Tractable Depth-Bounded Logics and the Problem of Logical Omniscience

Marcello D'Agostino*
Department of Human Sciences
University of Ferrara
dgm@unife.it

Published in H. Hosni and F. Montagna (eds), *Probability, Uncertainty and Rationality*, pp. 245–275. Springer, 2010

1 Introduction

Theories of rationality, in their broad sense, are notoriously marred by highly unrealistic assumptions about the reasoning power of agents. This applies also to several logical theories that are supposed to play a significant role in the development of models of rationality. For example, any such model is bound to make a pervasive use of notions such as "information", "knowledge" and "belief" which prompt for conceptual clarification. Any attempt at such a clarification must involve an analysis of the logical rules that govern their use, and these rules must interact in a meaningful way with those that govern the relation of logical consequence in a usual propositional or first-order language. However, if the relation of logical consequence is taken to be that of classical logic, the interaction turns out to be highly problematic. 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),[1] if an agent i knows, or believes, or is informed that a sentence A is true, and B is a logical consequence of A, then i 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 $\Box_i$ express any of the propositional attitudes at issue, referred to the agent i. Then, the "logical omniscience" assumption can be expressed by saying that, for any finite set Γ

---

[*] I wish to thank the organizers and all the participants of the "One-day workshop on the logical foundations of rational interaction" hosted in Pisa by the Centro di Ricerca Matematica Ennio de Giorgi, on 4 November 2009, and in particular Hykel Hosni who pointed me at the potential connections between my work and the problem of logical omniscience.

[1] For a survey on epistemic and doxastic logic see [20, 24]. For information logic, or "the logic of being informed", see [17, 29].

of sentences,

$$\text{if } \Box_i A \text{ for all } A \in \Gamma \text{ and } \Gamma \vdash B, \text{ then } \Box_i B, \tag{1}$$

where $\vdash$ stands for the relation of logical consequence. Observe that, letting $\Gamma = \emptyset$, it immediately follows from (1) that any rational agent i 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 (1) emerges from the combined effect of the "distribution axiom", namely,

$$\Box_i(A \to B) \to (\Box_i A \to \Box_i B) \tag{\textbf{K}}$$

and the "necessitation rule":

$$\text{if } \vdash A, \text{ then } \vdash \Box_i A. \tag{\textbf{N}}$$

On the other hand, the paradoxical flavour of (1) 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_i is the "accessibility" relation for the agent i. Intuitively, if s is the actual world and sR_it, then i would regard t as "possible". Then, a forcing relation $\models$ is introduced to define truth for the complex sentences of the language, starting from the initial assignment to the atomic sentences. The forcing relation incorporates the usual semantics of classical propositional logic and defines the truth of $\Box_i A$ as "A is true in all the worlds that i 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, (1) appears to be both compelling and, at the same time, counter-intuitive.

Now, under this reading of the consequence relation $\vdash$ as based on classical logic, (1) may perhaps be satisfied by an "idealized reasoner", in some sense to be made more precise,[2] but it is out of the question that it is not satisfied, and is not likely to ever be satisfiable, in practice. 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 [3], that is, among the hardest problems in co-NP. Although not a proved theorem, it is a widely accepted conjecture that there exists no decision procedure for such problems that runs in polynomial (i.e., feasible) time. 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 (1) and the classical notion of logical consequence, which arises in any

---

[2] 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.

real application context, may only be solved either by waiving the assumption stated in (1), 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 $\Box_i$ is closed under logical consequence for any practical reasoner.

Both options have been discussed in the literature.[3] Observe that, according to the latter, the problem of logical omniscience does not lie in assumption (1) 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 ([18]) and Levesque ([23]) 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" ([23], p. 355). A more recent (and related) proposal can be found in [13], 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 ([13, 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.[4]

---

[3] See [24] (Section 4), [20] (Section 4) and [22] for a survey and proper references. See also: [26] for an interesting third view that draws on the tradition of subjective probability, and [1] 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 [32, 31] for an insightful discussion of this framework and proper references to the literature.

[4] The idea of approximating full classical propositional logic via a converging sequence of tractable subsystems has received considerable attention in the field of automated deduction (see, for instance, [30, 8, 9, 4, 14, 15]). See also [2] for an early paper that proposes to use approximated Boolean reasoning to deal with the problem of logical omniscience. In this work we are interested in tractable approximations that can themselves be described as *consequence relations* in some of the currently accepted senses (mainly, in Tarski's sense), equipped with a uniform, and possibly intuitive, account of the meaning of the logical operators. A brief discussion of the main differences between our approach and the cited ones can be found in [7] and [6]. Here we observe only that some of the systems proposed in the literature are not consequence relations, in that they do not satisfy unrestricted transitivity, while others (like the notion of "limited inference" of [23] and the NPL system of [13]) do not satisfy

In this paper we set out to make a contribution towards the solution of the logical omniscience problem that is much in the spirit of [13]. We too maintain that the problem can be properly solved by restricting the classical notion of logical consequence rather than by waiving assumption (1). We suggest, however, that an interesting alternative solution could be based on *Depth-Bounded Boolean Logics*, a novel incremental approach to the characterization of classical propositional logic that construes it as the limit of an infinite sequence of weaker *tractable* logics.[5] Agents committed to these logics can be seen as *approximations* to the idealized reasoning agent of standard epistemic, doxastic and information logic. The full decision problem for each of the approximating logics is solvable in polynomial time — although its complexity grows as we proceed along the approximation sequence — with *no restriction* to any particular syntactic fragment. Moreover, the meaning of the logical operators is *the same for all logics* and is explained in purely informational terms — that is, in terms of informational interpretations of "true" and "false" — in such a way that the most basic inference principles of classical propositional logic, including disjunctive syllogism, are preserved throughout the sequence.

Although this paper does not contain any new technical result on Depth-Bounded Boolean Logics (being parasitical on [7] and [6] in this respect), it sets out to shed new light on their potential applications, trying them on an open problem that is widely debated in the logical literature. When technical results are mentioned, the proofs are omitted and the reader is referred to the source papers for more details. In Section 2 we make a diagnosis of the problem of logical omniscience and list some desiderata that a characterization of logical consequence should meet to be a liable candidate for replacing classical logic in the attempt to solve the problem. Then, in Section 3, we outline an "informational interpretation" of the logical operators that is more in tune than their classical interpretation with the propositional attitudes investigated in epistemic, doxastic and information logic. In Section 4 we present the sequence of Depth-Bounded Boolean Logics from the semantical viewpoint, and in Section 5 we illustrate a variant of the proof-theoretic presentation introduced in [7] in the style of "natural deduction". Finally, in the concluding section, we outline future developments of the ideas presented in this paper.

2 A diagnosis and a wish list

What goes wrong in the interplay between classical logic and the propositional attitudes investigated in epistemic, doxastic or information logic? A possible clue points at the well-known philosophical fact that the inner working of clas-

---

disjunctive syllogism, namely, one of the most basic inference principles of classical logic that is immediately recognized as sound, *pace* relevance logicians, by any far from idealized agent, including Chrysippus' dog (see [16]).

[5] This approach was introduced in [7] and is still under development. The denomination "Depth-Bounded Logics" itself was not used in the original paper and has been put forward in [6], which contains a more detailed semantical account and a network-based proof-theoretic characterization that adjoins the "natural deduction" system presented in [7].

sical logic cannot be fully explicated by referring only to an agent's information state. More precisely, the classical meaning of the logical operators makes an essential reference to recognition-transcendent notions of truth and falsity as properties of sentences that apply to them quite independently of the information available to us. This is the concession that classical logic makes to metaphysical realism and is expressed by the principle of bivalence: every sentence is determinately either true or false, no matter whether we are able or not to recognize its truth or falsity on the basis of the available evidence. However, it seems plausible that a consequence relation that is likely to turn assumption (1) into a compelling property should be explicable in purely informational terms. Rather than being truth-preserving with respect to *possible worlds*, requiring that the conclusion be true in all possible worlds in which all the premisses are true, it should be truth-preserving with respect to *information states*, requiring that the conclusion be true in all the information states in which all the premisses are true. Now, if a notion of information state is suitably defined — to the effect that an agent can be *realistically* assumed to be totally in control of his own information state — it will be a truism that the propositional attitude of "being informed" be preserved under logical consequence, once this latter notion has been taken in the informational sense that we have just outlined. Moreover, it will be quite plausible that an agent should also know (or believe) all the logical consequences, still in the informational sense, of the sentences that he knows (or believes). For, the problem with (1) is that the agent might not be *aware* that the truth of a sentence A logically follows from the truth of the sentences in Γ that he knows or believes as true: were he aware of this, the propositional attitudes would certainly apply also to A. But, if the consequence relation $\vdash$ and the notion of information state are defined in such a way that

D1 $\Gamma \vdash A$ means that A is true in all information states in which all the sentences in Γ are true,

D2 any agent can be assumed to be totally in control of his own information state, namely, to have feasible access to the information contained in it,

then assumption (1) becomes unassailable even when $\Box_i$ is interpreted in terms of knowledge or belief. Observe also that, taken together, D1 and D2 imply that the notion of logical consequence we are looking for should be feasible.

Hence, a good reason for distributing the modalities $\Box_i$ over logical consequence is that the latter — defined in accordance with D1, let us call this *the informational sense of logical consequence* — does not add anything *new* to an agent's information state, it does not provide any genuinely new information. This is possible only because a sentence's being a logical consequence of a set of assumptions means that its truth can be established, given the truth of the assumptions, by virtue only of the meaning of the logical operators occurring in the sentences, and so it depends solely on the conventions governing our use of language. For example, an agent may establish by external means that $A \vee B$ is true and A is false, which will force her information state to include the truth of B solely by virtue of the accepted meaning of $\vee$. Once such meaning has been properly explicated, one should literally be able to "see" that the truth of the conclusion is contained in the truth of the premisses. Accordingly, Wittgenstein dreamt of a logically perfect language in which "we can in fact recognize the formal properties of propositions by mere inspection of the propositions themselves" (Tractatus 6.122) and "every tautology itself shows that it is a tautology" (6.127(b)). However, if we take for granted the classical meaning of the logical operators — as defined by the usual semantic clauses, in terms of recognition-transcendent notions of truth and falsity satisfying the principle of bivalence — the explication process required to reveal the soundness of an arbitrary inference turns out to be be unfeasible, and the theory of NP-completeness strongly suggests that this situation cannot be improved upon: there cannot be any perfect language in Wittgenstein's sense.

At this point, one may conjecture that intuitionistic logic could be apt to provide a suitable solution to our problem. First, the intuitionistic 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.[6] Moreover, 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".[7] However, replacing classical with intuitionistic logic would be no solution, because the decision problem for the propositional fragment of the latter is PSPACE-complete (that is, among the hardest problems in PSPACE)[8] and, again, it is highly plausible that there is no polynomial time decision procedure for solving such problems. Prima facie, this may sound like a refutation of our initial diagnosis. Here we have a notion of logical consequence that is indeed construed in terms of information states; and yet, by adopting it, the logical omniscience problem would even be aggravated, since PSPACE-completeness is usually regarded as stronger evidence of intractability than co-NP-completeness. But a closer analysis reveals that the notion of truth at an "information state" that is embodied in Kripke's semantics has features that make it clash with our desideratum D2.

The first 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.

---

[6] 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 [12] for an overall picture. On the informational view of intuitionistic logic and its relations with epistemic logic see [37].

[7] See [21], p. 100.

[8] See [33].

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 $\{\rightarrow\}$-fragment of intuitionistic logic is also PSPACE complete ([33, 34]). The necessity 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:[9] we essentially need to go *beyond the data*, using "virtual information", i.e., simulating situations in which we hold information that, in fact, we do not hold. Although all virtual information is eventually removed, to the effect that the conclusion 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.

The 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.[10] 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.[11]

Beth's semantics for intuitionistic logic[12] seems to offer a more natural account of the truth of disjunctive statements based on an information state: a disjunction $A \vee B$ is true at the actual state provided we have the means of recognizing that necessarily one or the other disjunct will eventually become true at some future information state. However, at least in non-mathematical

---

[9] See [7] on this point.

[10] See [12] (Chapter 11) and [28] for a thorough discussion.

[11] 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.

[12] For an exposition, see [10].

reasoning, we may well be in a position to assert a disjunction even if we have no means of recognizing that we will eventually reach an information state that will enable us to establish the truth of one of the two disjuncts.[13] If we do have such means, this should be regarded as additional information, not as information that is incorporated in the original assertion of the disjunction. So, it seems inappropriate to assume, as part of the ordinary meaning of $\vee$, that the assertion of $A \vee B$ can be licensed only if we know in advance that some future information state will infallibly put us in a position to assert one of the two disjuncts. As Michael Dummett puts it:

> I may be entitled to assert "A or B" because I was reliably so informed by someone in a position to know, but if he did not choose to tell me which alternative held good, I could not apply an or-introduction rule to arrive at that conclusion. [...] Hardy may simply not have been able to hear whether Nelson said "Kismet hardy" or "Kiss me Hardy", though he heard him say one or the other: once we have the concept of disjunction, our perceptions themselves may assume an irremediably disjunctive form ([12], pp. 266–267). [...]
>
> Unlike mathematical information, empirical information decays at two stages: in the process of acquisition, and in the course of retention and transmission. An attendant directing theatre-goers to different entrances according to the colours of their tickets might even register that a ticket was yellow or green, without registering which it was, if holders of tickets of either colours were to use the same entrance; even our observations are incomplete, in the sense that we do not and cannot take in every detail of what is in our sensory fields. That information decays yet further in memory and in the process of being communicated is evident. In mathematics, any effective procedure remains eternally available to be executed; in the world of our experience, the opportunity for inspection and verification is fleeting ([12], pp. 277–278).

One way of going around this difficulty consists in postulating that, whenever the assertion of a disjunctive statement can be made at all, one or the other disjunct *could have* been asserted by a properly informed agent. Again, such a way out of the difficulty requires what we have called "virtual information", going *beyond* the information that is actually held by the agent that is making the assertion.[14]

The rôle played by virtual information is apparent in the so-called "discharge rules" of natural deduction, a proof-theoretic presentation of logical consequence

---

[13] In classical terms, the situation could be easily described as follows: according to our information we are able to assert that at least one of P and Q is true, but we are not able, and we may never be able, to assert that either of the two alternatives is true. However, even this description is inadequate for our purposes, since the sentence "at least one of P and Q is true" can be understood only with reference to the classical notion of truth as independent of our information state.

[14] On this point see [12], especially Chapter 12. and [11], especially Chapter 14.

that is very close to the intuitionistic explanations of the logical operators.[15] For example, in the $\vee$-elimination rule,

$$\frac{\Gamma \vdash A \vee B \quad \Gamma, A \vdash C \quad \Gamma, B \vdash C}{\Gamma \vdash C,}$$

each of the discharged assumptions A and B represents a piece of information that needs not be included in all information states that verify the sentences in Γ (when it is, the rule application is indeed redundant).

The above considerations strongly suggest that the notion of information state underlying both Kripke's and Beth's semantics, and the way the meaning of the logical operators is explicated thereby, is inadequate for our purposes, since it essentially requires a reference to "virtual information" that is not actually contained in the agent's information state with respect to which this meaning is being explicated. As a consequence, the contents of an information state are somehow overstated, in such a way that the desideratum D2 above cannot be guaranteed. So, in the light of the analysis carried out in this section, we can add to D1 and D2 a further desideratum for a putative consequence relation that is likely to solve the problem of logical omniscience. Given that the truth or falsity of a sentence in an information state is established either by external means or by virtue of the meaning of the logical operators, we may request that:

D3 the meaning of the logical operators should be fixed by appropriate conventions expressed exclusively in terms of an agent's information state, so that the agent in question can be assumed to be totally in control of this meaning; in particular, any reference to virtual information should be avoided.

Because of D3, we expect the logical operators that may emerge from this effort to bear a *weaker* meaning than the one they usually bear in classical semantics: we may call this *the informational meaning of the logical operators*. The underlying guess is that the inference steps that cannot be justified by virtue of such weaker informational meaning, i.e., without appealing to what we have called "virtual information", are exactly those that essentially increase the computational complexity of deductive inference. In other words, a meaning-theory satisfying D3 is likely to yield a notion of information state satisfying D2 and so, via D1, a feasible notion of logical consequence that may dissolve the paradox that lurks in the usual reading of (1).

This does not imply that one cannot define tractable logics in which the use of a certain amount of virtual information is tolerated. It implies, however, that if virtual information is essentially used *to define the meaning of the logical operators*, so that its unbounded use *has* to be tolerated if one wants to fully exploit this meaning in drawing logical conclusions, then the resulting logic is likely to be intractable. On the other hand, a weaker definition of this meaning,

---

[15] See [19, 27]; see also [36] for an excellent exposition

in accordance with D3, would pave the way for gradually re-introducing virtual information, by imposing an upper bound on its recursive use. Then, we may look for some purely *structural principle*, expressing such bounded manipulation of virtual information, which can allow us to define a sequence of tractable logics, depending on the chosen upper bound, that converge to classical propositional logic and can legitimately be said to share *the same logical operators.* In this way the idealized reasoning agent of epistemic, doxastic and information logic could be arbitrarily approximated by realistic agents of increasing deductive power, but *all speaking the same language.* These considerations lead to our final desideratum:

D4 classical propositional logic should be characterized as the infinite union $\bigcup_{k\in\mathbb{N}}\{\vdash_k\}$ of approximating logics $\vdash_k$, such that (i) $\vdash_k\subset\vdash_{k+1}$, for every $k \in \mathbb{N}$, (ii) the meaning of the logical operators should be the same for all the approximating logics $\vdash_k$, and (iii) each logic $\vdash_k$ should be defined by an upper bound on the recursive use of some purely structural principle expressing the manipulation of virtual information.

It is apparent that a key role, in this road map, is played by D3. So, we now turn to the main question it raises: is there such a thing as "the informational meaning of the logical operators" and how can it be properly defined?

3 The informational meaning of the logical operators

Whatever the nature of the information concerned may be, an information state should provide a *partial* valuation v of the sentences of a standard propositional language into $\{0, 1\}$, where "0" stands for "false" and "1" stands for "true", describing the effect of the information held by an agent on the assertion or rejection of sentences. Intuitively, $v(A) = 1$ means that the agent's information state licenses the assertion of A as true, while $v(A) = 0$ means that it licenses the rejection of A as false. The sentences for which v is undefined are those that, on the basis of the information available, can neither be asserted as true, nor rejected as false. We shall write "$v(A) = \bot$" for "A is undefined in v".[16] Within an information state, we can distinguish a set of basic sentences whose truth-value is established by external means, the truth-values of the other sentences being established, starting from the basic ones, by virtue of the very meaning of the logical operators. The latter is usually fixed by defining, within the set of all possible valuations, those which are *admissible.* In classical semantics this is done by specifying the following set of if-and-only-if conditions:

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

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

---

[16] Although $\bot$ is often used as a logical constant standing for the "absurd", here we use it in the sense of "undefined", as customary in domain theory.

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

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 is *upward saturated*, if it satisfies the "if-part" of the above conditions, and *downward saturated* if it satisfies the"only-if" parts. A *Boolean valuation* is a saturated valuation that satisfies the additional condition of being *total*, i.e. defined for all sentences.[17]

Now, not only cannot it be assumed that valuations representing information states be total, but the discussion in the previous section shows that some of the only-if conditions should also be dropped. In particular, the only-if parts of C3, C4 and C6 cannot be justified. That $A \vee B$ is assigned the value 1 by an agent's information state, does not guarantee that either A is assigned the value 1 or B is assigned the value 1, and so the only-if part of C3 fails. Dually, if $A \wedge B$ is assigned the value 0, this does not imply that either A is assigned the value 0 or B is assigned the value 0, and so the only-if part of C6 fails too.[18] Since we are dealing with classical (Boolean) conditional, the failure of the only-if part of C4 is a side-result of the definition of $A \to B$ as $\neg A \vee B$. Thus, if the meaning of $\vee$ and $\wedge$ (and, derivatively, of $\to$) has to be understood, in accordance with D3, by exclusive reference to an agent's *actual* information state (i.e., without allowing for any use of virtual information) there is no way in which one can express this meaning by means of if-and-only-if clauses like the ones usually employed in classical semantics. Moreover, the standard way in which this meaning is fixed in the so-called proof-theoretic semantics, by specifying suitable introduction and elimination rules — the former playing a prominent rôle in this task and the latter being, in some sense, derivative[19] — is also inadequate for our purposes because some of the rules that are taken as part of the meaning of $\vee$ and $\to$ would make essential use of virtual information, as the usual natural deduction rules for $\to$-introduction and $\vee$-elimination.[20]

---

[17] Observe that, for total valuations, conditions C5–C8 are redundant, in that they can be derived from conditions C1–C4. On the other hand, every valuation satisfying C1–C8 that is total over the *atomic* sentences of the language, is total over the whole language and, therefore, is a Boolean valuation.

[18] Observe that the intuitionistic meaning of $\vee$, via the notion of canonical verification, satisfies the only-if part of C3 (under the intuitionistic interpretation of truth and falsity) and so essentially agrees with its classical meaning. By way of contrast, the intuitionistic meaning of $\wedge$ does not satisfy the only-if part of C6.

[19] Proof-theoretic semantics dates back to [19]. See [28] for a discussion and proper references. See also [25] for interesting remarks on this topic.

[20] See also [7] on this point.

$\neg A$	A
1	1
0	0

$A \vee B$	A	B
1	0	0
0	1	1
0	1	0
0	1	$\bot$
0	0	1
0	$\bot$	1

$A \wedge B$	A	B
1	0	1
1	0	0
1	0	$\bot$
1	1	0
1	$\bot$	0
0	1	1

$A \to B$	A	B
1	1	0
0	1	1
0	0	1
0	$\bot$	1
0	0	0
0	0	$\bot$

Figure 1: Constraints on admissible partial valuations. Each line represents a *forbidden* local configuration of values.

The problem is then: how do we fix what we have called "the informational meaning of the logical operators"? What kind of conventions can be distilled from linguistic practice in order to determine this weaker kind of meaning, other than the standard if-and-only-if conditions on admissible valuations or the standard Gentzen-style intelim rules? If we maintain that to grasp the meaning of a logical operator consists in acquiring some kind of linguistic information, it seems plausible that the latter can be unfolded by determining which possibilities are *ruled out* if one wants to use that operator consistently. From this point of view, a solution to our problem may consist in taking the informational meaning of a logical operator to be fixed by a set of *negative constraints* on the valuations describing an agent's information state, specifying which ones should be ruled out as *inadmissible*. For example, a valuation that assigns the value 1 to $A \vee B$ and, at the same time, the value 0 to both A and B, should be excluded as inadmissible, revealing a mismatch between the valuation and the accepted (informational) meaning of $\vee$. An agent cannot, at the same time, hold the information that $A \vee B$ is true and the information that A and B are both false, without immediately realizing that this information is inconsistent. On the other hand, a valuation that assigns 1 to $A \vee B$, while being undefined for both A and B, is perfectly admissible and corresponds to a legitimate use of the word "or". Similarly, an agent cannot, at the same time, consistently hold the information that $A \wedge B$ is false and the information that A and B are both true, while a valuation that assigns 0 to $A \wedge B$, while being undefined for both A and B, is admissible and complies with the ordinary use of $\wedge$. A set of negative constraints that formally agree with the classical truth-tables, and can therefore be taken as distilling their informational content, is shown in Figure 1, where A and B stand for sentences of arbitrary complexity and each line, in the table for a given operator, represents a forbidden assignment, so that any valuation containing such an assignment is inadmissible.

Although these constraints may be said to reflect the classical meaning of the logical operators, to the extent that this meaning can be expressed in terms of an agent's actual information state, they do not immediately justify *any* logical inference. In the next section, however, we shall observe the gradual emergence of inference rules once these constraints are combined with *purely structural principles* that can be naturally associated with the "depth" of the inferential process involved.

4 The semantics of Depth-Bounded Boolean Logics

We start by observing how some basic inferences can be recognized as valid, on the grounds of the meaning-constraints, by virtue of a minimal structural principle that we call "the single candidate principle", or SCP for short, after the well-known strategy for solving simple sudoku puzzles:

> *Infer that A is true (false) if the other option is* immediately *ruled out by some of the accepted constraints that define the meaning of the logical operators.*

In what follows, we shall use the lower case letters p, q, r, etc. as variables for atomic sentences and continue using the capital letters A, B, C,etc. as variables for arbitrary sentences. As before, we shall use capital Greek letters Γ, Δ, Λ, etc. as variables for sets of sentences. Consider a valuation v such that $v(p \vee q) = 1$ and $v(p) = 0$, while $v(q) = \bot$. We can legitimately say that the value of q in this valuation is implicitly determined by the values of $p \vee q$ and p and by our understanding of the meaning of $\vee$ based on the constraints specified in Figure 1. For, there is *no admissible refinement* of v, that is, a refinement compatible with the meaning constraints, such that $v(q) = 0$: given the actual values of the other sentences, the assignment of 0 to q would be *immediately*[21] recognized as inadmissible by any agent that understands $\vee$ via the specified meaning constraints. In other words, the value 1 is *deterministically dictated* by v, since it is the only defined value that q can possibly take. The value 1 can therefore be assigned to q by exclusion of the other defined value.

By contrast, consider a typical example of "reasoning by cases": if our information state, described by a partial valuation v, is such that $v(p \vee q) = 1$, $v(p \to r) = 1$ and $v(q \to r) = 1$, then the piece of information that r is true is also, in some sense, implicitly contained in the information currently available, but we cannot specify this sense without introducing virtual information concerning p or q. We can reason as follows:

1. p must be either objectively true or objectively false (although this cannot be established on the basis of the current information);
2. assuming that we were informed about the objective truth-value of p
   (a) if we were informed that p is true, then the constraints on the meaning of $\to$ would rule out the possibility that r is false and so, by the SCP, r should be assigned the value 1;
   (b) if we were informed that p is false, then

---

[21] This means that the immediate refinement of v obtained by assigning 0 to q is inadmissible.

        i. the constraints on the meaning of $\vee$ would rule out the possibility that q is false and so, by SCP, q should be assigned the value 1;
        ii. if q were assigned the value 1, by the meaning constraints on $\rightarrow$ and SCP, r should be assigned the value 1.

3. Hence, r must be assigned the value 1 whatever the objective truth-value of p may be; therefore, the information that r is true is "implicitly contained" in our initial information state.

It is apparent that this sense of "implicitly contained" is essentially different from the sense in which the conclusion of disjunctive syllogism is implicitly contained in any information state that verifies the premisses, because it requires the introduction of virtual information in steps 2(a) and 2(b). These steps cannot be internally justified on the basis of the agent's actual information state, but involve simulating the possession of definite information about the objective truth-value of p, by enumerating the two possible outcomes of the process of acquiring such information, *neither of which* is deterministically dictated by v. The inference displays, intuitively, a deeper reasoning process than the one displayed by disjunctive syllogism, and we relate this depth to the necessity of manipulating virtual information concerning p. An even deeper inference process would be displayed if steps 2(a) and 2(b) themselves contained, in a recursive fashion, further use of virtual information in order to obtain the common conclusion. In the remains of this section, we shall elaborate on these intuitive remarks in a more systematic way, leading to a classification of inferences according to their logical depth, starting from the most basic ones that do not require any use of virtual information.

Let $\mathcal{L}$ be a standard language for propositional logic and let $\mathcal{A}$ be the set of all admissible partial valuations of $\mathcal{L}$, namely, all the partial functions $\mathcal{L} \to \{0,1\}$ that do not violate the negative constraints in Figure 1. The set $\mathcal{A}$ is 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 formulas for which v is defined. Being a partial function, a partial valuation v is a set of pairs of the form $\langle A, i\rangle$, where A is a sentence of the given language and i is equal to 0 or 1, subject to the restriction that, for no $A \in \mathcal{L}$, $\langle A, 1\rangle$ and $\langle A, 0\rangle$ are both in v. Each pair in v 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 that is compatible with the intended meaning of the logical operators. The partial ordering $\sqsubseteq$ is a meet-semilattice with a bottom element equal to $\emptyset$, the valuation which is undefined for all formulas of the language. It fails to be a lattice because the join of two admissible valuations may be inadmissible.

Let $\mathcal{L}^*$ be the *evaluated language* based on $\mathcal{L}$, i.e. the set of all ordered pairs $\langle A, i\rangle$ such that A is a formula of $\mathcal{L}$ and $i \in \{0,1\}$. Let $\Vdash_0$ be a relation $\mathcal{A} \times \mathcal{L}^*$ satisfying the following condition:

$$v \Vdash_0 \langle A, i\rangle \text{ if and only if } v \cup \{\langle A, |i-1|\rangle\} \notin \mathcal{A}. \tag{SCP}$$

Clearly, (SCP) expresses the structural property that we have called "single candidate principle". Notice that, by definition, $v \Vdash_0 \langle A, i\rangle$ for all $\langle A, i\rangle \in v$.

The image of $\Vdash_0$ under a partial valuation v represents all the information that can be recognized as "implicitly contained" in v without any need to introduce virtual information or, as we also say, *at depth* 0. This information may be seen to stem immediately from the (informational) meaning of the logical operators via a basic structural principle such as SCP. Now, we define a 0-depth information state as an admissible partial valuation that is closed under $\Vdash_0$:

Definition 4.1 *A partial valuation v is a* 0-depth information state *if and only if $v \in \mathcal{A}$ and, for all $A \in \mathcal{L}$ and $i \in \{0, 1\}$,*

$$\textit{if } v \Vdash_0 \langle A, i\rangle, \textit{ then } \langle A, i\rangle \in v.$$

Remark 4.2 *It may be the case that v is an admissible valuation, but v cannot be embedded into any 0-depth information state. For example, suppose that v is such that $v(A \vee B) = 1$, $v(A) = 0$, $v(B \wedge C) = 0$, $v(C) = 1$ and $v(B) = \perp$; then v is admissible (it does not violate any of the meaning constraints), but $v \Vdash_0 \langle B, 1\rangle$ and $v \cup \langle B, 1\rangle$ is not in $\mathcal{A}$, because it violates one of the meaning constraints for $\wedge$. So, there is no 0-depth information state v' such that $v \sqsubseteq v'$.*

We can now define the consequence relation $\vdash_0$ as truth-preserving over 0-depth information states.

Definition 4.3 *For every finite set Γ of formulas and every formula A, $\Gamma \vdash_0 A$ if and only if $v(A) = 1$ for all 0-depth information states v such that $v(B) = 1$ for all $B \in \Gamma$.*

The reader can check that $\vdash_0$ is a consequence relation in Tarski's sense, that is, it satisfies the following conditions for all formulas A, B and all finite sets Γ of formulas:

$$A \vdash_0 A \qquad \text{(Reflexivity)}$$

$$\text{if } \Gamma \vdash_0 A, \text{ then } \Gamma, B \vdash_0 A \qquad \text{(Monotonicity)}$$

$$\text{if } \Gamma \vdash_0 A, \text{ and } \Gamma, A \vdash_0 B, \text{ then } \Gamma \vdash_0 B. \qquad \text{(Transitivity)}$$

Moreover, $\vdash_0$ is also *substitution-invariant*, that is, it satisfies:

$$\text{if } \Gamma \vdash_0 A, \text{then } \sigma(\Gamma) \vdash_0 \sigma(A), \qquad \text{(SubInv)}$$

for every uniform substitution σ.

The logic $\vdash_0$ is the basic element in our hierarchy of depth-bounded logics approximating full Boolean logic, namely, the one that allows for no use of virtual information and therefore contains all the logical inferences that can be validated by virtue only of the informational meaning of the logical operators, as fixed by the negative meaning constraints, and of the *purely structural* principle SCP. We call it *the Boolean Logic of depth* 0. As a result of Theorems 5.1 and

5.3 below, *this logic is tractable*. We may regard it as a minimum requirement on a reasoning agent that she is able to recognize 0-depth logical consequences.[22]

Say that a set Γ of formulas is 0-*depth inconsistent* if there is no 0-depth information state that verifies all the formulas in Γ. (Since $\vdash_0$ is a proper subsystem of classical logic, a set of formulas may be classically inconsistent and 0-depth consistent at the same time.) Observe that the logic $\vdash_0$ validates, by definition, the controversial *ex-falso quodlibet principle*: if Γ is a 0-depth inconsistent set of formulas, then $\Gamma \vdash_0 A$ for every formula A. In this logic, however, the principle in question is not nearly as dangerous as it is in full classical logic because the 0-depth inconsistency of Γ can be feasibly detected (see Theorem 5.3 below). On the other hand, if Γ is classically inconsistent, but 0-depth consistent, the ex-falso principle does not apply to Γ, because there is some 0-depth information state that verifies all the formulas in Γ.

It is also worth noticing that $\vdash_0$, like Belnap's four-valued logic and the NPL system of [13], *has no tautologies*. This is not surprising, however, since a tautology is a sentence that is a "logical consequence of the empty set of assumptions" and so, in order to establish its truth in any information state, we must make essential use of virtual information, to the effect that the information state itself cannot be of depth 0. On the other hand, $\vdash_0$ validates a good deal of classical inference schemes, including *modus ponens* ($A \to B, A \vdash_0 B$), *modus tollens* ($A \to B, \neg B \vdash_0 \neg A$), *disjunctive syllogism* ($A \vee B, \neg A \vdash_0 B$) and its dual ($\neg(A \wedge B), A \vdash_0 \neg B$). Moreover, the transitivity of $\vdash_0$ and its being based on a systematic view of the meaning of the logical operators, are features that make it a respectable, although minimalist, logical system. This system achieves tractability in a natural way, which is expressed in its semantics, rather than by tampering procedurally with a system for classical logic.

Given two admissible valuations v and v', we say that *v' is a refinement of v on P* if (i) $v \sqsubseteq v'$ and (ii) P is defined in v'. Let now $\Vdash_k$, for $k = 1, 2, \ldots$, be a relation $\mathcal{A} \times \mathcal{L}^*$ satisfying the following condition:

$$\begin{array}{l} v \Vdash_k \langle A, i\rangle \text{ if and only if there exists an atomic } p \text{ such that} \\ v'(A) = i \text{ for every information state } v' \text{ of depth } k-1 \\ \text{that refines } v \text{ on } p. \end{array} \qquad (\mathrm{PB}(k))$$

An information state of depth k can be simultaneously defined as an admissible valuation closed under $\Vdash_k$.

Definition 4.4 *A partial valuation v is a k-*depth information state *if and only if $v \in \mathcal{A}$ and, for all $A \in \mathcal{L}$ and $i \in \{0, 1\}$,*

$$\text{if } v \Vdash_k \langle A, i\rangle, \text{ then } \langle A, i\rangle \in v.$$

The image of $\Vdash_k$ under a partial valuation v represents all the information that can be recognized as "implicitly contained" in v by means of the meaning constraints augmented with a *purely structural* principle, PB(k), that allows for

---

[22] One could say that the agent's "awareness" should be closed under 0-depth consequences.

bounded use of virtual information. The parameter k represents the maximal number of nested introductions of atomic virtual information that are allowed at each step of the process. From a classical viewpoint, PB(k) allows for expansions of the current information state by means of at most k nested applications of the classical "Principle of Bivalence".

Remark 4.5 *Given an admissible valuation v, it may be the case that, for some atomic p and some k, there is no information state of depth $k-1$ that refines v on p. Under these circumstances PB(k) implies that $v \Vdash_k \langle A, i\rangle$, whatever A and i may be. Then, if a refinement of v were closed under $\Vdash_k$, it could not be admissible and so v cannot be embedded in an information state of depth k. Moreover, if $v \cup \{\langle p, i\rangle\}$ cannot be embedded in an information state of depth $k-1$, while $v \cup \{p, |i-1|\rangle\}$ can, then (PB(k)) implies that $v \Vdash_k \langle A, i\rangle$ whenever $v \cup \{\langle p, |i-1|\rangle\} \Vdash_{k-1} \langle A, i\rangle$.*

The consequence relation $\vdash_k$ can then be defined as truth-preserving over information states of depth k.

Definition 4.6 *For every finite set Γ of formulas and every formula A, $\Gamma \vdash_k A$ if and only if $v(A) = 1$ for all k-depth information states v such that $v(B) = 1$ for all $B \in \Gamma$.*

For example, the argument given above for the inference expressing the principle of "reasoning by cases" shows the validity of the sequent $p \vee q, p \to r, q \to r \vdash_1 r$. We call $\vdash_k$ *the Boolean Logic of depth k* and it follows from Definitions 4.3 and 4.6 that $\vdash_k \subset \vdash_{k+1}$ for every $k \in \mathbb{N}$.

Each $\vdash_k$ is a consequence relation in Tarski's sense. Moreover, in each $\vdash_k$ the meaning of the logical operators is fixed by the same negative constraints, namely, those of Section 3. It is not difficult to show that:

Theorem 4.7 *Let $\vdash_C$ be the consequence relation of classical propositional logic. Then:*

$$\vdash_C = \bigcup_{k \in \mathbb{N}} \vdash_k .$$

A set Γ of sentence is *k-depth inconsistent* if there is no k-depth information state that verifies all the formulas in Γ. As for k-depth tautologies, unlike $\vdash_0$, every k-depth logic with $k > 0$ has its related set of tautologies. In particular, if A is a formula containing k atomic sentences, A is a classical tautology if and only if A is true in all information states of depth k, that is, if and only if $\vdash_k A$. (This is a straightforward consequence of the fact that every k-depth information state can simulate the complete truth table of any formula with k atomic sentences.)

Again, each Boolean Logic of depth k is tractable (this follows from Theorems 5.5 and 5.6 in the next section), although the complexity of the decision procedure essentially grows with k. Hence, each $\vdash_k$ can be regarded as a feasible approximation to the unrealistic deductive power of classical propositional logic.[23]

5 Natural deduction for Depth-Bounded Boolean Logics

Let $\mathcal{L}_s$ be *the signed language* based on $\mathcal{L}$, namely, the set of all expressions of the form $T\,A$ and $F\,A$, for $A \in \mathcal{L}$. The elements of $\mathcal{L}_s$ are called *signed formulas*. The intuitive interpretation of signed formulas is the usual one: $T\,A$ means "A is true" and $F\,A$ means "A is false". We shall use the lower case Greek letters, φ, ψ, χ, etc. as variables for arbitrary signed formulas and the capital letters X, Y, Z, etc. as variables for sets of signed formulas. We construe a deduction of the signed formula φ from the assumptions in X as a *sequence of signed formulas* starting with signed formulas in X and ending with φ, such that every intermediate element instantiates the conclusion of some schematic inference rule whose premisses are instantiated by previous elements of the sequence. In the context of $\vdash_0$ we do not need any device for discharging hypothesis, since any use of virtual information is banned. So, our inference rules will be of the simplest type, namely, principles licensing the assertion of a signed sentence of a certain form given the prior assertion of a finite number of other sentences of related forms.[24] Such simple inference rules will therefore be represented as follows:

$$\frac{\begin{array}{c} S_1\,A_1 \\ \vdots \\ S_{n-1}\,A_{n-1} \end{array}}{S_n\,A_n}$$

where each S_i (with $i = 1, \ldots, n$) is either "T" or "F".

Little reflection shows that, by the combined action of the meaning constraints of Section 3 and of the structural principle that we have named "Single Candidate Principle", the introduction and elimination rules of Figures 2 and 3 are all sound for $\vdash_0$. By the same means, the "mingle" elimination rules of Figure 4 are also shown to be sound for $\vdash_0$ and cannot be derived from the other rules. We shall use the expression "intelim rules" to refer to all these inference rules collectively.

An *intelim sequence* based on a set X of signed formulas is a sequence $\varphi_1, \ldots, \varphi_n$ of signed formulas such that each element φ_i of the sequence either (i) is an element of X, or (ii) results from preceding elements of the sequence by an application of one of the intelim rules in Figures 2, 3 and 4. An intelim sequence is *closed* if it contains both $T\,A$ and $F\,A$ for some formula A, otherwise

---

[23] The same remarks concerning the *ex-falso quodlibet* principle made above for $\vdash_0$ apply to $\vdash_k$ as well: if $\vdash_k$ is regarded as realistic, then an agent can detect the k-depth inconsistency of any finite set Γ of sentences and act upon it. On the other hand, if the inconsistency can be detected only at depth $m > k$, the *ex-falso quodlibet* principle does not apply to Γ at depth k.

[24] Clearly, the assertion of a signed formula $F\,A$ is tantamount to the rejection of A.

$$\frac{F\,A}{T\,A \to B}\;T \to \mathcal{I} \qquad \frac{T\,B}{T\,A \to B}\;T \to \mathcal{I}2 \qquad \frac{\begin{array}{l}T\,A\\ F\,B\end{array}}{F\,A \to B}\;F \to \mathcal{I}$$

$$\frac{T\,A}{T\,A \vee B}\;T \vee \mathcal{I}1 \qquad \frac{T\,B}{T\,A \vee B}\;T \vee \mathcal{I}2 \qquad \frac{\begin{array}{l}F\,A\\ F\,B\end{array}}{F\,A \vee B}\;F \vee \mathcal{I}$$

$$\frac{\begin{array}{l}T\,A\\ T\,B\end{array}}{T\,A \wedge B}\;T \wedge \mathcal{I} \qquad \frac{F\,A}{F\,A \wedge B}\;F \wedge \mathcal{I}1 \qquad \frac{F\,B}{F\,A \wedge B}\;F \wedge \mathcal{I}2$$

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

Figure 2: Introduction rules for signed sentences.

$$\frac{\begin{array}{l}T\,A \to B\\ T\,A\end{array}}{T\,B}\;T \to \mathcal{E}1 \qquad \frac{\begin{array}{l}T\,A \to B\\ F\,B\end{array}}{F\,A}\;T \to \mathcal{E}2 \qquad \frac{F\,A \to B}{T\,A}\;F \to \mathcal{E}1 \qquad \frac{F\,A \to B}{F\,B}\;F \to \mathcal{E}2$$

$$\frac{\begin{array}{l}T\,A \vee B\\ F\,A\end{array}}{T\,B}\;T \vee \mathcal{E}1 \qquad \frac{\begin{array}{l}T\,A \vee B\\ F\,B\end{array}}{T\,A}\;T \vee \mathcal{E}2 \qquad \frac{F\,A \vee B}{F\,A}\;F \vee \mathcal{E}1 \qquad \frac{F\,A \vee B}{F\,B}\;F \vee \mathcal{E}2$$

$$\frac{T\,A \wedge B}{T\,A}\;T \wedge \mathcal{E}1 \qquad \frac{T\,A \wedge B}{T\,B}\;T \wedge \mathcal{E}2 \qquad \frac{\begin{array}{l}F\,A \wedge B\\ T\,A\end{array}}{F\,B}\;F \wedge \mathcal{E}1 \qquad \frac{\begin{array}{l}F\,A \wedge B\\ T\,B\end{array}}{F\,A}\;F \wedge \mathcal{E}2$$

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

Figure 3: Elimination rules for signed sentences.

$$\frac{T\,A \vee A}{T\,A}\;T \vee \mathcal{E}3 \qquad \frac{F\,A \wedge A}{F\,A}\;F \wedge \mathcal{E}3$$

Figure 4: “Mingle” rules for $\vee$ and $\wedge$.

it is *open*. It can be easily shown that every closed intelim sequence can be extended to an *atomically closed* one, i.e. one that contains both $T\,p$ and $F\,p$ for some *atomic* formula p. An *intelim deduction* of a *signed* formula φ from the set of *signed* formulas X is an intelim sequence based on X ending with φ. An *intelim deduction* of an *unsigned* formula A from the set of *unsigned* formulas Γ is an intelim deduction of $T\,A$ from $\{T\,B \mid B \in \Gamma\}$. We say that a *signed* formula φ is *intelim-deducible* from a set of *signed* formulas X, if there is an intelim deduction of φ from X. We also say that an *unsigned* formula A is *intelim-deducible* from the set Γ of *unsigned* formulas if $T\,A$ is intelim-deducible from $\{T\,B \mid B \in \Gamma\}$. Figure 5 contains an example of an intelim deduction that proves the sequent:

$$\neg u \vee s, u, s \to r \vee \neg u, \neg(q \wedge r), p \to q, \neg p \to t, t \wedge r \vee z \to v \vdash_0 v.$$

Finally, an *intelim refutation* of a set of formulas Γ is a closed intelim sequence based on $\{T\,B \mid B \in \Gamma\}$. When there is an intelim refutation of Γ, we say that Γ is *intelim-inconsistent*.

The notion of intelim deduction is adequate for the logic $\vdash_0$ semantically presented in the previous section and allows for a particularly strong normalization procedure. These properties are stated in Theorems 5.1 and 5.2 whose proofs can be adapted from [7] and [6].

Theorem 5.1 *For every finite set* Γ *of formulas and every formula* A*:*

1. $\Gamma \vdash_0 A$ *if and only if* A *is intelim-deducible from* Γ*;*
2. Γ *is* 0*-depth inconsistent if and only if* Γ *is intelim-inconsistent.*

Say that a signed formula ψ occurring in an intelim deduction π of A from Γ is *redundant* if $\psi \neq T\,A$ and ψ is not used as a premise of any application of an intelim rule in π. Call *non-redundant reduction* of π the intelim deduction of A from Γ obtained from π by removing the redundant signed formulas. Then, we say that an intelim deduction of A from Γ is *regular* if its non-redundant reduction is open, and *irregular* otherwise. In other words, irregular deductions of A from Γ are deductions in which information which is explicitly inconsistent, and recognized as such, has been used to obtain a given conclusion. Finally, say that an intelim deduction π of A from Γ enjoys the *subformula property* if every signed formula occurring in π has the form $T\,B$ or $F\,B$, where B is a subformula of A or of some formula in Γ.

The following theorem states a basic normalization property for intelim deductions (the *length* $|\pi|$ of an intelim sequence π is defined as the total number of symbols occurring in π):

Theorem 5.2 *Let* Γ *be a finite set of formulas. Then:*

1. *every regular intelim deduction* π *of* A *from* Γ *can be transformed into an intelim deduction* π' *of* A *from* Γ *such that (i)* π' *enjoys the subformula property, and (ii)* $|\pi'| \leq |\pi|$*;*

1	$T\,\neg u \vee s$	
2	$T\,u$	
3	$T\,s \rightarrow r \vee \neg u$	
4	$T\,\neg(q \wedge r)$	
5	$T\,p \rightarrow q$	
6	$T\,\neg p \rightarrow t$	
7	$T\,t \wedge r \vee z \rightarrow v$	
8	$F\,\neg u$	$F\neg\mathcal{I}(2)$
9	$T\,s$	$T \vee \mathcal{E}1(1, 8)$
10	$T\,r \vee \neg u$	$T \rightarrow \mathcal{E}1(3, 9)$
11	$T\,r$	$T \vee \mathcal{E}2(10, 8)$
12	$F\,q \wedge r$	$T\neg\mathcal{E}(4)$
13	$F\,q$	$F \wedge \mathcal{E}2(12, 11)$
14	$F\,p$	$T \rightarrow \mathcal{E}2(5, 13)$
15	$T\,\neg p$	$T\neg\mathcal{I}(14)$
16	$T\,t$	$T \rightarrow \mathcal{E}1(6, 15)$
17	$T\,t \wedge r$	$T \wedge \mathcal{I}(11, 16)$
18	$T\,t \wedge r \vee z$	$T \vee \mathcal{I}1(17)$
19	$T\,v$	$T \rightarrow \mathcal{E}1(7, 18)$

Figure 5: An example of intelim deduction.

1	$T\,p$	
2	$T\,\neg p$	
3	$F\,p$	$T\neg\mathcal{E}(2)$
4	$T\,p \vee q$	$T\vee\mathcal{I}1(1)$
5	$T\,q$	$T\vee\mathcal{E}1(4,3)$

Figure 6: An intelim deduction of q from $\{p, \neg p\}$.

2. *every intelim refutation π of Γ can be transformed into an intelim refutation π' of Γ such that (i) π' enjoys the subformula property, and (ii) $|\pi'| \leq |\pi|$.*

Theorem 5.2 suggests that irregular intelim deductions may not be normalizable. And this is indeed the case, as shown by the intelim deduction of q from $\{p, \neg p\}$ shown in Figure 6, which cannot be normalized. In some sense, however, normalization fails exactly when it ought to, that is, when we have already obtained a closed intelim sequence and try to use two signed formulas that explicitly contradict each other in order to obtain a certain "conclusion" from them. But, to quote Michael Dummett again, "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".[25] On the other hand, intelim refutations are always normalizable.[26] Notice that Theorem 5.2 marks a clear distinction from normalization theorems that can be proved for full classical (or intuitionistic) logic, where normal proofs may be longer, and sometimes exponentially longer, than non-normal ones (the same holds true for cut-free proofs versus cut-based proofs in the sequent calculus).

Theorem 5.2 paves the way for efficient decision procedures. One of them is presented in [7], and improved on in [6], where it is used to show the following:

Theorem 5.3 *Intelim-deducibility and intelim-refutability are tractable problems. Whether a formula A is intelim-deducible from a finite set Γ of formulas and whether a finite set Γ of formulas is intelim-refutable are both questions that can be decided in time $\mathcal{O}(n^2)$.*

We now turn to the general proof-theoretic presentation of the logics $\vdash_k$ with $k > 0$.

Given a signed formula SA, with S equal to T or F, let us denote by $\overline{S}A$ its *conjugate*, that is, $F\,A$ if $S = T$ and $T\,A$ if $S = F$. An *intelim sequence of depth k*, with $k > 0$ based on a set X of signed formulas is a sequence of

---

[25] [12], p. 209.

[26] If we add a new structural rule corresponding to the *ex-falso quodlibet* principle — to the effect that an arbitrary signed formula can be deduced from two contradictory ones — and modify the deducibility relation accordingly, then the subformula property holds in general.

signed formulas $\varphi_1, \dots, \varphi_n$ such that for each element φ_i, with $i \leq n$, one of the following conditions is satisfied:

1. $\varphi_i \in X$,
2. φ_i is intelim($k-1$)-deducible from $\varphi_1, \dots, \varphi_{i-1}, T\,p$ and $\varphi_1, \dots, \varphi_{i-1}, F\,p$, for some *atomic* formula p,
3. φ_i is intelim($k-1$)-deducible from $\varphi_1, \dots, \varphi_{i-1}, S\,p$, for some *atomic* formula p such that $\varphi_1, \dots, \varphi_{i-1}, \overline{S}\,p$ is intelim($k-1$)-inconsistent.[27]

Remark 5.4 *Observe that, since intelim deducibility (at any depth) is clearly monotonic, clause 2 in the above definition covers the case in which φ_i is intelim($k-1$) deducible from $\varphi_1, \dots, \varphi_{i-1}$ alone, with no need for the virtual information concerning the atomic formula p. As a result any intelim(k) sequence, with $k \in \mathbb{N}$, is also an intelim($k+1$) sequence.*

An *intelim(k) deduction* of a *signed* formula φ from the set of *signed* formulas X is an intelim(k) sequence based on X ending with φ. An *intelim(k) deduction* of an *unsigned* formula A from the set of *unsigned* formulas Γ is an intelim(k) deduction of $T\,A$ from $\{T\,B \mid B \in \Gamma\}$. We say that a *signed* formula φ is *intelim(k)-deducible* from a set of *signed* formulas X, if there is an intelim(k) deduction of φ from X. We also say that the *unsigned* formula A is *intelim(k)-deducible* from the set Γ of *unsigned* formulas if $T\,A$ is intelim(k)-deducible from $\{T\,B \mid B \in \Gamma\}$.

An intelim(k) sequence $\varphi_1, \dots, \varphi_n$ is *closed* if, for some atomic p, there are closed intelim($k-1$) sequences for both $\varphi_1, \dots, \varphi_n, T\,p$ and $\varphi_1, \dots, \varphi_n, F\,p$. (The notion of a closed intelim(0) sequence is the same as that of a closed intelim sequence.) Finally, an *intelim(k) refutation* of a set of formulas Γ is a closed intelim(k) sequence based on $\{T\,B \mid B \in \Gamma\}$. When there is an intelim(k) refutation of Γ, we say that Γ is *intelim(k) inconsistent*.

An intelim(k) sequence can be conveniently represented using boxes for the auxiliary sequences of depth greater than 0 that may be needed to establish that a signed formula φ_i can be appended to the sequence by virtue of clauses 2–3 of the above definition. To represent such a step, two parallel boxes may be opened (when necessary, see Remark 5.4) that contain, respectively, the auxiliary ($k-1$)-depth sequences based on $\varphi_1, \dots, \varphi_{i-1}, T\,p$ and $\varphi_1, \dots, \varphi_{i-1}, F\,p$, for some atomic p. The usual scoping rules for boxes are employed here: each formula occurring in a box can be used, as premiss of a rule application, in every box contained in it and *cannot* be used in any other box. The whole deduction should be regarded as being contained in a root box that contains all the others, although we shall not usually draw the borders of this most external box. Fig. 7 and Fig. 8 show two examples of intelim deductions of depth 1 and 2 respectively.

Theorem 5.5 *For every finite set Γ of formulas and every formula A:*

---

[27] See Remark 4.5 above.

1	$T\,p \vee q$	Assumption		
2	$T\,p \to r \vee s$	Assumption		
3	$T\,q \to r \vee s$	Assumption		
4	$T\,r \to t$	Assumption		
5	$T\,s \to u$	Assumption		
6	$T\,p$		$F\,p$	
7	$T\,r \vee s$	$T \to \mathcal{E}1(2,6)$	$T\,q$	$T \vee \mathcal{E}2(1,6)$
			$T\,r \vee s$	$T \to \mathcal{E}1(3,7)$
9	$T\,r \vee s$			
10	$T\,r$		$F\,r$	
11	$T\,t$	$T \to \mathcal{E}1(10,4)$	$T\,s$	$T \vee \mathcal{E}1(9,10)$
12	$T\,t \vee u$	$T \vee \mathcal{I}1(11)$	$T\,u$	$T \to \mathcal{E}1(5,11)$
			$T\,t \vee u$	$T \vee \mathcal{I}2(12)$
14	$T\,t \vee u$			

Figure 7: An intelim deduction of depth 1.

1. *$\Gamma \vdash_k A$ if and only if A is intelim(k)-deducible from Γ;*
2. *Γ is k-depth inconsistent if and only if Γ is intelim(k)-inconsistent.*

A normalization theorem analogous to Theorem 5.2 can be shown for intelim(k) deductions. The restriction of virtual information to *atomic* formulas, in the definition of intelim(k) sequences, as well as in the semantical definition of $\vdash_k$, ensures that the normalization procedure is essentially the same as in the 0-depth case and does not increase the length of proofs, since spurious atomic formulas are easily removed with no loss of deductive power. The price is that, unlike $\vdash_0$, the consequence relations$\vdash_k$, with $k > 0$, are not substitution-invariant. On the other hand, the family $\{\vdash_k\}_{k \in \mathbb{N}}$, as a whole, *is* substitution invariant in the sense that $\Gamma \vdash_k A$ implies $\sigma(\Gamma) \vdash_{k+j} \sigma(A)$ for some j depending on the substitution σ (which is obvious, given Theorem 4.7 above and the fact that classical logic is substitution-invariant). To make each single $\vdash_k$ substitution-invariant, one should modify the definitions so as to allow for the use of virtual information of arbitrary complexity. In this case, normalization can still be achieved, but the length of normal proofs can be exponentially longer than that of non-normal ones. A normalization theorem for an intelim system with unbounded use of virtual information, which is therefore sound and complete for

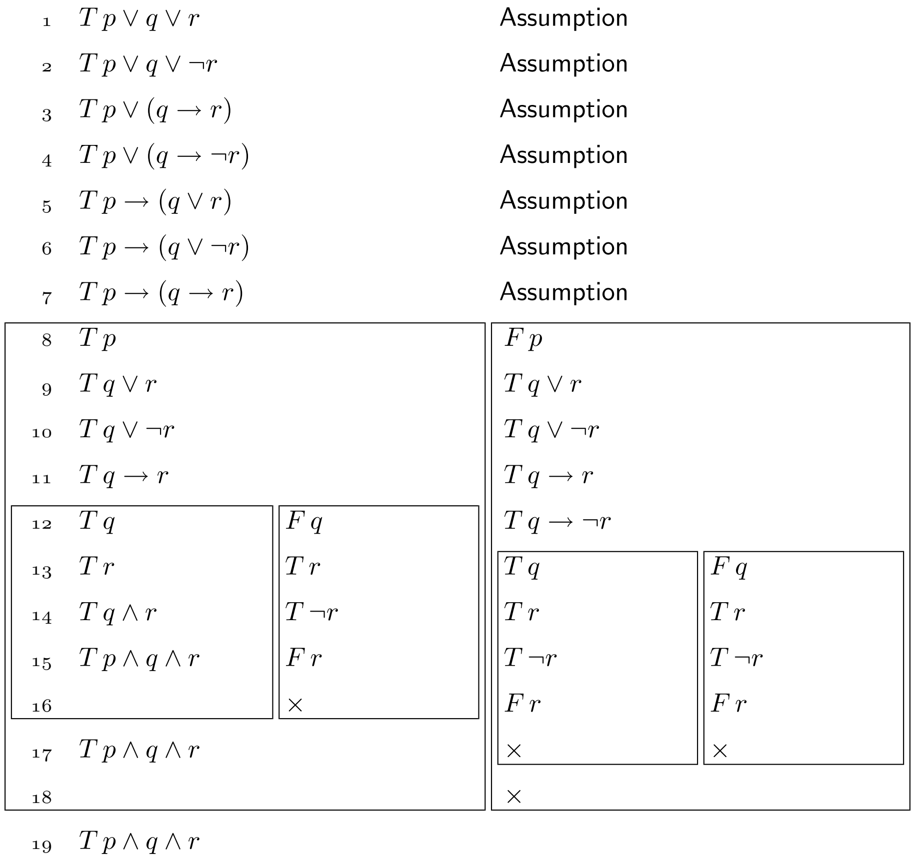

Figure 8: An intelim deduction of depth 2. The justification of the inference steps is omitted for reasons of space. The symbol "×" indicates a closed intelim sequence.

full classical propositional logic, is given in [5].[28] Finally, it can be shown that the relations of intelim(k)-deducibility and intelim(k)-refutability admit of a polynomial time decision procedure.

Theorem 5.6 *Intelim(k)-deducibility and intelim(k)-refutability are tractable problems for every fixed k. Whether a formula A is intelim(k)-deducible or whether a finite set* Γ *of formulas is intelim(k)-refutable can be decided in time* $\mathcal{O}(n^{2k+2})$*.*

---

[28] In the cited paper the intelim rules are presented for *unsigned* formulas. The difference is largely a matter of taste, but we now believe that the use of signed formulas is conceptually more transparent when dealing with a classical-like negation operator.

6 Concluding remarks

We have argued that Depth-Bounded Boolean Logics can be profitably adopted as a basis for constructing epistemic, doxastic and information logics that are not affected by the problem of logical omniscience, but a good deal of work has yet to be done in this direction. First, one should devise an alternative to the standard characterization of the propositional attitudes $\Box_i$ based on possible worlds. Our analysis suggests that, for this task, possible worlds could be replaced by feasible depth-bounded information states, as defined in Section 4, and that the semantics of $\Box_i$ could be fixed via some sort of invariance over information states, depending on the propositional attitude under consideration. For example, knowledge could be characterized as information that is persistent through information states that may change in a non-monotonic way. Moreover, the incremental characterization of classical logical consequence that is typical of Depth-Bounded Boolean Logics may allow for the design of multimodal systems where different (human or artificial) agents may be endowed with unequal deductive power, maybe on the basis of the unequal computing resources available to them, but share the same understanding of the logical operators (displayed in the common 0-depth logic). Another interesting topic of future research is the treatment of inconsistent information. As remarked in Sections 4 and 5, Depth-Bounded Boolean Logics seem likely to provide an interesting way out of the qualms about the *ex-falso quodlibet* principle of classical logic. Furthermore, the failure of substitution-invariance for the logics of depth greater than 0 prompts for an exploration of the possible remedies outlined in Section 5. Last, but not least, the possibility of augmenting Depth-Bounded Boolean Logics with suitable quantifier rules, in a way that preserve their useful properties (first of all tractability), is an important topic of future investigation.

References

[1] Sergei Artemov and Roman Kuznets. Logical omniscience as a computational complexity problem. In *TARK '09: Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 14–23, New York, NY, USA, 2009. ACM.

[2] Marco Cadoli and Marco Schaerf. Approximate reasoning and non-omniscient agents. In *TARK '92: Proceedings of the 4th conference on Theoretical aspects of reasoning about knowledge*, pages 169–183, San Francisco, CA, USA, 1992. Morgan Kaufmann Publishers Inc.

[3] Stephen A. Cook. The complexity of theorem-proving procedures. In *STOC. Conference Record of Third Annual ACM Symposium on Theory of Computing, 1971, Shaker Heights, Ohio, USA*, pages 151–158. ACM, 1971.

[4] James M. Crawford and David W. Etherington. A non-deterministic semantics for tractable inference. In *AAAI/IAAI*, pages 286–291, 1998.

[5] Marcello D'Agostino. Classical natural deduction. In Sergei N. Artëmov, Howard Barringer, Artur S. d'Avila Garcez, Luís C. Lamb, and John Woods, editors, *We Will Show Them! (1)*, pages 429–468. College Publications, 2005.

[6] Marcello D'Agostino, Marcelo Finger, and Dov M. Gabbay. Depth-bounded boolean logics. Submitted. Preliminary version available on request.

[7] Marcello D'Agostino and Luciano Floridi. The enduring scandal of deduction. Is propositionally logic really uninformative? *Synthese*, 167:271–315, 2009.

[8] Mukesh Dalal. Anytime families of tractable propositional reasoners. In *Proceedings of the Fourth International Symposium on AI and Mathematics (AI/MATH-96), 42–45.*, 1996.

[9] Mukesh Dalal. Anytime families of tractable propositional reasoners. *Annals of Mathematics and Artificial Intelligence*, 22:297–318, 1998.

[10] Michael Dummett. *Elements of Intuitionism.* Clarendon Press, Oxford, 1977.

[11] Michael Dummett. *Truth and Other Enigmas.* Duckworth, 1978.

[12] Michael Dummett. *The Logical Basis of Metaphysics.* Duckworth, 1991.

[13] Ronald Fagin, Joseph Y. Halpern, and Moshe Y. Vardi. A nonstandard approach to the logical omniscience problem. *Artificial Intelligence*, 79:203–240, 1995.

[14] Marcelo Finger. Polynomial approximations of full propositional logic via limited bivalence. In *9th European Conference on Logics in Artificial Intelligence (JELIA 2004)*, volume 3229 of *Lecture Notes in Artificial Intelligence*, pages 526–538. Springer, 2004.

[15] Marcelo Finger and Dov Gabbay. Cut and pay. *Journal of Logic, Language and Information*, 15(3):195–218, 2006.

[16] Luciano Floridi. Scepticism and animal rationality: the fortune of chrysippus' dog in the history of western thought. *Archiv für Geschichte der Philosophie*, 79(1):27–57, 1997.

[17] Luciano Floridi. The logic of being informed. *Logique et Analyse*, 49(196):433–460, 2006.

[18] Alan M. Frisch. Inference without chaining. In *IJCAI'87: Proceedings of the 10th international joint conference on Artificial intelligence*, pages 515–519, San Francisco, CA, USA, 1987. Morgan Kaufmann Publishers Inc.

[19] G. Gentzen. Unstersuchungen über das logische Schliessen. *Math. Zeitschrift*, 39:176–210, 1935. English translation in [35].

[20] Joseph Y. Halpern. Reasoning about knowledge: a survey. In Dov M. Gabbay, C. J. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 4, pages 1–34. Clarendon Press, Oxford, 1995.

[21] Saul A. Kripke. Semantical analysis of intuitionistic logic I. In J. Crossley and M. A. E. Dummett, editors, *Formal Systems and Recursive Functions*, pages 92–130. North-Holland Publishing, Amsterdam, 1965.

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

[23] Hector J. Levesque. Logic and the complexity of reasoning. *Journal of Philosophical Logic*, 17(4):355–389, November 1988.

[24] John-Jules C. Meyer. Modal epistemic and doxastic logic. In Dov M. Gabbay and Franz Guenthner, editors, *Handbook of Philosophical Logic*, volume 10, pages 1–38. Kluwer Academic Publishers, 2nd edition, 2003.

[25] Enrico Moriconi. Normalization and meaning theory. *Epistemologia*, XXIII:281–304, 2000.

[26] Rohit Parikh. Sentences, belief and logical omniscience, or what does deduction tell us? *The Review of Symbolic Logic*, 1(4), 2008.

[27] D. Prawitz. *Natural Deduction. A Proof-Theoretical Study.* Almqvist & Wilksell, Uppsala, 1965.

[28] Dag Prawitz. Meaning approached via proofs. *Synthese*, 148(3):507–524, February 2006.

[29] Giuseppe Primiero. An epistemic logic for being informed. *Synthese*, 167(2):363–389, 2009.

[30] Marco Schaerf and Marco Cadoli. Tractable reasoning via approximation. *Artificial Intelligence*, 74(2):249–310, April 1995.

[31] Giacomi Sillari. Quantified logic of awareness and impossible possible worlds. *Review of Symbolic Logic*, 1(4):1–16, 2008.

[32] Giacomo Sillari. Models of awareness. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Games and Decisions*, pages 209–240. University of Amsterdam, 2008.

[33] Richard Statman. Intuitionistic propositional logic is polynomial-space complete. *Theoretical Computer Science*, 9:67–72, 1979.

[34] Vitezslav Svejdar. On the polynomial-space completeness of intuitionistic propositional logic. *Archive for Mathematical Logic*, 42(7):711–716, 2003.

[35] M. Szabo, editor. *The Collected Papers of Gerhard Gentzen.* North-Holland, Amsterdam, 1969.

[36] N. Tennant. *Natural Logic.* Edinburgh University Press, 1990.

[37] Johan van Benthem. The information in intuitionistic logic. *Synthese*, 167(2):251–270, March 2009.