessary and sufficient conditions determined by the truth-tables unchanged for the two determinate values. This extended semantics for the Boolean operators is shown in Table 1. Under the standard definition of logical consequence (Γ entails φ if and only if the value of φ is designated for all valuations that give a designated value to all the sentences in Γ) this semantics yields Kleene’s 3-valued logic when 1 is taken as the only designated value, but collapses into classical logic again if 1 and $1/2$ are both taken as designated. However, not even this 3-valued semantics qualifies as an informational semantics in our sense.

Table 1. 3-valued truth-tables.

$\wedge$	1	0	$\perp$	$\vee$	1	0	$\perp$	$\rightarrow$	1	0	$\perp$	$\neg$	
1	1	0	$\perp$	1	1	1	1	1	1	0	$\perp$	1	0
0	0	0	0	0	1	0	$\perp$	0	1	1	1	0	1
$\perp$	$\perp$	0	$\perp$	$\perp$	1	$\perp$	$\perp$	$\perp$	1	$\perp$	$\perp$	$\perp$	$\perp$

The problem becomes apparent when one considers the necessary and sufficient condition for the truth of a disjunction or the necessary and sufficient condition for the falsity of a conjunction. Under the informational interpretation of 1 and 0, these conditions would read, respectively, as follows:

- we actually hold the information that $\psi \vee \varphi$ is true if and only if we actually hold the information that ψ is true or we actually hold the information that φ is true;

- we actually hold the information that $\psi \wedge \varphi$ is false if and only if we actually hold the information that ψ is false or we actually hold the information that φ is false.

While these clauses are sound in the “if” direction, they are obviously unsound in the “only-if” direction, at least under the commonsense notion of “holding the information”. It is sensible to say “we hold the information that the sentence ‘either the roulette ball will fall into a red pocket or it will fall into a black pocket’ is true” even when we do not hold any information about the truth or falsity of its immediate components. And, under the same circumstances, it is equally sensible to say “we hold the information that the sentence ‘the ball will fall into a red pocket and it will fall into a black pocket’ is false”. Now, the problem cannot be solved by stipulating that a disjunction takes the value 1, and a conjunction the value 0, whenever both their components take the value 1/2. When we hold no information about ψ and φ , whether or not we hold the information that a disjunction $\psi \vee \varphi$ is true, or the information that a conjunction $\psi \wedge \varphi$ is false, clearly depends on the content of ψ and φ and not only on their truth-values. This means that our sought informational semantics cannot be truth-functional.

5.2. Informational Meaning and Surface Semantic Information

It can be verified that the only meaning-conditions that can be justified in accordance with the founding principle of informational semantics (p. 50 above) are those expressed by the inference rules shown in Tables 2 and 3, where $T_a\varphi$ and $F_a\varphi$ are abbreviations for “the agent a (actually) holds the information that φ is true (respectively false)” and the conditional operator $\rightarrow$ is defined, as customary, in terms of the other operators (e.g., $\varphi \rightarrow \psi =_{def} \neg\varphi \vee \psi$).

Table 2. Sufficient conditions (introduction rules) for the standard Boolean operators.

$\frac{F_a\varphi}{T_a\neg\varphi} T\neg\mathcal{I}$	$\frac{T_a\varphi}{F_a\neg\varphi} F\neg\mathcal{I}$	
$\frac{T_a\varphi}{T_a\varphi \vee \psi} T\vee\mathcal{I}1$	$\frac{T_a\psi}{T_a\varphi \vee \psi} T\vee\mathcal{I}2$	$\frac{F_a\varphi}{F_a\varphi \vee \psi} F\vee\mathcal{I}$
$\frac{F_a\varphi}{F_a\varphi \wedge \psi} F\wedge\mathcal{I}1$	$\frac{F_a\psi}{F_a\varphi \wedge \psi} F\wedge\mathcal{I}2$	$\frac{T_a\varphi}{T_a\varphi \wedge \psi} T\wedge\mathcal{I}$
$\frac{F_a\varphi}{T_a\varphi \rightarrow \psi} T\rightarrow\mathcal{I}1$	$\frac{T_a\psi}{T_a\varphi \rightarrow \psi} T\rightarrow\mathcal{I}2$	$\frac{F_a\psi}{F_a\varphi \rightarrow \psi} F\rightarrow\mathcal{I}$

Table 3. Necessary conditions (elimination rules) for the four standard Boolean operators.

$\frac{T_a \neg \varphi}{F_a \varphi} T \neg \mathcal{E}$	$\frac{F_a \neg \varphi}{T_a \varphi} F \neg \mathcal{E}$		
$\frac{T_a \varphi \vee \psi}{F_a \varphi} T \vee \mathcal{E}1$	$\frac{T_a \varphi \vee \psi}{F_a \psi} T \vee \mathcal{E}2$	$\frac{F_a \varphi \vee \psi}{F_a \varphi} F \vee \mathcal{E}1$	$\frac{F_a \varphi \vee \psi}{F_a \psi} F \vee \mathcal{E}2$
$\frac{F_a \varphi \wedge \psi}{T_a \varphi} F \wedge \mathcal{E}1$	$\frac{F_a \varphi \wedge \psi}{T_a \psi} F \wedge \mathcal{E}2$	$\frac{T_a \varphi \wedge \psi}{T_a \varphi} T \wedge \mathcal{E}1$	$\frac{T_a \varphi \wedge \psi}{T_a \psi} T \wedge \mathcal{E}2$
$\frac{T_a \varphi \rightarrow \psi}{T_a \varphi} T \rightarrow \mathcal{E}1$	$\frac{T_a \varphi \rightarrow \psi}{F_a \varphi} T \rightarrow \mathcal{E}2$	$\frac{F_a \varphi \rightarrow \psi}{T_a \varphi} F \rightarrow \mathcal{E}1$	$\frac{F_a \varphi \rightarrow \psi}{F_a \psi} F \rightarrow \mathcal{E}2$

As required by the informational semantics, these rules specify the necessary (elimination rules) and sufficient (introduction rules) conditions for an agent a to hold the information that a compound sentence is true/false in terms of the information that a (actually) holds concerning its components. All the inferences that can be justified by means of these rules are to be regarded as analytic in the sense that, according to Quine's suggestions (see Section 2.2), their correctness is learned in the very process of learning the logical words.

Since the introduction and elimination (intelim) rules are intended as valid for any agent, reference to the agent can be omitted, unless otherwise required, by stripping the label a off the signs T and F .

Let us call *surface information state* any set X of signed sentences of the form $T \varphi$ or $F \varphi$ satisfying the following conditions:

1. for no sentence φ , $T \varphi$ and $F \varphi$ are both in X ;
2. if $S \psi$ (where S is either T or F) follows from the signed sentences in X by a chain of applications of the basic intelim rules, then $S \psi$ also belongs to X .

The first condition requires that no agent may actually hold information that is *explicitly inconsistent* and can be seen as an informational version of the more metaphysically flavoured principle of non-contradiction (no sentence can be at the same time true and false). Although this requirement is far from being uncontroversial, it seems to be in line with Kant's comment in his *Critique of Pure Reason*, where he regarded the Principle of Non-Contradiction as "the supreme principle of all analytical judgements":

[This principle] is a universal but purely negative criterion of all truth. But it belongs to logic alone, because it is valid of all cognitions, merely as cognitions and without respect to their content, and declares that the contradiction entirely nullifies them [37].

Let us say that a set Γ of sentences is *analytically inconsistent* if there is a sequence of applications of the intelim rules that leads from $\{T\psi \mid \psi \in \Gamma\}$ to both $T\varphi$ and $F\varphi$ for some sentence φ . It follows from the definition of information state that if Γ is analytically inconsistent, $\{T\psi \mid \psi \in \Gamma\}$ cannot be included in any information state.

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

Now, the *surface semantic information* carried by a sentence φ , $\text{INF}(\varphi)$ can be defined as

$$\text{INF}(\varphi) = \bigcap \{Y \in \mathcal{S} \mid T\varphi \in Y\} \quad (3)$$

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

$$\text{INF}(\Gamma) = \bigcap \{Y \in \mathcal{S} \mid T\Gamma \subseteq Y\} \quad (4)$$

where $T\Gamma = \{T\varphi \mid \varphi \in \Gamma\}$.

Observe that, since $\bigcap \emptyset = \top$, (4) yields $\text{INF}(\Gamma) = \top$ whenever Γ is analytically inconsistent, for there is no $Y \in \mathcal{S}$ that may include $T\Gamma$. 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 empty set of signed sentences. 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 analytically consistent.

This requirement of *analytic* consistency (not classical consistency) can be seen as a substantial mitigation of Floridi’s “veridicality thesis”. Even if one is not willing to endorse the somewhat controversial view that “information encapsulates truth” ([6] and [8] (Chapters 4 and 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 analytically inconsistent, no agent a can “hold the information” that all the sentences in Γ are true, because adding $T\Gamma$ 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:

$$\iota_a(\Gamma) = \text{INF}(\text{INF}(X_a \cup \Gamma) \sim X_a) \quad (5)$$

where X_a is the current information state of a . Again, it follows from (5) that $\iota_a(\Gamma) = \top$ whenever Γ is analytically inconsistent.

Finally, we define the following deducibility relation:

$\Gamma \vdash \varphi$ if and only if $T \varphi$ results from $T \Gamma$ by a chain of applications of the intelim rules.

Observe that $\Gamma \vdash \varphi$ if and only if $T \varphi$ belongs to every information state that includes $T \Gamma$ and so:

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

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 does not qualify as genuine information ($\top$ is not an information state).

Do this meaning-theory for the logical operators and its associated notion of surface semantic information satisfy SMR? The answer is yes. The details can be found in [1], [38] and [40], where it is shown that (i) the deducibility relation $\vdash$ is a logic in Tarski's sense, (ii) it satisfies the subformula property, and (iii) whether or not $\Gamma \vdash \varphi$ can be decided in polynomial (quadratic) time.

It may be objected that the deducibility relation $\vdash$ is still “explosive” when Γ is analytically inconsistent, for there is no information state for a that contains $T_a \varphi$ for all $\varphi \in \Gamma$. So, if Γ is analytically inconsistent, $\Gamma \vdash \varphi$ for every sentence φ . Accordingly, if Γ is analytically inconsistent, $\text{INF}(\Gamma) = \top$ and so, for every φ , $\text{INF}(\varphi) \subseteq \text{INF}(\Gamma)$. However, the problem raised by this kind of explosivity is far less serious than the similar problem for the classical deducibility relation. For, we can detect that the premises are analytically inconsistent in feasible time and, therefore, we may as well abstain from drawing bizarre conclusions on their basis. As Michael Dummett once put it:

Obviously, once a contradiction has been discovered, no one is going to go through it: to exploit it to show that the train leaves at 11:52 or that the next Pope will be a woman [12] (p. 209).

Unlike hidden classical inconsistencies, which may be hard to discover even for agents equipped with powerful (but still bounded) computational resources, analytic inconsistency rests, as it were, on the surface and can be feasibly detected. So, we always have a feasible means to establish that our premises are analytically consistent and for consistent premises the consequence relation $\vdash$ is *not* explosive, even if these premises are classically inconsistent.

We stress again that our definition of information state and surface semantic information do not require that information “encapsulates truth”, nor do they even require that it “encapsulates consistency”, but only that information “encapsulates *analytic* consistency”. Unlike classical consistency, this kind of surface consistency is *feasible* and so, in accordance with SMR, any agent is in the position to tell, in practice and not only in principle, whether she is holding genuine surface information or not.

According to this characterization, $\vdash$ is informationally trivial by definition, and this is in accordance with the tenet that analytic inferences are utterly uninformative. However, now both the meaning theory

and the residual notion of semantic information do satisfy SMR and so the scandal of deduction is dissolved. The inferences that can be justified by $\vdash$ are only a subclass of the classically valid inferences and their validity can be recognized in feasible time. Moreover, as we have already remarked, the Tarskian logic $\vdash$ contains no *tautologies*, it only allows for inferring from a non-empty set of signed formulas describing the agent's initial information. These are the “easy” inferences that (nearly) everybody learns to make correctly in the very process of learning the meaning of the logical operators. (This claim has been in fact severely tested in over 15 years of teaching to first-year students.) What about the other, more complex, inferences that are still classically valid? How do we fill the gap? Can it be filled in a *graded* way, so as to make a significant step towards a useful logical theory that indefinitely approximates the inferential power of an ideal Laplacian agent?

5.3. Actual vs. Virtual Information

In [1] 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:

$$\frac{\begin{array}{l} T_a \varphi \vee \psi \\ T_a \varphi \rightarrow \theta \\ T_a \psi \rightarrow \theta \end{array}}{T_a \theta} \quad (6)$$

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 (6) based on these analytic rules will necessarily have to introduce temporary assumptions that are “discharged” when the conclusion is drawn, as in the following schematic argument:

1	$T_a \varphi \vee \psi$		
2	$T_a \varphi \rightarrow \theta$		
3	$T_a \psi \rightarrow \theta$		
4	$T_a \varphi$		$F_a \varphi$
5	$T_a \theta$	by $T \rightarrow \mathcal{E}1$	$T_a \psi$ by $T \vee \mathcal{E}1$
			$T_a \theta$ by $T \rightarrow \mathcal{E}1$
7	$T_a \theta$		

Here, the information expressed by the signed sentences $T_a \varphi$ and $F_a \varphi$ is *not*, in general, information that is actually held by the agent a . It is *virtual information* that goes *beyond what is “given”* in the premises. This use of virtual information, which 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.

[...]

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 [...] [37].

One could say, by analogy, that analytical inferences are those that are recognized as sound via steps that are all “explicative”, that is, descending immediately from the meaning of the logical operators, as given by the necessary and sufficient conditions expressed by the elimination and introduction rules, while synthetic ones are those that are “augmentative”, involving some intuition that goes beyond this meaning, *i.e.*, 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 (and subsequence discharge) of virtual information. The manipulation of virtual information can be governed, like in the example given above, by a single *structural* proof rule that, for every formula φ that is a subformula either of the premises or of the conclusion, allows to split the argument in two distinct subarguments, depending on new and complementary virtual information, respectively $T_a \varphi$ and $F_a \varphi$. Such a tree of arguments involving virtual information is a proof of φ from Γ whenever each subargument either ends with an explicit inconsistency or with the conclusion $T_a \varphi$. This is only the intuitive idea, the formal details are in [1,38–40]. For each $n \in \mathbb{N}$, $\vdash_k$ is the consequence relation that allows for bounded applications of this proof rule up to a given fixed depth k . It can then be shown that classical logic is the limit of the sequence of these bounded consequence relations $\vdash_k$ as k tends to infinity. 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 following argument from the empty set of premises:

1			
2	$T_a p$		$F_a p$
3	$T_a p \vee \neg p$	by $T \vee \mathcal{I}1$	$T_a \neg p$ by $T \neg \mathcal{I}$
			$T_a p \vee \neg p$ by $T \vee \mathcal{I}2$
5	$T_a p \vee \neg p$		

6. Conclusions

We have discussed some of Floridi's views on semantic information, paying special attention to the solutions they offer to the anomalies of the received view, such as the Bar-Hillel–Carnap Paradox, the Scandal of Deduction and the Problem of Information Overload (aka. “Logical Omniscience”). Most of the discussion has been oriented by the approach to the scandal of deduction put forward in a joint paper by Floridi and the present author, so it is inevitably biased. However, as Popper repeatedly stressed, there is no problem with holding a biased view provided that one is willing to honestly and severely test it. Being our theory a non-empirical one, it can be tested mainly for consistency, intuitive plausibility and, above all, for its heuristic value, its capability of raising new interesting problems. Here, the proof of the pudding consists in trying to develop a quantitative Theory of Bounded Semantic Information to complement the qualitative theory that has been outlined in [1] and, in a rather informal and sloppy mood, in the present paper.

Acknowledgements

I wish to thank Stefania Bandini and Marcelo Finger for valuable comments on a previous draft.

References

1. D'Agostino, M.; Floridi, L. The enduring scandal of deduction. Is propositionally logic really uninformative? *Synthese* **2009**, *167*, 271–315.
2. Floridi, L. Semantic Conceptions of Information. Stanford Encyclopedia of Philosophy, 2011. Available online: <http://plato.stanford.edu/entries/information-semantic/> (accessed on 26 December 2012).
3. Shannon, C. Abstract to “The Lattice Theory of Information”. In *Claude Elwood Shannon: Collected Papers*; Sloane, N., Wyner, A., Society, I.I.T., Eds.; IEEE Press: New York, NY, USA, 1993; p. 180.
4. Floridi, L. *The Philosophy of Information*; Oxford University Press: Oxford, UK, 2011.
5. Hintikka, J. *Logic, Language Games and Information. Kantian Themes in the Philosophy of Logic*; Clarendon Press: Oxford, UK, 1973.
6. Floridi, L. Outline of a Theory of Strongly Semantic Information. *Mind. Mach.* **2004**, *14*, 197–222.
7. Floridi, L. Is Information Meaningful Data? *Philos. Phenomen. Res.* **2005**, *70*, 351–370.
8. Floridi, L. *The Philosophy of Information*; Oxford University Press: Oxford, UK, 2011.
9. Tennant, N. *Natural Logic*; Edinburgh University Press: Edinburgh, UK, 1990.

10. Hahn, H.; Neurath, O.; Carnap, R. The scientific conception of the world [1929]. In *Empiricism and Sociology*; Neurath, M., Cohen, R., Eds.; Reidel: Dordrecht, The Netherlands, 1973.
11. Hempel, C. Geometry and Empirical Science. *Am. Math. Mon.* **1945**, *52*, 7–17.
12. Dummett, M. *The Logical Basis of Metaphysics*; Duckworth: London, UK, 1991.
13. Quine, W. Two dogmas of empiricism. In *From a Logical Point of View*; Harvard University Press: Cambridge, MA, USA, 1961; pp. 20–46.
14. Davidson, D. On the very idea of a conceptual scheme. *Proc. Address. Am. Philos. Assoc.* **1974**, *47*, 5–20.
15. Quine, W. Two dogmas in retrospect. *Can. J. Philos.* **1991**, *21*, 265–274.
16. Quine, W. *The Roots of Reference*; Open Court: LaSalle, IL, USA, 1973.
17. Bergström, L.; Føllesdag, D. Interview with Willard van Orman Quine, November 1993. *Theoria* **1994**, *60*, 193–206.
18. Decock, L. True by virtue of meaning. Carnap and Quine on the analytic-synthetic distinction. 2006. Available online: http://vu-nl.academia.edu/LievenDecock/Papers/866728/Carnap_and_Quine_on_some_analytic-synthetic_distinctions (accessed on 27 December 2012).
19. Shannon, C.; Weaver, W. *The Mathematical Theory of Communication*; University of Illinois Press: Urbana, IL, USA, 1949.
20. Bar-Hillel, Y.; Carnap, R. Semantic Information. *Br. J. Philos. Sci.* **1953**, *4*, 147–157.
21. Popper, K. *The Logic of Scientific Discovery*; Hutchinson: London, UK, 1959.
22. Carnap, R.; Bar-Hillel, Y. An Outline of a Theory of Semantic Information. In *Language and Information*; Bar-Hillel, Y., Ed.; Addison-Wesley: London, UK, 1953; pp. 221–274.
23. Cohen, M.; Nagel, E. *An Introduction to Logic and Scientific Method*; Routledge and Kegan Paul: London, UK, 1934.
24. Carapezza, M.; D’Agostino, M. Logic and the myth of the perfect language. *Log. Philos. Sci.* **2010**, *8*, 1–29.
25. Sequoia-Grayson, S. The Scandal of Deduction. Hintikka on the Information Yield of Deductive Inferences. *J. Philos. Log.* **2008**, *37*, 67–94.
26. Cook, S.A. The complexity of theorem-proving procedures. *STOC '71: Proceedings of the Third Annual ACM Symposium on Theory of Computing*; ACM Press: New York, NY, USA, 1971; pp. 151–158.
27. Finger, M.; Reis, P. On the predictability of classical propositional logic. *Information* **2013**, *4*, 60–74.
28. Primiero, G. *Information and Knowledge. A Constructive Type-Theoretical Approach*; Springer: Berlin, Germany, 2008.
29. Sillari, G. Quantified Logic of Awareness and Impossible Possible Worlds. *Rev. Symb. Log.* **2008**, *1*, 1–16.
30. Duží, M. The Paradox of Inference and the Non-Triviality of Analytic Information. *J. Philos. Log.* **2010**, *39*, 473–510.
31. Jago, M. The content of deduction. *J. Philos. Log.* **2013**, in press.

32. Halpern, J.Y. Reasoning About Knowledge: A Survey. In *Handbook of Logic in Artificial Intelligence and Logic Programming*; Gabbay, D.M., Hogger, C.J., Robinson, J., Eds.; Clarendon Press: Oxford, UK, 1995; Volume 4, pp. 1–34.
33. Meyer, J.J.C. Modal Epistemic and Doxastic Logic. In *Handbook of Philosophical Logic*, 2nd ed.; Gabbay, D.M., Guenther, F., Eds.; Kluwer Academic Publishers: Dordrecht, The Netherlands, 2003; Volume 10, pp. 1–38.
34. Floridi, L. The Logic of Being Informed. *Logique et Analyse* **2006**, *49*, 433–460.
35. Primiero, G. An epistemic logic for being informed. *Synthese* **2009**, *167*, 363–389.
36. Gabbay, D.M.; Woods, J. The New Logic. *Log. J. IGPL* **2001**, *9*, 141–174.
37. Kant, I. *Critique of Pure Reason* [1781], Book II, Chapter II, Section I. Available online: <http://ebooks.adelaide.edu.au/k/kant/immanuel/k16p/index.html> (accessed on 27 December 2012).
38. D’Agostino, M. Analytic inference and the informational meaning of the logical operators. *Logique et Analyse* **2013**, in press.
39. D’Agostino, M. Classical Natural Deduction. In *We Will Show Them!*; Artëmov, S.N., Barringer, H., d’Avila Garcez, A.S., Lamb, L.C., Woods, J., Eds.; College Publications: London, UK, 2005; pp. 429–468.
40. D’Agostino, M.; Finger, M.; Gabbay, D.M. Semantics and proof theory of depth-bounded Boolean logics. *Theor. Comput. Sci.* **2013**, in press.

© 2013 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).