

THE LOGIC OF BEING INFORMED

LUCIANO FLORIDI*

Abstract

One of the open problems in the philosophy of information is whether there is an *information logic* (*IL*), different from *epistemic* (*EL*) and *doxastic logic* (*DL*), which formalises the relation “*a* is informed that *p*” ($I_a p$) satisfactorily. In this paper, the problem is solved by arguing that the axiom schemata of the normal modal logic (*NML*) *KTB* (also known as *B* or *Br* or Brouwer’s system) are well suited to formalise the relation of “being informed”. After having shown that *IL* can be constructed as an informational reading of *KTB*, four consequences of a *KTB*-based *IL* are explored: information overload; the veridicality thesis ($I_a p \rightarrow p$); the relation between *IL* and *EL*; and the $Kp \rightarrow Bp$ principle or *entailment property*, according to which knowledge implies belief. Although these issues are discussed later in the article, they are the motivations behind the development of *IL*.

Introduction

As anyone acquainted with modal logic (*ML*) knows, epistemic logic (*EL*) formalises the relation “*a* knows that *p*” ($K_a p$), whereas doxastic logic (*DL*) formalises the relation “*a* believes that *p*” ($B_a p$). One of the open problems in the philosophy of information (Floridi, 2004c) is whether there is also an information logic (*IL*), different from *EL* and from *DL*, that formalises the relation “*a* is informed that *p*” ($I_a p$) equally well. The keyword here is “equally” not “well”. One may contend that *EL* and *DL* do not capture the relevant relations very well or even not well at all. Hocutt (1972), for example, provides an early criticism. Yet this is not the point here, since all I wish to argue in this paper is that *IL* can do for “being informed” what *EL* does for “knowing” and *DL* does for “believing”. If one objects to the last two, one may object to the first as well, yet one should not object to it more.

* Address for correspondence: St Cross College, OX1 3LZ, Oxford, UK..

The proposal developed in the following pages is that the normal modal logic (*NML*) *KTB* (also known as *B*, *Br* or Brouwer’s system¹) is well suited to formalise the relation of “being informed”, and hence that *IL* can be constructed as an informational reading of *KTB*. The proposal is in three sections.

In section one, several meanings of “information” are recalled, in order to focus only on the “cognitive” sense. Three main ways in which one may speak of a “logic of (cognitive) information” are then distinguished. Only one of them is immediately relevant here, namely, “*a* is informed that *p*” as meaning “*a* holds the information that *p*”. These clarifications are finally used to make precise the specific question addressed in the rest of the paper.

In section two, the analysis of the informational relation of “being informed” provides the specifications to be satisfied by its accurate formalisation. It is then shown that *KTB* successfully captures the relation of “being informed”.

In section three, once it is established that there is an *IL* different from *EL* and from *DL*, four consequences of a *KTB*-based *IL* are briefly explored: information overload; the veridicality thesis ($Ip \rightarrow p$); the relation between *IL* and *EL*; and the entailment property or $Kp \rightarrow Bp$ principle, according to which knowledge implies belief. Although they are discussed later in the article, these four issues are the motivations behind the development of *IL*.

In the conclusion, I sketch some of the work that lies ahead.

Throughout the paper the ordinary language of classical, propositional calculus (*PC*) and of normal, propositional modal logic (see for example: Girle, 2000) will be presupposed. Implication (\rightarrow) is used in its “material” sense; the semantics is Kripkean; Greek letters are metalinguistic, propositional variables ranging over well-formed formulae of the object language of the corresponding *NML*; and until section 2.6 attention is focused only on the axiom schemata of the *NMLs* in question.

1. *Three Logics of Information*

“Information” may be understood in many ways, e.g. as signals, natural patterns or nomic regularities, as instructions, as content, as news, as synonymous with data, as power or as an economic resource and so forth. It is notoriously controversial whether even most of these senses of “information”

¹The name was assigned by Becker (1930). As Goldblatt (2003) remarks: “The connection with Brouwer is remote: if ‘not’ is translated to ‘impossible’ ($\rightarrow\Diamond$), and ‘implies’ to its strict version, then the intuitionistically acceptable principle $p \rightarrow \neg\neg p$ becomes the Brouwersche axiom”. For a description of *KTB* see Hughes & Cresswell (1996).

might be reduced to a fundamental concept². However, the sort of “information” that interests us here is arguably the most important. It is “information” as semantic content that, on one side, concerns some state of a system, and that, on the other side, allows the elaboration of an agent’s propositional knowledge of that state of the system. It is the sense in which Matthew is informed that p , e.g. that “the train to London leaves at 10.30 am”, or about the state of affairs f expressed by p , e.g. the railway timetable. In the rest of the paper, “information” will be discussed only in this intuitive sense of declarative, objective and semantic content that p or about f (Floridi, 2005a). This sense may loosely be qualified as “cognitive”, a neutral label useful to refer here to a whole family of relations expressing propositional attitudes, including “knowing”, “believing”, “remembering”, “perceiving” and “experiencing”. Any “non-cognitive” sense of “semantic information” will be disregarded³. The scope of our inquiry can now be narrowed by considering the logical analysis of the cognitive relation “ a is informed that p ”. Three related yet separate features of interest need to be further distinguished, namely

a) how p may be informative for a . For example, the information that p may or may not be informative depending on whether a is already informed that ($p \rightarrow q$). This aspect of information — the informativeness of a message — raises issues of e.g. novelty, reliability of the source and background information. It is a crucial aspect related to the quantitative theory of semantic information (Bar-Hillel & Carnap, 1953; Bar-Hillel, 1964; Floridi, 2004d), to the logic of transition states in dynamic system, that is, how change in a system may be informative for an observer (Barwise & Seligman, 1997) and to the theory of levels of abstraction at which a system is being considered (Floridi & Sanders, 2004, forthcoming);

b) the process through which a becomes informed that p . The informativeness of p makes possible the process that leads from a ’s uninformed (or less informed) state A to a ’s (more) informed state B . Upgrading a ’s state A to a state B usually involves receiving the information that p from some external source S and processing it. It implies that a cannot be informed

²For an overview see Floridi (2004a, 2005b). Personally, I am very sceptical about attempts to find a unified theory of information and hence a unique logic that would capture all its interesting features.

³There are many plausible contexts in which a stipulation (“let the value of $x = 3$ ” or “suppose we discover the bones of a unicorn”), an invitation (“you are cordially invited to the college party”), an order (“close the window!”), an instruction (“to open the box turn the key”), a game move (“1.e2-e4 c7-c5” at the beginning of a chess game) may be correctly qualified as kinds of information understood as semantic content. These and other similar, non-cognitive meanings of “information” (e.g. to refer to a music file or to a digital painting) are not discussed in this paper, where semantic information is taken to have a declarative or factual value i.e. it is suppose to be correctly qualifiable alethically.

that p unless a was previously uninformed that p . And the logical relation that underlies this state transition raises important issues of timeliness and cost of acquisition, for example, and of adequate procedures of information processing, including introspection and metainformation, as we shall see. It is related to communication theory (Shannon & Weaver, 1949, rep. 98), temporal logic, updating procedures (Gärdenfors, 1988), and recent trends in dynamic epistemic logic (Baltag & Moss, 2004);

c) the state of the epistemic agent a , insofar as a holds the information that p . This is the *statal* condition into which a enters, once a has acquired the information (*actional* state of being informed) that p . It is the sense in which a witness, for example, is informed (holds the information) that the suspect was with her at the time when the crime was committed. The distinction is standard among grammarians, who speak of passive verbal forms or states as “statal” (e.g. “the door was shut (state) when I last checked it”) or “actional” (e.g. “but I don’t know when the door was shut (act)”)⁴. Here, we are interested only in the statal sense of “is informed”. This sense (c) is related to cognitive issues and to the logical analysis of an agent’s “possession” of a belief or a piece of knowledge.

Point (a) requires the development of a logic of “being informative”; (b) requires the development of a logic of “becoming informed”; and (c) requires the development of a logic of “being informed (i.e. holding the information)”. Work on (a) and (b) is already in progress. Mark Jago, in this issue, Allo (2005) and Sanders (forthcoming) develop lines of research complementary to this paper. In the following pages, I shall be concerned with (c) and seek to show that there is a logic of information comparable, for adequacy, flexibility and usefulness, to *EL* and *DL*.

Our problem can now be formulated more precisely. Let us concentrate our attention on the most popular and traditional *NML*, obtainable through the analysis of some of the well-known characteristics of the relation of accessibility (reflexivity, transitivity etc.). These fifteen⁵ *NMLs* range from the weakest *K* to the strongest *S5* (see below Figure 1). They are also obtainable through the combination of the usual axiom schemata of *PC* with the fundamental modal axiom schemata (see below Figure 2). Both *EL* and *DL* comprise a number of cognitively interpretable *NML*, depending on the sets of axioms that qualify the corresponding *NML* used to capture the relevant “cognitive” notions. If we restrict our attention to the six most popular *EL*

⁴I owe to Christopher Kirwan this very useful clarification; in a previous version of this paper I had tried to reinvent it, but the wheel was already there.

⁵The number of *NMLs* available is infinite. I am grateful to Timothy Williamson and John Halleck who kindly warned me against a misleading wording in a previous version of this paper.

and *DL* — those based on systems *KT*, *S4*, *S5* and on systems *KD*, *KD4*, *KD45* respectively — the question about the availability of an information logic can be rephrased thus: among the popular *NMLs* taken into consideration, is there one, not belonging to *KT*, *S4*, *S5*, *KD*, *KD4*, *KD45*, which, if cognitively interpreted, can successfully capture and formalise our intuitions regarding “*a* is informed that *p*” in the (c) sense specified above?

A potential confusion may be immediately dispelled. Of course, the logical analysis of the cognitive relation of “being informed” can sometimes be provided in terms of “knowing” or “believing”, and hence of *EL* or *DL*. This is not in question, for it is trivially achievable, insofar as “being informed” can sometimes be correctly treated as synonymous with “knowing” or “believing”. We shall also see in § 3.3 that *IL* may sometime overlap with *EL*. The interesting problem is whether “being informed” may show properties that typically (i.e., whenever the overlapping would be unjustified, see § 3.3) require a logic different from *EL* and *DL*, in order to be modelled accurately. The hypothesis defended in the following pages is that it does and, moreover, that this has some interesting consequences for our understanding of the nature of the relation between “knowing” and “believing”.

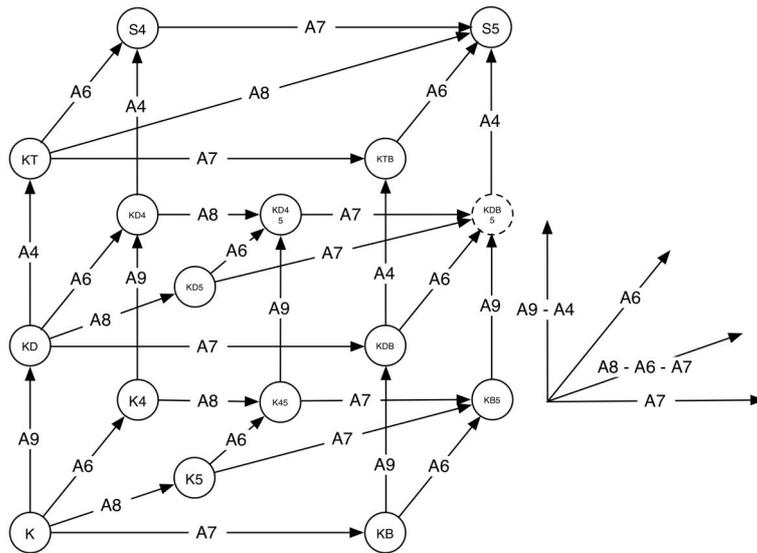


Figure 1: Fifteen Normal Modal Logics

Note that *KDB5* is a “dummy” system: it is equivalent to *S5* and it is added to the diagram just for the sake of elegance.

| | |
|--------------|--|
| Synonymous | Equivalent axiomatic systems |
| T = M = KT | B = TB |
| B = Br = KTB | KB5 = KB4, KB45 |
| D = KD | S5 = T5, T45, TB4, TB5, TB45, DB4, DB5, DB45 |

2. Modelling “Being Informed”

Let us interpret the modal operator \Box as “is informed that”. We may then replace the symbol \Box with I for “being informed”, include an explicit reference to the informed agent a , and write

$$\Box p = I_a p \quad (1)$$

to mean a is informed (holds the information) that p ⁶.

As customary, the subscript will be omitted whenever we shall be dealing with a single, stand-alone agent a . It will be reintroduced in § 2.4, when dealing with multiagent IL . Next, we can then define \Diamond in the standard way, thus

$$U_a p =_{def} \neg I_a \neg p \quad (2)$$

to mean a is uninformed (is not informed, does not hold the information) that $\neg p$; or for all a 's information (given a 's information base), it is possible that p .

Simplifying, a 's information base can be modelled by representing it as a dynamic set D_a of sentences of a language L ⁷. The intended interpretation

⁶ A *de re* interpretation is obtainable by interpreting $I_a p$ as “there is the information that p ”.

⁷ Dynamic sets are an important class of data structures in which sets of items, indexed by keys, are maintained. It is assumed that the elements of the dynamic set contain a field (called the key) by whose value they can be ordered. The phone directory of a company is a simple example of a dynamic set (it changes over time), whose key might be “last name”. Dynamic sets can change over the execution of a process by gaining or losing elements. Of the variety of operations usually supported by a dynamic set, three are fundamental and will be assumed in this paper:

Search(S, k) = given a set S and a key value k , a query operation that returns a pointer x to an element in S such that $\text{key}[x] = k$, or nil if no such element belongs to S .

Insert(S, x) = an operation that augments the set S with the element x .

Delete(S, x) = an operation that removes an element pointed to by x from S (if it is there).

is that D_a consists of all the sentences, i.e. all the information, that a holds at time t . We then have that $I_a p$ means that $p \in D_a$, and $U_a p$ means that p can be uploaded in D_a while maintaining the consistency of D_a , that is, $U_a p$ means $\diamond(p \in D_a)$ “salva cohaerentiae”⁸. Note that a need not be committed, either doxastically (e.g. in terms of strengths of belief, Lenzen (1978)) or epistemically (e.g. in terms of degrees of certainty) in favour of any element in D_a .

Given that IL might actually overlap and hence be confused with EL or DL , the most plausible conjecture is that an IL that can capture our intuitions, and hence satisfy our requirements regarding the proper formalisation of I_p , will probably bear some strong resemblance to EL and DL . If there is any difference between these three families of cognitive logics it is likely to be identifiable more easily in terms of satisfaction (or lack thereof) of one or more axioms qualifying the corresponding NML . The heuristic assumption here is that, by restricting our attention to the fifteen $NMLs$ in question, we may be able to identify the one which best captures our requirements. It is a bit like finding where, on a continuous map, the logic of information may be placed: even if we succeed in showing that KTB is the right NML for our task, there is still an infinite number of neighbouring $NMLs$ extending KTB ⁹.

For ease of reference, the axiom schemata in question are summarised and numbered progressively in Figure 2, where φ , χ and ψ are propositional variables referring to any wff of PC.

Following Hintikka’s standard approach (Hintikka, 1962), a systematic way to justify the choice of some axiom schemata is by trying to identify a plausible interpretation of a semantics for the corresponding NML . We shall now consider the 12 axiom schemata and show that IL shares only some of them with EL and DL .

⁸ As Patrick Allo has noted in a personal communication, this can also be expressed in terms of safety of inclusion of p in D_a .

⁹ Many thanks to John Halleck for calling my attention to this point and to Miyazaki (2005).

| Label | Definitions of Axiom Schemata | Name of the Axiom or Corresponding NML | Frame Property |
|----------|---|--|----------------|
| A_1 | $\varphi \rightarrow (\chi \rightarrow \varphi)$ | 1st axiom of PC | |
| A_2 | $(\varphi \rightarrow (\chi \rightarrow \psi)) \rightarrow ((\varphi \rightarrow \chi) \rightarrow (\varphi \rightarrow \psi))$ | 2nd axiom of PC | |
| A_3 | $(\neg\varphi \rightarrow \neg\chi) \rightarrow (\chi \rightarrow \varphi)$ | 3rd axiom of PC | |
| A_4 | $\Box\varphi \rightarrow \varphi$ | KT or M, K2, veridicality | Reflexive |
| A_5 | $\Box(\varphi \rightarrow \chi) \rightarrow (\Box\varphi \rightarrow \Box\chi)$ | K, distribution, deductive cogency | Normal |
| A_6 | $\Box\varphi \rightarrow \Box\Box\varphi$ | 4, S4, K3, KK, reflective thesis or positive introspection | Transitive |
| A_7 | $\varphi \rightarrow \Box\Diamond\varphi$ | KTB, B, Br, Brouwer's axiom or Platonic thesis | Symmetric |
| A_8 | $\Diamond\varphi \rightarrow \Box\Diamond\varphi$ | S5, reflective, Socratic thesis or negative introspection | Euclidean |
| A_9 | $\Box\varphi \rightarrow \Diamond\varphi$ | KD, D, consistency | Serial |
| A_{10} | $(\Box(\varphi \rightarrow \chi) \rightarrow (\Box(\chi \rightarrow \psi) \rightarrow \Box(\varphi \rightarrow \psi)))$ | Single agent transmission | |
| A_{11} | $(\Box_x\Box_y\varphi \rightarrow \Box_x\varphi)$ | K4, multiagent transmission, or Hintikka's axiom | |

Figure 2

2.1. *IL* Satisfies A_1, A_2, A_3, A_5

Trivially, we may assume that *IL* satisfies the axioms $A_1 - A_3$. As for A_5 , this specifies that *IL* is distributive, as it should be. If an agent a is informed that $p \rightarrow q$, then, if a is informed that p , a is also informed that q . Note that, although this is entirely uncontroversial, it is less trivial. Not all “cognitive” relations are distributive. “Knowing”, “believing” and “being informed” are, as well as “remembering” and “recalling”. This is why Plato is able to argue that a “mnemonic logic”, which he seems to base on K4, may replace *DL* as a foundation for *EL*¹⁰. However, “seeing” and other experiential relations, for example, are not: if an agent a sees (in a non metaphorical sense) or hears or experiences or perceives that $p \rightarrow q$, it may still be false that, if a sees (hears etc.) p , a then also sees (hears etc.) q .

The inclusion or exclusion of the remaining seven axioms is more contentious. Although logically independent, the reasons leading to their inclusion or exclusion are not, and they suggest the following clustering. In § 2.2, *IL* is shown to satisfy not only A_9 (consistency) but also A_4 (veridicality). In § 2.3, it is argued that *IL* does not have to satisfy the two “reflective” axioms, that is A_6 and A_8 . And in § 2.4, it is argued that *IL* should satisfy the “transmissibility” axioms A_{10} and A_{11} . This will leave us with A_7 , to be discussed in § 2.5.

¹⁰On Plato's interpretation of knowledge as recollection see especially *Phaedo* 72e-75 and *Meno* 82b-85.

2.2. Consistency and Truth: *IL* satisfies A_9 and A_4

In *DL*, A_9 replaces the stronger A_4 , which characterizes *EL*: whereas p must be true for the epistemic agent a to know that p , the doxastic agent a only needs to be consistent in her beliefs. There are at least four reasons why *IL* should be characterized as satisfying A_9 :

1) A_9 specifies that, in *IL*, the informational agent a is consistent, but so can be our ordinary informed agent in everyday life: $Ip \rightarrow Up$. If a holds the information that the train leaves at 10.30 am then, for all a 's information, it is possible that the train leaves at 10.30 am, in other words, p can be uploaded in a 's information base D_a while maintaining the consistency of D_a ;

2) even if (1) were unconvincing, *IL* should qualify a as consistent at least normatively, if not factually, in the same way as *DL* does. If a holds the information that the train leaves at 10.30 am, then a should not hold the information that the train does not leave at 10.30 am. The point is not that doxastic or informational agents cannot be inconsistent¹¹, but that A_9 provides an information integrity constraint: inconsistent agents should be disregarded. Again, to appreciate the non-trivial nature of a normative approach to A_9 , consider the case of a “mnemonic logic”: it might be factually implausible and only normatively desirable to formalise “ a remembers that p ” as implying that, if this is the case, then a does not remember that $\neg p$. Matthew may remember something that actually never happened, or he might remember both p (that he left the keys in the car) and $\neg p$ (that he left the keys on his desk) and be undecided about which memory is reliable. Likewise, if a database contains the information that p it might, unfortunately, still contain also the information that $\neg p$, even if, in principle, it should not, because this would seriously undermine the informative nature of the database itself (see next point 3), and although it is arguable (because of A_4 , see below) that in such case either p or $\neg p$ fail to count as information;

3) objections against *IL* satisfying A_9 appear to be motivated by a confusion between “becoming informed” and “being informed”, a distinction emphasised in § 2.1. In the former case, it is unquestionable that a may receive and hence hold two contradictory messages (e.g., a may read in a printed timetable that the train leaves at 10.30 am, as it does, but a may also be told by b that the train does not leave at 10.30 am). However, from this it only follows that a has the information that the train leaves at 10.30 am, but since p and $\neg p$ erase each other's value as pieces of information for a , a may be unable, subjectively, to identify which information a holds. It does

¹¹ It might be possible to develop a modal approach to QC (quasi-classical) logic in order to weaken the integrity constraint, see Grant & Hunter (forthcoming).

not follow that a is actually informed both that the train leaves at 10.30 am and that it does not;

4) if IL satisfies the stronger A_4 then, *a fortiori*, IL satisfies A_9 . Accepting that IL satisfies A_9 on the basis of (1)–(3) is obviously not an argument in favour of the inclusion of A_4 . At most, it only defuses any argument against it based on the reasoning that, if IL did not satisfy A_9 , it would fail to satisfy A_4 as well. The inclusion of A_4 requires some positive support of its own, to which we now turn.

According to A_4 , if a is informed that p then p is true. Can this be right? Couldn't it be the case that one might be qualified as being informed that p even if p is false? The answer is in the negative, for the following reason. Including A_4 as one of IL axioms depends on whether p counts as information only if p is true. Now, some critics (Colburn (2000), Fox (1983), Dodig-Crnkovic (2005) and, among situation theorists, Devlin (1991)) may still be unconvinced about the necessarily veridical nature of information, witness the debate between Floridi (2004d) and Fetzer (2004). However, more recently, it was shown in Floridi (2005a) that the Dretske-Grice approach to the so-called standard definition of information as meaningful data¹² remains by far the most plausible. In short, p counts as information only if p is true because:

“[...] false information and mis-information are not kinds of information — any more than decoy ducks and rubber ducks are kinds of ducks” (Dretske, 1981, 45).

“False information is not an inferior kind of information; it just is not information” (Grice, 1989, 371).

As in the case of knowledge, truth is a necessary condition for p to qualify as information. In Floridi (2005a) this is established by proving that none of the reasons usually offered in support of the alethic neutrality of information is convincing, and then that there are several good reasons to treat information as encapsulating truth and hence to disqualify misinformation (that is, “false information”) as pseudo-information, that is, as not (a type of) information at all. The arguments presented there will not be rehearsed here, since it is sufficient to accept the conclusion that either one agrees that information encapsulates truth or (at least) the burden of proof is on her side.

Once the veridical approach to the analysis of semantic information is endorsed as the most plausible, it follows that, strictly speaking, to hold (exchange, receive, sell, buy, etc.) some “false information”, e.g. that the train

¹²Other philosophers who accept a truth-based definition of information are Barwise & Seligman (1997) and Graham (1999).

leaves at 11.30 am when in fact it leaves at 10.30 am, is to hold (exchange, receive, sell, buy, etc.) no information at all, only some semantic content (meaningful data). But then, a cannot hold the information (be informed) that p unless p is true, which is precisely what A_4 states. Mathew is not informed but misinformed that Italy lost the world cup in 2006 because Italy won it. And most English readers will gladly acknowledge that Matthew is informed about who won the world cup in 1966 only if he holds that England did.

The mistake — arguing that a may be informed that p even if p is false, and hence that IL should not satisfy A_4 — might arise if one confuses “holding the information that p ”, which we have seen must satisfy A_4 , with “holding p as information”, which of course need not, since an agent is free to believe that p qualifies as information even when p is actually false, and hence counts as mere misinformation.

As far as A_4 is concerned, “knowing that p ” and “being informed that p ” work in the same way. This conclusion may still be resisted in view of a final objection, which may be phrased as a dilemma: either the veridical approach to information is incorrect, and therefore IL should not satisfy A_4 , or it is correct, and therefore IL should satisfy A_4 , yet only because there is no substantial difference between IL and EL (information logic becomes only another name for epistemic logic). In short, the inclusion of A_4 among the axiom schemata qualifying IL is either wrong or trivial.

The objection is interesting but mistaken. So far, IL shares all its axiom schemata with EL , but information logic allows truth-encapsulation without epistemic collapse because there are two other axiom schemata that are epistemic but not informational. This is what we are going to see in the next section.

2.3. *No reflectivity: IL does not satisfy A_6 , A_8*

Let us begin from the most “infamous” of EL axiom schemata, namely A_6 . One way of putting the argument in favour of A_4 and against A_6 , is by specifying that the relation of “informational accessibility”¹³ H in the system that best formalises “being informed/holding the information that p ” is reflexive without being reflective, reflectivity being here the outcome of a transitive relation in a single agent context, that is, “introspection”, a rather more common label that should be used with some caution given its psychologistic overtones.

¹³The choice of the letter H is arbitrary, but it may graphically remind one of the H in Shannon’s famous equation and in the expression “holding the information that p ”.

If H were reflective (if the informational agent were introspective), II should support the equivalent of the KK or BB thesis, i.e., $Ip \rightarrow IIp$. However, the II thesis is not merely problematic, it is unjustified, for it is perfectly acceptable for a to be informed that p while being (even in principle) incapable of being informed that a is informed that p , without adopting a second, meta-informational approach to Ip . The distinction requires some unpacking.

On the one hand, “believing” and “knowing” (the latter here understood as reducible to some doxastic relation, but see § 3.4) are mental states that, arguably, in the most favourable circumstances, could implement a “privileged access” relation, and hence be fully transparent to the agents enjoying them, at least in principle and even if, perhaps, only for Cartesian agents. Yet KK or BB remain controversial (see Williamson (1999, 2000) for arguments against them). The point here is that defenders of the inevitability of the BB or KK thesis may maintain that, in principle, whatever makes it possible for a to believe (or to know) that p , is also what makes it possible for a to believe (or to know) that a believes (or knows) that p . B and BB (or K and KK) are two sides of the same coin. More precisely, if a believes (or knows) that p , this is an internal mental fact that could also be mentally accessible, at least in principle, to a Cartesian a , who can be presumed to be also capable of acquiring the relevant, reflective mental state of believing (knowing) that a believes (or knows) that p . Translating this into information theory, we are saying that either there is no communication channel that allows a to have a doxastic (or epistemic) access to p , or, if there is, this is also the same channel that, in principle, allows a to have a doxastic (or epistemic) access to a 's belief (or knowledge) that p . So a defender of the BB or KK thesis may argue that the mental nature of doxastic and epistemic states may allow BB and KK to piggyback on B and K without requiring a second, meta-channel of communication. Call this the single-channel nature of doxastic and epistemic relations.

On the other hand, all this does not hold true for “being informed/holding the information”, because the latter is a relation that does not necessarily require a mental or conscious state. Beliefs and knowledge (again, analysed doxastically) are in the head, information can be in the pocket. Less metaphorically, artificial and biological agents may hold the information that p , even if they lack a mind or anything resembling mental states concerning p . As a result, “being informed” should be analysed as providing an unprivileged access to some p . A dog is informed (holds the information) that a stranger is approaching the house only if a stranger is actually approaching the house, yet this does not imply that the dog is (or can even ever be) informed that he is informed that a stranger is approaching the house. Indeed, the opposite is true: animals do not satisfy any of the KK , BB or II thesis.

There are no Cartesian dogs. Likewise, a computer may hold the information that “the train to London leaves at 10.30 am”, but this, by itself, does not guarantee, even in principle, that the computer also holds the information that it holds the information about the train timetable, or we might be much closer to true AI than anybody ever seriously claimed. Finally, Matthew might have the information that “the train to London leaves at 10.30 am” written in a note in his pocket, and yet not be informed that he holds the information that p . Actually, Matthew might even have it stored in his brain, like Johnny Mnemonic, who in William Gibson’s homonymous novel is a mnemonic data courier hired to carry in his brain 320 gigabytes of crucial information to safety from the Pharmacom corporation. Note the difference: Johnny holds the information that he holds some precious information, yet this is like a black box, for he does not hold the information that he holds the information that p .

The distinction may be further clarified if, once again, it is translated into information theory. We are saying that either there is no communication channel that allows a to have an informational access to p , or, if there is, it is such that, even with a Cartesian agent placed in favourable circumstances (no malicious demon etc.), it may still fail to allow a to have an informational access to a ’s information that p . The possibly non-mental nature of informational states impedes II from piggybacking on I through the same channel of communication. An II relation requires in fact a second, meta-channel that allows an I relation between a and Ip , but then this channel too is not, by itself, reflective, since any III relation requires a third channel between I and IIp , and so forth. As far as reflectivity is concerned, “being informed that p ” is not like “believing that p ” or “knowing that p ” but more like “having recorded that p ” or “seeing that p ”. The former two require mental states, whose nature is such as to allow the possibility in principle of the BB -thesis or KK -thesis. The latter two do not require mental states and hence do not include the possibility of a reflective state: information, records and perceptual sensations do not come with metainformation or metarecords or metasensations by default, even in principle, although there may be a second layer of memory, or another channel of communication or of experience, that refers to the first layer of memory or the first channel of information or the more basic experience. Call this the double-channel nature of the information relation.

The distinction between the single and double channel of information may be compared to the distinction between a reflective sentence that speaks of itself (single-channel, e.g. “this sentence is written in English”) and a meta-sentence that speaks of another sentence (double-channel, e.g. “the following sentence is written in English” “the cat is on the mat”). Natural languages normally allow both. Consider Matthew again. He may have in his pocket a note about the first note about the train timetable, yet this would

be irrelevant, since it would just be another case of double-channel condition or meta-information. As Wittgenstein succinctly put it: “nothing in the visual field allows you to infer that it is seen by an eye” (*Tractatus*, 5.633). Likewise, nothing in a piece of information p allows you to infer that an information system that holds p also holds the information that it holds p (compare this to the fact that nothing in Matthew’s ignorance allows you to infer that he is aware of his ignorance), whereas nothing in a belief or in a piece of knowledge allows you to infer that a doxastic or epistemic agent holding that belief or enjoying that piece of knowledge does not also believe that she believes that p , or does not also know that she knows that p . Knowledge and beliefs are primed to become reflective, information is not.

Consider now the following two objections against the distinction between the single-channel (or reflective or conscious or introspective, depending on the technical vocabulary) nature of epistemic and doxastic states and the double-channel (or opaque, or unreflective, or unconscious) nature of informational states.

First, one may point out that the *II* thesis seems to be implemented by some artificial systems. Actually, there are so-called “reflective” artificial agents capable of proving the classic knowledge theorem (Brazier & Treur, 1999), variously known as the “muddy children” or the “three wise men” problem, the drosophila of epistemic logic and distributed AI¹⁴. The description, however, is only evocative. Artificial agents may appear to be “reflective” only because of some smart tricks played at the level of interfaces and human-computer interactions, or because of a multi-layer structure. In particular, architectures or programs for computational systems (of AI) and systems for machine learning are technically called “reflective” when they contain an accessible representation of themselves that can be used (by themselves) e.g. to monitor and improve their performance. But what is known as reflective computing is only a case of metaprogramming or a communication channel about another communication channel, precisely as expected¹⁵. It is what has been labelled above the double-channel nature of the *II* states. One may compare it to a dog being informed that (or barking because) another dog is informed that (or is barking because) a stranger is approaching.

¹⁴The classic version of the theorem is related to the Conway-Paterson-Moscow theorem and the Conway paradox (see Groenendijk et al. (1984), pp. 159–182 and Conway & Guy (1996)) and was studied, among others, by Barwise & Seligman (1997). For some indications on its history see Fagin & Halpern (1988), p. 13.

¹⁵Barklund (1995) and Costantini (2002) are two valuable surveys with further references to the “three wise men” problem. Note that, for those who object to *EL*, the axiomatization of the reasoning involved in the classic knowledge game may be done in standard (i.e. non-modal) FOL (McCarthy, 1971–87; McCarty, 1990); at the same time, it is amenable to a treatment in terms of BDI (Belief, Desire, Intention) architecture (Rao & Georgeff, 1991).

At a higher level of abstraction, the two dogs may form a single security system, but the possibility of multiagent (e.g. n dogs or n computational) informational systems does not contradict the deflationist view that “being informed” is not a reflective relation.

Second, the *II* thesis seems to be implemented at least by some human agents. In this case, the reply is that this is so only because information relations can be implemented by human agents by means of mental states, which can then lend their reflective nature to H . It is not H to be reflective; rather, if an agent a can manage Ip through some epistemic or conscious state, for example, then, if the corresponding relation of accessibility is reflective the *II* thesis may become acceptable.

To summarise with a slogan: information entails no iteration. The point concerning the rejection of A_6 is not that “being informed” cannot appear to be a reflective relation: this is possible because Ip may be the object of a second relation I (double-channel nature of *II*), when a is a multiagent system, or because Ip may be implemented mentally, when a is a human agent, and hence be subject to reflection, consciousness or introspection. The point concerning the rejection of A_6 is that doxastic and epistemic accessibility relations, interpreted as mental states, may require in principle only a single-channel communication to become reflective, so the *BB* and *KK* theses may be justifiable as limit cases; whereas H , by itself, is not necessarily mental, and requires a double-channel communication to become reflective. But then (a) the second channel may be absent even in the most idealised, animal or artificial agents, even in principle, and (b) in any case, we are developing a logic of the communication channel represented by the information relation between a and p , and this channel is not reflective. The conclusion is that adopting A_6 to formalise Ip would be a misrepresentation.

There is a further objection to the latter conclusion, but we shall see it in the next section, since it is connected to A_{10} . Before, we may briefly look at a consequence of the exclusion of A_5 by considering A_8 . This axiom too is reflective, and therefore equally inappropriate to qualify *IL*. From the fact that an artificial agent does not hold the information that $\neg p$ it does not follow that it holds the information that it is missing the information that $\neg p$. We shall return to this point in § 2.5. In this case too, the previous considerations regarding the possibility of meta-information (two-channel) or mental implementation of the information relation apply, but do not modify the conclusion.

2.4. *Transmissibility*: *IL* satisfies A_{10} and A_{11}

The exclusion of A_6 from the group of axiom schemata characterizing *IL* might still be opposed on the basis of the following reasoning: if the relation of informational accessibility is not interpreted as transitive, then it becomes

impossible to transfer information, but this is obviously absurd, so A_6 must be included.

The objection is flawed for three reasons. First, transmission does not necessarily depend on transitivity: in the KD-based DL , a belief may be transferred from a to b despite the fact that the axiom schema $(B_a\varphi \rightarrow B_aB_a\varphi)$ and the corresponding relation of accessibility do not characterize KD. Second, the exclusion of A_6 does not concern the exclusion of the transitivity of modal inferences formulated in A_{10} , which can easily be shown to be satisfied by IL . A_{10} is a theorem in all NML and, being a weaker version of the K -principle, it formulates a very weak property, unlike the KK -principle¹⁶. Third, the exclusion of A_6 concerns the transitive nature of H when a single, standalone agent is in question. It does not preclude the inclusion of A_{11} (Hintikka's axiom of transmission) in a multiagent context. On the contrary, in this case, A_{11} correctly characterizes IL , as it is perfectly reasonable to assume that $(I_aI_b\varphi \rightarrow I_a\varphi)$: if Matthew is informed that Jenny is informed that the train to London leaves at 10.30 am, then he is also informed that it does. Note that this is made possible also thanks to A_4 , i.e. the assumption that to be informed that p the latter must be true.

2.5. Constructing the information base: IL satisfies A_7

A_7 is the defining axiom schema of the system KTB . IL satisfies A_7 in the sense that, for any true p , the informational agent a not only cannot be informed that $\neg p$ (because of A_4), but now is also informed that a does not hold the information that $\neg p$.

The inclusion of A_7 in IL does not contradict the anti-reflective (i.e., zero introspection) constraint supported in § 2.3. True, the conclusion IUp can be inferred both from Up and from p . However, in the former case (A_8), one would have to assume some form of negative reflection (introspection), in order to allow the agent a to draw the inference from an informational state Up to the relevant, meta-informational state IUp . Whereas in the latter case (A_7) the inference is drawn externally, by an observer, who concludes that, for any piece of information p , one can attribute to the agent a the information that a does not have the information that $\neg p$, irrespective of whether a lacks any kind of reflection on a 's informational states. This holds true for theorems such as $II(p \vee \neg p)$, which are demonstrable in $KTB-IL$: as we saw in 2.3, the point here is not denying the possibility of meta-information — it is trivially true that computers can have information about their information that p , for example — but objecting against the reflective (introspective, single-channel) nature of it.

¹⁶I am very grateful to Patrick Allo for having called my attention to this point.

The distinction may be better appreciated if we look at a second objection against the inclusion of A_7 , which actually turns in its favour. It concerns the provability of $\diamond\Box\varphi \rightarrow \varphi$ in KTB. Ontologically, this is known to be a rather controversial result. Yet, informationally, $UI\varphi \rightarrow \varphi$ has a very intuitive reading. We already know from A_9 that a is an informationally consistent agent and, from A_4 , that a is informed that p only if p , so we only need now an axiom of constructability of a 's information base: if, for all a 's information it is possible that a holds the information that p (if, according to a 's information base D_a , D_a can be consistently extended to include the information that p) then p must be the case. In other words, the negation of $UI\varphi \rightarrow \varphi$ would make no sense: if φ is false, then no coherent incrementation of the information database is possible by uploading the information that φ . This shows, quite interestingly, that the connection between the intuitionistically-inspired KTB and IL is not accidental. What lies behind both is a concern for direct methods to expand the information base.

It might seem that, by satisfying A_7 , IL embeds a closed-world assumption¹⁷. The similarity is indeed there, but there is also a fundamental difference. In any interesting formalisation of “being informed”, it is plausible to assume that the agent has only incomplete information about the world. This precludes, as inappropriate, the assumption that, if a is not informed that φ then φ is false¹⁸. What A_7 guarantees is that any possible extension of a 's information base corresponds to a genuine state of the world. Since the dual (A_7d) $\varphi \rightarrow \diamond\Box\varphi$ can replace A_7 as the characterizing axiom schema of any KTB-based system, in the next section we shall adopt it as a more intuitive alternative.

2.6. KTB-IL

We have now completed the analysis of all the axiom schemata. The result is a KTB-based information logic (KTB- IL). Compared to EL and DL , KTB- IL satisfies the following minimal set of axiom schemata and inference rules (modus ponens and necessitation):

- A_1 $\varphi \rightarrow (\chi \rightarrow \varphi)$
- A_2 $(\varphi \rightarrow (\chi \rightarrow \psi)) \rightarrow ((\varphi \rightarrow \chi) \rightarrow (\varphi \rightarrow \psi))$
- A_3 $(\neg\varphi \rightarrow \neg\chi) \rightarrow (\chi \rightarrow \varphi)$

¹⁷I am grateful to Daniel Lemire for having called my attention to this point. I agree with Patrick Allo that an elegant way of reading Lemire's suggestion is by explaining the weakening of the closed-world assumption by saying that being informed is ‘prospectively or purposefully consistent / true’, and hence ‘closed for the limiting case’.

¹⁸For a qualified assumption, in terms of local closed-world, see Golden et al. (1994).

| | |
|----------|---|
| A_4 | $I\varphi \rightarrow \varphi$ |
| A_5 | $I(\varphi \rightarrow \chi) \rightarrow (I\varphi \rightarrow I\chi)$ |
| A_{7d} | $UI\varphi \rightarrow \varphi$ |
| MP | $\vdash \varphi, \vdash \varphi \rightarrow \chi \Rightarrow \vdash \chi$ |
| Nec | $\vdash \varphi \Rightarrow \vdash I\varphi$ |

Two birds with the same stone, as the saying goes: we have a *NML*-based logic for “being informed” and a cognitive reading of *KTB*.

3. Four epistemological implications of *KTB-IL*

The debate on information overload, the veridical nature of information, the unsatisfactory state of the $Kp \rightarrow Bp$ principle, and more generally the “Gettierisable” nature of the tripartite definition of knowledge as justified true belief, are what motivated the search for a logic of information. In this second part of the article, we shall be concerned with these issues.

3.1. Information Overload in *KTB-IL*

KTB-IL is not immune from the classic difficulty of information overload, generated by the inevitable inclusion of the rule of necessitation together with *IL*’s closure under implication through the axiom schema A_5 ($I_a(p \rightarrow q) \rightarrow (I_ap \rightarrow I_aq)$). The informational agent a is informed about all theorems provable in *PC* as well as in *KTB-IL*. This is a lot of information, perhaps too much to be realistically attributed to a .

The difficulty has long been recognised in *EL* as a problematic consequence (Hintikka, 1962), to the point of being sometimes deployed as a *reductio ad absurdum*.

A first reply, of course, is to bite the bullet and argue that, in *IL*, the rule of necessitation describes only an ideal agent (Lemmon, 1959), one who is strongly logical omniscient, to adopt Girle’s appropriate classification (Girle, 2000). One may then stress that cognitive overload — whether informational, epistemic or doxastic — is a problem common to all cognitive modal logics anyway, not just *KTB-IL*. This is not a solution, of course, but “a problem shared is a problem halved”: *KTB-IL* is not less successful than *DL* or *EL*, and any argument usable to limit the damage of cognitive overload in those logics (again, see Girle (2000) for an overview) can be adapted to try to rescue *KTB-IL* as well. With an extra advantage: the informational agent a could be an ideal artificial agent, a Turing Machine for example, and one may argue that, in this case (but the case is of course generalisable insofar as a Turing Machine is not computationally more powerful than a human

agent provided with the same boundless resources), the rule of necessitation is stating the conversion of φ from being a theorem to being inferable by an agent who, through the relevant axioms, could eventually deduce the information that φ without any external input (a priori, using a classic Kantian vocabulary), at least in principle.

The last suggestion is related to a second, more interesting reply, referring to the non-informative nature of logical truths (Floridi, 2004d, 2005b). In information theory, the “Inverse Relationship Principle” states that the probability P of p — which may range over sentences of a given language (as in Bar-Hillel & Carnap, 1953) or events, situations or possible worlds (as in Dretske, 1981) — is inversely proportionate to the amount of semantic information carried by p . Information goes hand in hand with unpredictability. It follows that when p is a logical truth, we have $P(p) = 1$ and the informativeness of p is 0, that is $\text{Inf}(p) = 0$. Recall now the distinction introduced in § 1 between p being informative and a holding the information that p . If the information that p is “empty”, i.e. entirely uninformative, as it is the case of e.g. a tautology $(q \vee \neg q)$, then a can hold the (empty) information that $(q \vee \neg q)$, but cannot be informed by receiving it, i.e., a 's deficit of information cannot be filled by receiving $(q \vee \neg q)$. 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. The next step consists in realising that inputting a logical theorem $\vdash \varphi$ into a is indistinguishable from assuming that a already holds the information (is already informed) that φ , which is exactly what is stated in $\vdash \varphi \Rightarrow I\varphi$. It turns out that the apparent difficulty of information overload can be defused by interpreting $\vdash \varphi \Rightarrow I\varphi$ as an abbreviation for $\vdash \varphi \Rightarrow P(\varphi) = 1 \Rightarrow \text{Inf}(\varphi) = 0 \Rightarrow \vdash I\varphi$, which does not mean that a is actually informed about all theorems provable in PC as well as in KTB-IL — 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 . Recall that a might be a Turing Machine, and note the difference: we are not saying that a cannot hold the information that φ .

One may object that we have assumed the availability of boundless resources. The reply is that this is a useful abstraction and the approach is neatly consistent with the “implicit knowledge” strategy developed to solve the logical omniscience problem when this affects resource-bounded agents (Levesque, 1984; Fagin & Halpern, 1988).

3.2. *In Favour of the Veridicality Thesis*

One of the counterintuitive consequences of the “Inverse Relationship Principle” is that the less probable p is the more informative it becomes, with the result that the most informative p is a contradiction, since $P(\text{contradiction}) = 0$. In Floridi (2004d) I defined this as the Bar-Hillel-Carnap paradox (Bar-Hillel & Carnap, 1953). I then argued that the paradox may be solved by assuming that factual semantic information encapsulates truth: Matthew is informed that milk contains calcium if and only if Matthew holds that milk contains calcium and it is true that it does. Were milk not to contain calcium we would deem Matthew disinformed or uninformed. The details of the solution are not relevant here. What matters is that this approach to semantic information has been criticised for being too strong (see for example Fetzer, 2004). I have answered such criticism elsewhere (Floridi, 2005a), but it must be acknowledged that any strongly semantic theory of information (i.e. one that defines information as necessarily veridical) faces a difficulty, namely the lack of a logic that may allow truth-encapsulation without facing epistemic collapse (i.e. the transformation into an epistemic logic). We have seen that this is the difficulty solved by the availability of *KTB-IL*, which shows that a modal logic that captures the relation of “being informed” by interpreting it on the basis of a strongly semantic interpretation of information is possible.

3.3. *The Relations Between DL, IL and EL*

As Lemmon (1959) rightly remarked, “With different interpretations in mind, and with generically different justifications, one may accept as in some way correct any of the formal systems [...] *M*, *S4* and *S5*. Once the complexity of the notion of correctness here is made clear, there is little temptation to view these (and other) modal systems as if they were rival competitors in the same field, of which only one can win. The very multiplicity of modal systems is precisely an advantage, because it gives opportunities for choice.” (p. 40.) *Mutatis mutandis*, a similar temptation should be resisted in any “cognitive” interpretation of *NML*. Let us briefly look at the variety of alternatives.

The exclusion of A_4 from *KTB-IL* yields a *KDB*- or *KB*-based logic, which may be confused with some kind of *DL* (see Figure 3). Yet both systems still include A_{7d} , which makes a doxastic interpretation unfeasible. *KTB-IL* is not based on a more basic, doxastic logic, not even when *DL* is constructed using the “logic of strong belief” as in Lenzen (2002). For in this case, Cp formalises “ a is firmly convinced that p ”, but axiom A_4 still fails to apply, so UIp cannot be interpreted as being equivalent to Cp . We shall see the importance of this conclusion in the next section.

On the other hand, the exclusion of A_{7d} from $KT\text{-}IL$ yields a KT -based IL , which is modally equivalent to, and hence subjectively indistinguishable for a from, the corresponding EL : KT may be equally used to formalise a weak IL or a weak EL , with at least three significant consequences.

First, $KT\text{-}IL$ may be generated by adding A_4 to a KD -based DL . This is interesting because it allows a different interpretation of DL as a logic of (well-formed and meaningful) data holding, free from any mental component. Moving from K to KD to KT , one may read each system as formalising increasingly stringent logics of “ a holds that p ”, where p is some well-formed and meaningful data, i.e. some semantic content, expressed propositionally.

Second, $KT\text{-}IL$ can then be used to generate an $S4$ -based (by adding A_6) or an $S5$ -based EL (through $S4$ or by adding A_8), which is also obtainable from $KT\text{-}IL$, through A_6 or A_8 . All this goes some way towards explaining why conceptual analyses of knowledge, belief and information may move rather freely, and hence sometime confusingly, between DL , IL and EL .

Third, the partial overlapping between IL and EL in KT points out that there is something missing in EL itself. In the epistemological context, the relation of “knowing” is normally expected to include more than just true (doxastic or informational) content. Normally, “being justified” (or some similar relation of well-foundedness, e.g. Plantinga’s “warranty” or Nozick’s “truth-tracking”) plays a significant role. Yet, in EL there is no reference to any further condition. This reinforces the point made in § 3.2: there is room for IL between DL and EL at the very least because EL is just reflective (introspective, in more psychologistic vocabulary) IL . The present state of EL may therefore finally look unsatisfactory, insofar as a crucial feature of the “knowing” relation escapes the formalizations offered by the various versions of EL . EL needs to be augmented by a logic of a relation of well-foundedness. Luckily, work in this direction has been in progress for some time¹⁹.

¹⁹ van Benthem (1991) has called attention to the importance of developing an epistemic logic reinforced by a logic of justification. Research in this direction includes Voorbraak (1991, 1992) and more recently several papers by Artemov and colleagues, see especially Artemov & Nogina (2005), which provides an overview as well.

| | A_4 $\Box_x \varphi \rightarrow \varphi$ | A_6 $\Box_x \varphi \rightarrow \Box_x \Box_x \varphi$ | A_7 $\varphi \rightarrow \Box_x \Diamond_x \varphi$ $\Diamond_x \Box_x \varphi \rightarrow \varphi$ | A_8 $\Diamond_x \varphi \rightarrow \Box_x \Diamond_x \varphi$ | A_{11} $\Box_x \Box_y \varphi \rightarrow \Box_x \varphi$ |
|--------------|---|---|---|---|--|
| Frame | Reflexive | Transitive | Symmetric | Euclidean | (Multi-Agent) |
| S5-based EL | ✓ | ✓ | ✓ | ✓ | ✓ |
| S4-based EL | ✓ | ✓ | | | ✓ |
| KT-based EL | ✓ | | | | ✓ |
| KD5-based DL | | ✓ | ✓ | ✓ | |
| KD4-based DL | | | ✓ | ✓ | |
| KTB-based IL | ✓ | | ✓ | | ✓ |
| KDB-based IL | | | ✓ | | ✓ |
| KT-based IL | ✓ | | | | ✓ |

Figure 3: Summary of the main “cognitive” modal logics discussed in § 3.3

3.4. *Against the Untouchable*

The commonly held principle²⁰ that knowledge necessarily implies belief — known as the entailment property (Lenzen, 1978) or the $Kp \rightarrow Bp$ principle (Girle, 2000) — although reasonable, can now be shown not to be the only alternative. This is because $\neg Bp \rightarrow \neg Kp$ is no longer necessarily true, since we might have $Kp \rightarrow Ip$ and it is not true that $Bp \rightarrow Ip$ or that the relation of “being informed that p ” is necessarily reducible to a weaker relation of “believing that p ” (see § 2.6). This means that a may know that p only if a is informed that p and a being informed that p may not be based on any doxastic state or process. *EL* may be based on *IL* bypassing any *DL* entirely, because “believing” may bear no relation to coherence or truth, but does require a mental attitude of assent and may encompass the *BB* thesis, whereas “being informed” may not be a mental condition and hence avoids the *II* thesis, but cannot be decoupled from the veridical nature of the content, through A_{7d} or A_4 .

Rejecting the $Kp \rightarrow Bp$ principle requires a new approach in epistemology, yet the reader should not be scandalised. “Doxasticism” in epistemology is a recent phenomenon, despite some pervasive propaganda. The Greeks, and especially Plato, could not have mistrusted “doxa” more; modern philosophers were equally concerned with epistemic processes involving ideas (Descartes, Locke, Hume) or judgements (Kant) not necessarily beliefs; and many philosophers of science have always been suspicious of

²⁰ In the literature on epistemic logic, the principle is usually introduced as uncontroversial and unproblematic, see for example Schlesinger (1985), Girle (2000) or Lenzen (2002). The same holds true in analytic epistemology, where it is often attributed to Plato, mistakenly.

“doxasticism”, considering it far too close to forms of armchair psychologism to provide a reliable starting point (one may recall Popper’s “epistemology without a knowing subject”). As Plato forcefully argues in the *Theaetetus*:

“[208b] Socrates: So my friend there is such a thing as right belief together with justification, which is not entitled to be called knowledge.

Theaetetus: I am afraid so.

[210a] Socrates: [...] So, Theaetetus, neither perception, nor true belief, nor the addition of a ‘justification’ to true belief can be knowledge.”

It was an uncritical revival, between the wars, of a psychologistic reading of the Cartesian tradition (Floridi, 2003) that brought the $Kp \rightarrow Bp$ principle to the forefront and silently transformed it into a dogma. The time has come to approach it with a pinch of critical attitude. The invitation is not entirely new. Recent research (Voorbraak, 1991; Halpern, 1996) has raised substantial doubts on the indiscriminate acceptability of the principle, although for reasons different from those expounded here. Voorbraak (1991) has proved, for example, that objective knowledge as formalised in *S5-EL* does not imply rational belief as formalised in *KD45-DL*. A critical, if not suspicious, attitude towards “doxasticism” is a healthy outcome of epistemological investigations used to deal with artificial epistemic agents.

Dethroning the $Kp \rightarrow Bp$ principle from its safe position as a de facto axiom²¹ has a crucial consequence: it opens up the possibility of a non-doxastic but informational approach to the definition and conceptual understanding of knowledge. This is important. Since the Gettier problem is demonstrably unsolvable (Floridi, 2004b), it follows not only that the tripartite account is logically inadequate as it is, but also that it is irretrievably so in principle. The Gettier problem is not a mere anomaly, requiring the rectification of an otherwise stable and acceptable account of propositional knowledge. It is proof that something in the core of the approach needs to be abandoned. Now, of the conditions required by the tripartite definition of knowledge, once we exclude the possibility of fiddling with the truth requirement, it has always been the justification relation that has come under investigation, to be revised or augmented by a fourth condition, depending on the verdict. It can be proved, however (Floridi, 2004b), that the relation of justification is

²¹ Kraus & Lehmann (1986) and van der Hoek (1991) have developed epistemic systems that include $Kp \rightarrow Bp$ among the axioms.

not guilty, i.e. that nothing one can do about it can actually change the outcome: the Gettier problem remains unsolvable. Where else could we look then? The culprit might have been in front of our eyes, unsuspected, all along: it may be the doxastic condition, the conviction that if a knows that p then, necessarily, a must believe that p . This seems to be far from obvious now. We have been blinded by the uncritical assumption of the $Kp \rightarrow Bp$ principle as a dogma: a may know that p because a may be informed that p (plus other conditions of well-foundedness) and “being informed” requires a different, non-doxastic analysis.

Conclusion

The results just seen pave the way to a better understanding of the relations between “knowing”, believing” and “being informed”, to a non-doxastic foundation of knowledge, and to the possibility of a non-psychologistic, non-mentalistic and non-anthropomorphic approach to epistemology, which can easily be applied to artificial or synthetic agents such as computers, robots, webbots, companies, and organizations. There is, admittedly, quite a lot of work to be done. For example, if an informational analysis of knowledge is possible then the strategy to defuse the problem of information overload proposed in § 3.1 could be extended to try to solve the problem of strongly logical omniscience in *EL* as well. More generally, the agenda includes the development of, on the one hand, a clear analysis of the connections between *KTB-IL* and the logics of “becoming informed” and of “being informative” and, on the other, of an informational (as opposed to doxastic) approach to the definition of knowledge. Luckily, however, these are topics that can be left to future research. For as far as the task of this article is concerned, we have come to an end of our toil.

ACKNOWLEDGEMENTS

I discussed several drafts of this paper at many meetings: at the 10th Workshop of Topika, Mesta, Chios (September, 2002); at the Department of Philosophy, University of Hertfordshire (January, 2005); at the Department of Computer Science, UCL (February 2005); at the Dipartimento di Scienze della Formazione, Università degli Studi di Salerno (Marzo, 2005); at the Departamento de Filosofia and the Faculdade de Letras, Universidade do Porto (June 2005); the Unité d’enseignement et de recherche, Sciences humaines, Lettres et Communication, Télé-Université (Teluq) Université du Québec (May 2005); Regensburg Universität (November 2005). I wish to thank Gerassimos Kouzelis; Brendan Larvor; Mel Slater; Roberto Cordeschi; Armando Malheiro da Silva and Fernanda Ribeiro; Jean Robillard and

Rainer Hammwöhner for their kind invitations, which gave me the opportunity to discuss and clarify several ideas contained in this paper. I am also grateful to the participants in these meetings for their helpful discussions. In particular, I would like to acknowledge the help, useful comments and criticisms by Patrick Allo, Fred Dretske, Rod Girle, John Halleck, Daniel Lemire, Paul Oldfield, Gianluca Paronitti, Claudio Pizzi, Hans Rott, Jeff Sanders, Sebastian Sequoiah-Grayson, Timothy Williamson. As usual, they are responsible only for the improvements and not for any remaining mistakes.

Dipartimento di Scienze Filosofiche
 Università degli Studi di Bari
 Faculty of Philosophy and IEG
 Computing Laboratory
 Oxford University

E-mail: luciano.floridi@philosophy.oxford.ac.uk

REFERENCES

- Allo, P., 2005, “Being Informative – Information as Information Handling”, in: *WM2005: Professional Knowledge Management Experiences and Visions*, K.-D. Althoff, A. Dengel, R. Bergmann, M. Nick and T. Roth-Berghofer, ed., DFKI GmbH, Kaiserslautern: 579–586.
- Artemov, S., and Nogina, E. 2005, “On Epistemic Logic with Justification”, in: *Proceedings of the 10th conference on Theoretical aspects of rationality and knowledge*, Singapore, edited by Ron van der Meyden (National University of Singapore), 279–294.
- Baltag, A. and L. S. Moss, 2004, Logics for Epistemic Programs, *Synthese*, 139(2): 165–224.
- Bar-Hillel, Y., and Carnap, R. 1953, “An Outline of a Theory of Semantic Information” repr. in Bar-Hillel (1964), pp. 221–74.
- Bar-Hillel, Y., 1964, *Language and Information. Selected Essays on Their Theory and Application*, Addison-Wesley, Reading / London.
- Barklund, J. 1995, “Metaprogramming in Logic”, in: *Encyclopedia of Computer Science and Technology*, edited by A. Kent and J. G. Williams (New York: Marcel Dekker), vol. 33, 205–227.
- Barwise, J., and Seligman, J. 1997, *Information Flow: The Logic of Distributed Systems* (Cambridge: Cambridge University Press).
- Becker, O. 1930, “Zur Logik Der Modalitäten”, *Jahrbuch für Philosophie und phänomenologische Forschung*, 11, 497–548.

- Brazier, F. M. T., and Treur, J. 1999, “Compositional Modelling of Reflective Agents”, *International Journal of Human-Computer Studies*, 50, 407–431.
- Colburn, T. R. 2000, *Philosophy and Computer Science* (Armonk, N.Y.: M.E. Sharpe).
- Conway, J. H., and Guy, R. K. 1996, *The Book of Numbers* (New York: Copernicus).
- Costantini, S. 2002, “Meta-Reasoning: A Survey”, in: *Computational Logic: Logic Programming and Beyond – Essays in Honour of Robert A. Kowalski*, edited by A. C. Kakas and F. Sadri (Springer-Verlag).
- Devlin, K. J. 1991, *Logic and Information*, (Cambridge: Cambridge University Press).
- Dodig-Crnkovic, G. 2005, “System Modeling and Information Semantics”, in: *Proceedings of the Fifth Promote IT Conference*, Borlänge, Sweden, edited by Janis Bubenko, Owen Eriksson, Hans Fernlund, and Mikael Lind (Studentlitteratur: Lund).
- Dretske, F. I. 1981, *Knowledge and the Flow of Information* (Oxford: Blackwell). Reprinted in 1999 (Stanford, CA: CSLI Publications).
- Fagin, R., and Halpern, J. Y. 1988, “Belief, Awareness and Limited Reasoning”, *Artificial Intelligence*, 34, 39–76.
- Fetzer, J. H. 2004, “Information, Misinformation, and Disinformation”, *Minds and Machines*.
- Floridi, L. 2003, “The Renaissance of Epistemology: 1914–1945”, in: *Cambridge History of Philosophy 1870–1945*, edited by Thomas Baldwin (Cambridge: Cambridge University Press).
- Floridi, L. 2004a, “Information”, in: *The Blackwell Guide to the Philosophy of Computing and Information*, edited by L. Floridi (Oxford – New York: Blackwell), 40–61.
- Floridi, L. 2004b, “On the Logical Unsolvability of the Gettier Problem”, *Synthese*, 142(1), 61–79.
- Floridi, L. 2004c, “Open Problems in the Philosophy of Information”, *Metaphilosophy*, 35(4), 554–582.
- Floridi, L. 2004d, “Outline of a Theory of Strongly Semantic Information”, *Minds and Machines*, 14(2), 197–222.
- Floridi, L. 2005a, “Is Information Meaningful Data?”, *Philosophy and Phenomenological Research*, 70(2), 351–70.
- Floridi, L. 2005b, “Semantic Conceptions of Information”, *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, <http://plato.stanford.edu/entries/information-semantic>.
- Floridi, L., and Sanders, J. W. 2004, “The Method of Abstraction”, in: *Yearbook of the Artificial. Nature, Culture and Technology. Models in Contemporary Sciences*, edited by M. Negrotti (Bern: Peter Lang), 177–220.

- Floridi, L., and Sanders, J. W. forthcoming, "Levellism and the Method of Abstraction". The final draft of this paper is available as IEG – Research Report 22.11.04, see www.wolfson.ox.ac.uk/~floridi/pdf/latmoa.pdf
- Fox, C. J. 1983, *Information and Misinformation: An Investigation of the Notions of Information, Misinformation, Informing, and Misinforming* (Westport, Conn: Greenwood Press).
- Gärdenfors, P. 1988, *Knowledge in Flux: Modeling the Dynamics of Epistemic States* (Cambridge, Mass; London: MIT).
- Girle, R. 2000, *Modal Logics and Philosophy* (Teddington: Acumen).
- Goldblatt, R. 2003, "Mathematical Modal Logic: A View of Its Evolution", *Journal of Applied Logic*, 1(5–6), 309–392.
- Golden, K., Etzioni, O., and Weld, D. 1994, "Omnipotence without Omniscience: Efficient Sensor Management for Software Agents", in: *Proceedings of the twelfth national conference on Artificial intelligence*, Seattle, Washington, United States, edited by O. Etzioni (AAAI Press), 1048–1054.
- Graham, G. 1999, *The Internet: A Philosophical Inquiry* (London: Routledge).
- Grant, J., and Hunter, A. forthcoming, "Measuring Inconsistency in Knowledgebases", *Journal of Intelligent Information Systems*.
- Grice, H. P. 1989, *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press).
- Groenendijk, J. A. G., Janssen, T. M. V., and Stokhof, M. J. B. (ed.) 1984, *Truth, Interpretation, and Information: Selected Papers from the Third Amsterdam Colloquium*, (Dordrecht, Holland; Cinnaminson, U.S.A: Foris Publications).
- Halpern, J. Y. 1996, "Should Knowledge Entail Belief?" *Journal of Philosophical Logic*, 25(5), 483–494.
- Hintikka, J. 1962, *Knowledge and Belief: An Introduction to the Logic of the Two Notions* (Ithaca: Cornell University Press).
- Hocutt, M. 1972, "Is Epistemic Logic Possible?" *Notre Dame Journal of Formal Logic*, 13(4), 433–453.
- Hughes, G. E., and Cresswell, M. J. 1968, *A New Introduction to Modal Logic* (London: Routledge).
- Kraus, S., and Lehmann, D. 1986, "Knowledge, Belief and Time", in: *Proceedings of the 13th ICALP*, edited by C. Krott (Springer), 186–195.
- Lemmon, E. J. 1959, "Is There Only One Correct System of Modal Logic?", *Proceedings of the Aristotelian Society*, 23, 23–40.
- Lenzen, W. 1978, *Recent Work in Epistemic Logic* (Amsterdam: North-Holland).

- Lenzen, W. 2002, “Epistemic Logic” in: *Handbook of Epistemology*, edited by I. Niiniluoto, M. Sintonen, and J. Wolenski (Dordrecht: Kluwer), chap. 25.
- Levesque, H. 1984, “A Logic of Implicit and Explicit Belief”, in: *Proceedings AAAI-84*, Austin, TX, 198–202.
- McCarthy, J. 1971–1987, “Formalization of Two Puzzles Involving Knowledge”.
- McCarthy, J. 1990, *Formalizing Common Sense: Papers by John McCarthy* (Norwood, NJ: Ablex).
- Miyazaki, Y. 2005, “Normal Modal Logics Containing KTB with Some Finiteness Conditions”, in: *Advances in Modal Logic, Vol. 5*, edited by R. Schmidt, I. Pratt-Hartmann, M. Reynolds, and H. Wansing (London: College Publications), 171–190.
- Rao, A., and Georgeff, M. 1991, “Modeling Rational Agents within a BDI-Architecture” in: *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, edited by J. Allen, R. Fikes, and E. Sandewall (San Mateo, CA: Morgan Kaufmann), 473–484.
- Sanders, J. W. forthcoming, “On Information”.
- Schlesinger, G. N. 1985, *The Range of Epistemic Logic* (Aberdeen and Atlantic Highlands, N.J: Aberdeen University Press – Humanities Press).
- Shannon, C. E., and Weaver, W. 1949 rep. 1998, *The Mathematical Theory of Communication* (Urbana: University of Illinois Press). Foreword by Richard E. Blahut and Bruce Hajek.
- van Benthem, J. 1991, “Reflections on Epistemic Logic”, *Logique & Analyse*, 34(133–134), 5–14.
- van der Hoek, W. 1991, “Systems for Knowledge and Beliefs” in: *Logics in AI*, edited by J. van Eijck (Berlin: Springer), 267–281.
- Voorbraak, F. 1991, “The Logic of Objective Knowledge and Rational Belief”, in: *Logics in AI, European Workshop JELIA '90*, edited by Jan van Eijck (Springer), 499–515.
- Voorbraak, F. 1992, “Generalised Kripke Models for Epistemic Logic”, in: *Theoretical Aspects of Reasoning about Knowledge – Proceedings of the Fourth Conference*, edited by Y. O. Moses (Morgan Kaufmann), 214–228.
- Williamson, T. 1999, “Rational Failures of the KK Principle” in: *The Logic of Strategy*, edited by C. Bicchieri, R. Jeffrey, and B. Skyrms (Oxford: Oxford University Press), 101–118.
- Williamson, T. 2000, *Knowledge and Its Limits* (Oxford: Oxford University Press).