

The Logic of ‘Being Informed’ Revisited and Revised

Abstract

The logic of ‘being informed’ gives a formal analysis of a cognitive state that does not coincide with either belief, or knowledge. To Floridi, who first proposed the formal analysis, the latter is supported by the fact that unlike knowledge or belief, being informed is a factive, but not a reflective state. This paper takes a closer look at the formal analysis itself, provides a pure and an applied semantics for the logic of being informed, and tries to find out to what extent the formal analysis can contribute to an information-based epistemology.

1 INTRODUCTION

In Floridi [2006] it is argued that the statal relation of an agent a holding the information that p could correctly be formalised as $l_a p$, where l_a is a **KTB** box-operator. This means that being informed is, in the first place, characterised as a state which is veridical (one is informed that p only if p is true; in short: $l_a p \rightarrow p$), but not reflective or introspective (being informed does not imply being informed that one is informed). In addition, the interpretation is claimed to satisfy a sort of constructibility principle which is associated with the Brouwerian axiom $p \rightarrow l_a \neg l_a \neg p$. Given this analysis, it is concluded that the logic of l_a does not coincide with any of the logics that are traditionally thought of as logics of knowledge or belief. The main value of this insight is, according to Floridi, that it shows that being informed is neither knowing (yet) nor (just) believing. Following the rejection of any tri-partite analysis of knowledge in Floridi [2004b], this insight is a first step towards a new information-based and non-doxastic analysis of knowledge.

The case for a **KTB**-analysis of being informed is strictly speaking independent from the broader suggestion that being informed is equivalent to neither knowing nor believing. In other words, one can assent to the latter while rejecting the former. The other direction is, at least in its general form, harder to resist. The view that knowledge, belief, and being informed are three different cognitive states is Floridi’s main target. Call this the *basic independence thesis*. Unless common logical analyses of knowledge and belief are fundamentally mistaken, this thesis is implied by the **KTB**-analysis. Still, there are two other independence theses (or conjectures) that are as relevant as the basic one that is implicitly endorsed by Floridi.

The first one aims at the difference between first-order knowledge and first-order information; that is, the difference in extension between knowledge and information of non-epistemic facts. The second one aims (with roughly the same first-order focus) at the difference between information and true belief. If we're interested in how the **KT**B-analysis relates to these independence theses, it is the second one which poses the most serious challenge.¹ Call this one the *strong independence thesis*.

On a more general level, we could say that the philosophical defence of an information-based reconstruction of epistemology should at least include (a) the basic independence thesis, (b) the strong independence thesis, and (c) the more general non-mentalistic motivations for such an epistemology. Ideally, each of these should at least get some support from the **KT**B-analysis, but (a*) while the original version supports the basic independence thesis, (b*) it does not at all support the strong independence thesis, and (c*) its implementation of being informed as a non-mentalistic state is questionable as well. The main reason why the original **KT**B-analysis does a poor job is its purely syntactical formulation. In section 2, this diagnosis of the mismatch between the **KT**B-analysis and its philosophical ambition is analysed in a two-pronged fashion. First, it is concluded that a mere syntactic formulation of the **KT**B-analysis provides an insufficient basis for the more interesting independence theses; second, it is argued that a semantic presentation of the **KT**B-analysis requires a non-standard interpretation of informational accessibility. Fortunately, neither of these constitutes an insurmountable obstacle.

The move to a model-theoretic presentation of the logic of being informed is only a first step towards the further integration of an information-based epistemology and the **KT**B-analysis. That is, given the model-theory or pure semantics for a slightly weakened—non-normal—version of the logic of being informed that I propose in section 3, it remains to be shown that such a model theory admits an interpretation that is itself in line with the notion it intends to model. A first stab at an applied or so-called de-praved semantics is based on the *de dicto* reading of data and a notion of information based on it. The resulting sentential approach is based on the suggestion that information can only be derived from one's truthful data. This proposal is described, and ultimately rejected in section 4. A more refined interpretation of being informed that is based on a *de re* reading of data is introduced in section 5. This final proposal succeeds where the others fail: it shows that the **KT**B-analysis can be brought in line with an information based epistemology by (a) exploiting the open-endedness of *de re* data to account for the opacity of being informed and (b) imposing a weak modal constraint on being informed (the impossibility to rule out

¹One reason for this is that one can still maintain that knowledge implies belief whereas being informed does not. This move is, at least in principle, available to anyone interested in an information-based epistemology because even if the tri-partite analysis of knowledge is beyond repair, knowledge can still as a matter of fact imply belief.

the actual world) over and above the non-modal condition that information cannot be derived from false data. While this is sufficient to defend the **KT**B-analysis, it also reveals that the logic of being informed and the philosophical program it is meant to support do not always interact as envisaged in Floridi [2006].

2 BEING INFORMED AND THE **KT**B-ANALYSIS

The original case for the **KT**B-analysis, as it is given in Floridi [2006], is exclusively based on the examination of a number of characteristic axioms. As a consequence, all the argument shows is that the logic for being informed is at least as strong as **KT**B, but strictly weaker than **S5**. This is important to keep in mind, as it reveals that while an appropriate model-theory for the logic of being informed should validate all **KT**B-theorems and provide counter-examples for **S5**-theorems that are **KT**B-*invalid*, there is no need for **KT**B to be complete with respect to that model-theory.

These requirements are still too strong. A closer examination of the reasons why being informed should be formalised by means of a normal modal logic reveals that only the weak necessitation-rule (all theorems of propositional logic are necessarily true) is in fact required. As a result, an appropriate model-theory for being informed should still warrant a notion of being informed for so-called *Cartesian Agents*, but there is no need for an informational counterpart of the principle of strong logical omniscience (Girle [2000: 10.4]). The admissibility of a particular kind of non-normal modal logics is not just a frivolous option; a natural way of modelling being informed (section 3) gives us just that. This is why the logic for being informed not only needs to be revisited (in virtue of the lack of a proper model-theory), but is also in need of revision.



One might point out that, independently of how we model higher-order information (iterated I_a 's), there is a more obvious sense in which the logic of being informed would differ from more common logics for belief (and, *a fortiori*, for true belief). After all, it is well known that if, as Dretske [1999] has it, information requires the satisfaction of some dependency conditions, it cannot be closed under known (alternatively, informed) implication.² If, unlike being informed, true belief does satisfy this form of closure, the stronger independence thesis we're after immediately follows.

Still, it becomes quite clear in Floridi [forthcoming] that—apart from the veridicality thesis—the logic of being informed was never meant to capture this more demanding notion. Indeed, being informed that p as holding the

²See Adams [forthcoming] where the logic of being informed is contrasted with the Dretskeian program, and Arlo Costa & Parikh [2006] where a neighbourhood semantics is used to model this kind of knowledge. Note also that the rejection of closure under known implication goes back to Dretske [1970] where the notion of information isn't even mentioned.

information that p is just holding a piece of information; where the latter is to be understood as a truthful and meaningful well-formed datum. But what is then the difference between being informed that p and having a true belief that p ? The main reason for keeping these distinct is that the latter includes a mental state whereas the former doesn't.³ As emphasised by Floridi (pc), this is a good reason to refrain from modelling one after the other. While the philosophical motivation for this difference is easily grasped, it is not as easy to see how this difference should be reflected within a formal system. In addition, this does not have to mean that the difference can be reduced to telling mental and non-mental states apart. There could also be a difference in ways of being truthful.

According to Floridi [2005a], information is truth-constituted. Unlike mere propositional contents, on which the truth-value supervenes, pieces of information encapsulate their veridicality. One challenge for the logic of being informed is to explain how this relates to the **KT**B-analysis (or some close relative of this analysis). The relevance of this idea of truth-encapsulation draws on the fact that being informed is meant to be a (the) stepping stone to knowledge that, in view of the Gettier-counterexamples, mere (true) belief cannot be (Floridi [2004b]). In spite of its broader importance, the original formulation of the **KT**B-analysis does not sufficiently clarify how being informed can play this role. The lesson we learn from the **KT**B-analysis is that information is truthful whereas mere beliefs need not be, and that beliefs are introspective whereas information is not. Neither of these, however, leads to a positive characterisation of information vis-à-vis true belief. This much needed positive characterisation largely coincides with the stronger independence thesis mentioned in the introduction.⁴



One problem with the **KT**B-analysis of being informed is that superficial matches between the philosophical and the formal side of the story can be misleading. Here, I want to look at a first sense in which the original **KT**B-analysis can only scratch at the surface of what it means to be informed, and reveal how this affects the prospects of that analysis. Put simply, the problem is that the original analysis cannot see beyond the syntactical surface of the axiomatic presentation. What this hints at is that Floridi came to the **KT**B-analysis through the examination of modal axioms. He proceeded by elimination, and settled on a particular set of axioms in virtue of the mutual agreement between what one could call the surface behaviour of

³By way of comparison, Dunn thinks of information as “what is left from knowledge when you subtract, justification, truth, belief (...) [and] the thinker” [2008: 581]. Floridi would presumably agree with all but the omission of truth. Still, the contrast is crucial, for on Dunn’s account information is any kind of semantic content (which is fine for a logician), whereas for Floridi it is much less generic (but more useful for doing epistemology).

⁴One should, however, keep in mind that it is not all that obvious that the framework of basic modal logic in which the **KT**B-analysis is formulated has the necessary resources to establish the stronger result.

being informed and an imprecise interpretation of each of these axioms. Proceeding along these lines is like working in the dark. This is typical for the practice of reverse-engineering (in other papers, Floridi uses exactly that term to describe his practice), where one takes an existing approach (modal logics for cognitive states like knowledge and belief) and tries to apply it to a new domain (being informed).

Initially, this kind of reverse-engineering does not seem all that different from the usual practice of formal modelling where one can either opt for a syntactically, or for a semantically driven approach.⁵ There is nothing fundamentally wrong with either of these options, but what I want to point out is that—at least in their standard form—none of these options is particularly well-suited to establish the intended independence theses.

First, consider the syntactical approach. Here, there are two kinds of problems. One is due to the fact that a difference in logical axioms for the notions of being informed, (true) belief and knowledge indicates a difference, but not yet a way to positively characterise being informed versus truly believing. A second one is due to the fact that since we're working in the dark, we do not know where the difference comes from. Here, one could reply that there is a more encompassing motivation behind the **KT**B-analysis which does explain this difference, and adds a positive characterisation. The problem with the syntactical approach,⁶ is that there is no way to find out if the difference that arises from the logical analysis does in fact reflect the kind of difference that is presupposed by this more encompassing story.

This last point can be clarified by appealing to the difference between a pure and an applied semantics (Plantinga [1974], Copeland [1983]), and in particular to the fact that only applied semantics can play a philosophical (as opposed to a merely formal) role. Pure semantics come in many forms, but applied semantics are constrained by the need to give an interpretation of the formal apparatus in terms of non-formal terms that are directly related to the intended use of the logic in question. In this case the intended use as well as the relevant non-formal terms are to a certain extent already in place. The **KT**B-analysis includes arguments for each of the relevant axioms and thereby implicitly fixes the non-formal terms. The main difficulty at this point is that if we don't even have a formal or pure semantics, there is no way to find out if the **KT**B-analysis can be used as a basis for an applied semantics for the logic in question. As a result, we cannot be sure that the agreement between the intended interpretation of being informed and its formal analysis extends below the surface of the axiomatic presentation.

⁵This distinction is relevant for and used by formally oriented logicians (see Blackburn, De Rijke & Venema [2001: 1.7]) as well as more philosophically oriented ones (see Girle [2000: 10.4]).

⁶I assume here that the syntactical approach cannot provide the kind of positive characterisation of information versus true belief, and therefore further ignore that aspect and focus on the negative characterisation induced by the lack of introspection.

The semantical approach has the formal resources to avoid the problems of the syntactical approach, but it also has problems of its own. Consider, for that purpose, the standard epistemic interpretation of Kripke-style semantics. On that account, a knows that p just if p is true in all alternative ‘worlds’ that a cannot distinguish from the actual one. The state of being informed, by contrast, cannot as easily be related to the informed agent’s ability to discriminate alternative states of the world.⁷ Indeed, if being informed that p is just a matter of having true, well-formed and meaningful data for p , there is no reason to assume that such a state has any of the relevant modal properties. This insight poses a serious problem if we want to make sense of a semantic approach to (the **KTB**-analysis of) being informed, but it does not yet imply that such an approach is impossible. What it does show is that, first, we need an account of *informational accessibility* that does not directly refer to an informee’s own ability to tell alternative states of the world apart. Second, once we have such an account, it is not yet clear whether it can do the explanatory work we would like it to do. Namely, to characterise the difference between information, knowledge and belief (as it is presupposed by the **KTB**-analysis) from a model-theoretic perspective. What I mean by this is that if informational accessibility is substantially different from epistemic and doxastic accessibility, it immediately follows that iterated I ’s cannot or should not be compared to the iterated K ’s or B ’s in the respective logical systems, and that therefore the invalidity of $I_a p \rightarrow I_a I_a p$ is not directly related to the non-reflective or non-introspective nature of being informed.



At this point, one might be tempted to base the semantic interpretation of being informed on information-sets, and thus map each axiom to a postulate about what does and does not belong to an agent’s information-set I_a . As such, the validity of $I_a p \rightarrow p$ is understood as $p \in I_a \Rightarrow p$ (p belongs to a ’s information-set only if p is true), and the invalidity of $I_a p \rightarrow I_a I_a p$ as $p \in I_a \not\Rightarrow I_a p \in I_a$. This approach is familiar from the belief-revision literature where the generalised inclusion of *reflective modalities* in belief-sets is shown to clash with the presumed conservativity of revisions (Fuhrmann [1989]). When it comes to the Brouwerian axiom, this kind of approach turns out to be problematic for entirely different reasons. The dilemma it poses is the following. Either $p \rightarrow I_a \neg p$ is understood as $p \Rightarrow \neg I_a \neg p \in I_a$, or as $p \Rightarrow I_a \neg p \notin I_a$. Neither is, however, acceptable. The former is consistent with our interpretation of veridicality and non-reflectivity, but comes with an unwanted introspective flavour. The latter avoids this introspective flavour, but is based on a weaker interpretation of the Brouwerian axiom which cannot be generalised to the interpretation of the other axioms.

⁷For otherwise the “margin for error principle” given in Williamson [2000: Chapt. 5 & Appendix 2] would, given that it warrants a **KTB**-analysis of knowledge, perfectly serve our purpose.

As a consequence, if we assume that there should be a single procedure to transform axioms of the **KTB**-analysis into constraints on information-sets, it turns out that there is no uniform information-set interpretation that vindicates the whole **KTB**-analysis.

Before we move on, let me just stress that the above is not meant to imply that the **KTB**-analysis is itself beyond repair. Rather, the suggestion is that its syntactical formulation does not decisively support either form of the independence thesis, while its semantical formulation cannot be a simple modification of an epistemic alternatives interpretation. The main reason for this is that information is no longer the generic notion most logicians like to think about, and which is based on the identification of the information an agent holds with its semantic content (a proposition or intention), but something more specific. Unlike other discussions of strongly semantic information, the original **KTB**-analysis does not account for this specificity, and (as already suggested in fn. 4) it is not yet clear whether that specificity is even expressible in the language of basic modal logic. Even if it lacked the relevant resources, the **KTB**-analysis would perhaps not yet be refuted, but it would for sure lose much of its appeal. The inclusion of the Brouwerian axiom would come under pressure, and the logic for being informed (presumably weakened to the logic **T** or even **S0.5**) would be reduced to what knowledge and true belief have in common.

3 A ‘PURE’ SEMANTICS FOR THE **KTB**-ANALYSIS

Using formal methods for philosophical purposes requires us to be attentive to the contrast between so-called ‘pure semantics’ and ‘applied semantics’ (cfr. the remarks in the previous section). Even if it is not our goal to characterise the meaning of ‘being informed’ in natural language, the fact that our semantics is meant to support a philosophical analysis of being informed is itself sufficient to warrant the demand for an applied semantics. With this proviso in mind, I shall still proceed by first formulating a formal model-theory for the logic of being informed. The latter will then be used as a guide to obtain the intended real semantics for being informed.

The language \mathcal{L} for the logic of being informed is inductively defined as:

$$\phi ::= p \mid \perp \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid I_a\phi$$

where p ranges over a set of propositional parameters Prop , and a ranges over a set of agents \mathcal{A} .

A model for being informed is a 4-tuple $\mathfrak{M} = (S, @, \mathcal{I}, \|\cdot\|_{\mathfrak{M}})$ where S is a non-empty set of points, $@$ a designated element of S , \mathcal{I} a set of binary relations I_a on S for each $a \in \mathcal{A}$, and $\|\cdot\|_{\mathfrak{M}}$ a map $\text{Prop} \rightarrow \mathcal{P}(S)$. Finally, we say that a formula of \mathcal{L} is satisfied at a point in a model iff it complies with the usual recursive clauses:

1. $\mathfrak{M}, s \models \phi$ iff $s \in \|\phi\|_{\mathfrak{M}}$, for $\phi \in \text{Prop}$,

2. $\mathfrak{M}, s \Vdash \perp$ never,
3. $\mathfrak{M}, s \Vdash \neg\phi$ iff $\mathfrak{M}, s \nVdash \phi$,
4. $\mathfrak{M}, s \Vdash \phi_1 \vee \phi_2$ iff $\mathfrak{M}, s \Vdash \phi_1$ or $\mathfrak{M}, s \Vdash \phi_2$,
5. $\mathfrak{M}, s \Vdash I_a\phi$ iff $sI_a s'$ implies $\mathfrak{M}, s' \Vdash \phi$.

This means that the specificity of this logic has to surface in the properties of \mathcal{I} as well as in the role we impute on $@$. The minimal properties of informational accessibility are the following:

$$@I_a@ \quad (R)$$

$$@I_a s \rightarrow sI_a@ \quad (S)$$

If, in addition, validity is defined as satisfaction at $@$ for all models, it is a standard exercise to check that the **KTB**-axioms come out valid, while $I_a\phi \rightarrow I_a I_a\phi$ and $\neg I_a\phi \rightarrow I_a \neg I_a\phi$ remain invalid. By contrast, since informational accessibility is only reflexive and symmetric where $@$ is involved, the informational counterpart of the rule of strong necessitation is not sound on this semantics. The rule of weak necessitation is retained in virtue of the fact that validity is only evaluated at $@$, which by definition is a normal world (cfr. the contrast between the non-normal **S**-systems and the even weaker **E**-systems).

From this, it follows that I_a is a box-like modality of a non-normal logic that is strictly stronger than Lemmon's **S0.5** wherein $\Box\phi$ means that ϕ is tautological (Lemmon [1959], Cresswell [1966]). The resulting logic is also the non-normal pendant of the normal modal logic **KTB**. In view of the comments made in section 2, we may thus conclude that the present model-theory agrees with the (appropriately weakened) **KTB**-analysis of being informed (further results as well as an axiomatic presentation of this system are not required for our purposes).

4 A SENTENTIAL INTERPRETATION

Being informed that p not only means having data for p , but requires having veridical data for p (Floridi [2006: 2.2]). This simple remark, together with the assumption that data can be identified with their syntactical representations (see e.g. Floridi [2005b: 1.3]), is the basic insight we need to turn our pure semantics into an applied semantics. As for most other modal logics, we have an applied semantics when the notions of (*informational*) *alternative*, (*informational*) *accessibility*, and *actual world* are not just the set-theoretical constructs used in the previous section, but do correspond to something real. In other words, it does not suffice to say that when $sI_a s'$ we say that for a the world s' is informationally accessible from s , and that therefore s' is an informational alternative to s for a . The proposal set out in this section is based on three distinct considerations.

First, when we say that validity is evaluated at $@$ we mean that validity is only evaluated at the actual world, which in its turn means that the

properties of being informed are identified with its properties at the actual world. This does not say much yet, but for now it suffices to mention that this agrees with the idea that being informed has only few modal properties; being informed is hardly more than having true data. To this we only need to add that what counts as the actual world can be understood quite loosely—all that matters is that it includes the situation or system the relevant agent is said to be informed about.

Second, when we say that s is an informational alternative to $@$ for the agent a , the related notion of informational accessibility has to be understood from a third-person perspective. This point is related to the suggestion in Floridi [forthcoming] that knowledge and information might almost coincide from a first-person perspective, but is best generalised to the idea that being informed does not require the access and/or the transparency that is typically associated with first-person perspectives. The latter can then be translated into the claim that the information held by an agent (who does not necessarily have access to the content of said information) does not have to be reflected in what, from a first-person perspective, that agent considers to be possible given the information he holds.

The third and final consideration is that when information is truthful data,⁸ informational accessibility should be a function of the truthful data an agent holds. This suggestion does not yet pick out a single relation of informational accessibility, for the expression “the truthful data an agent holds” remains ambiguous. To resolve this ambiguity we need to decide whether truthful data that are inferred from false data still counts as information. If the relation of being informed is immune to Gettier-style counterexamples—which it purportedly is—, one would think that information can only be inferred from other information. This “no false premises” commitment (see e.g. Shope [1979])⁹ implies that the sentential approach to being informed needs to distinguish between basic and derivative data. I won't argue in favour of that distinction, but just note that from a formal point of view this leads to something that's closer to belief-bases than to belief-sets whereas from a philosophical point of view this does not have to commit us to a foundationalist rather than to a coherentist notion of justification because, quite simply, that distinction does not apply here.

Using the above, we may say that an agent a is informed that ϕ when the truthful part of his basic data (a 's basic information, for short) implies ϕ (compare with Shin [1993]). Perhaps, due to the requirement that ϕ should be a logical consequence of a 's basic information this is, at least

⁸Strongly semantic information is standardly defined as true, well-formed and meaningful data. Given the assumption that all basic data are sentences of a suitable formal language, all such data are well-formed and meaningful by definition.

⁹We might expect that objections to this kind of restrictions as a means to avoid Gettier-counterexamples (Feldman [1974]) carry over to the present proposal. I shall ignore this issue here, but only remark that the analogy is only partial: Gettier-counterexamples presuppose a fallible account of justification, but being informed is entirely independent from this justificatory aspect.

in general, all too restrictive. Recall, however, that we only want to gain a better understanding of the **KTB**-analysis of being informed, and that this restricted aim does not include solving the sceptical problem. What matters is that the basic information held by a immediately determines the relation of informational accessibility. Where Δ is a 's basic information, each world where all of Δ is true is an informational alternative to the actual world. If $\|\cdot\|_{\mathfrak{M}}$ is used to denote the set of worlds (in a model) where a formula or sets of formulae is true,¹⁰ the sentential interpretation of being informed can be conveniently summarised as follows:

1. Where Γ is the *basic data* held by a ,
2. the basic information held by a is $\Delta = \{\phi \in \Gamma : @ \Vdash \phi\}$,
3. s is informationally accessible from $@$ when $s \in \|\Delta\|_{\mathfrak{M}}$, and
4. a is informed that ϕ when $\|\Delta\|_{\mathfrak{M}} \subseteq \|\phi\|_{\mathfrak{M}}$.

Strictly speaking, the fourth condition above does not coincide with the requirement that ϕ has to be a logical consequence of Δ . To capture the latter, the relevant condition would have to be true in all models. Defining being informed relative to a single model is nevertheless a sensible option; especially if we don't want to consider all worlds that are consistent with an agent's basic information. A single model will do for present purposes, and we thereby weaken the initial proposal by stipulating that a is informed that ϕ if (in the appropriate model) ϕ follows from a 's basic information. That is, if a 's basic information excludes all non ϕ worlds. As we can ignore the sceptical problem, we do not have to specify which models are appropriate.



As an interpretation of the pure semantics of the previous section this approach has two virtues, but also one major drawback. A first virtue is that we have a natural interpretation which links the definition of strongly semantic information to the reflexivity of the relation of informational accessibility. The second virtue is that it supports the strong independence thesis. Finally, the main drawback of a sentential approach is that it is incomplete: the four conditions given above do not settle what counts as an informational alternative to s where $s \neq @$. What is worse is that the most plausible conditions for informational accessibility at non-actual worlds are at odds with the **KTB**-analysis. The reflexivity of informational accessibility is straightforward, but the second virtue as well as the problem of informational accessibility at non-actual worlds deserve a closer look.

Consider the strong independence thesis. How do information and true belief compare? The present model does not as such specify what a believes, but if a 's basic data are understood as the set of a 's basic beliefs it is reasonable to assume that if ϕ follows from these basic beliefs, ϕ is

¹⁰I.e. $\|\cdot\|_{\mathfrak{M}}$ is ‘lifted’ from a function from propositional parameters to sets of worlds to a function from (sets of) formulae to sets of worlds.

at least an implicit belief we may ascribe to a . Finally, for ϕ to be a true belief it only has to be true at the actual world. This agrees with the difference we already aimed at: true beliefs may be grounded in false ones, but information may not be grounded in false data. How does this difference surface within the modal logics that are adequate for each of these notions? A rough answer is that they both satisfy the same T-axiom because informational as well as (lets call it) t-doxastic accessibility are reflexive. A more specific answer is that s is t-doxastically accessible from $@$ whenever it belongs to $\|\Gamma\|_{\mathfrak{M}} \cup @$,¹¹ and that it is informationally accessible when it belongs to $\|\Delta\|_{\mathfrak{M}}$. This settles the strong independence thesis in, admittedly, a rather trivial way. Still, it also makes the more subtle point that at least this difference has no effect on what counts as the right logic for both notions, and thereby enforces our previous remark that the strong independence thesis cannot be settled in a purely axiomatic setting.

Now, consider the issue of what is informationally accessible at non-actual worlds. Expanding on how we settled this question for the actual world, we should expect this to be a function of the basic information held by an agent at the relevant non-actual world. Despite its simplicity, the question of what qualifies as information at a non-actual (but informationally accessible) world is not as easily answered. The straightforward solution is that it should be exclusively determined by which data are true at that non-actual world. An equally defensible alternative is that the agent's information remains fixed; only truth at the actual world matters. Fortunately, we do not need to debate about the respective virtues of each of these proposals, but only need to note that where Δ and Δ' are, respectively, a 's information at the actual world and at the non-actual world, both these proposals still ensure that $\Delta \subseteq \Delta'$ and hence $\|\Delta'\| \subseteq \|\Delta\|$ which is equivalent to the transitivity of the relation of informational accessibility.¹²

The main insight we gain from the above result is that the **KTB**-analysis requires that, for the failure of transitivity of informational accessibility, there be informationally accessible non-actual worlds where some of the basic information a holds at the actual world no longer counts as information. By the same token, the symmetry of informational accessibility requires that one's information at these non-actual worlds would not exclude the actual world; i.e. that at those worlds one could not be informed of something that is actually false (or inconsistent with what is actually true). Using this description, one indeed obtains an interpretation of the pure semantics set out in the previous section.

One way to make this more precise is to specify non-actual information very tightly along the just described lines. As such, a 's information at s can be any subset of his data which (i) are true at s , and (ii) are consistent

¹¹To see that this duly captures the above description of true beliefs, just note that $(\|\Gamma\|_{\mathfrak{M}} \cup @) \subseteq \|\phi\|_{\mathfrak{M}}$ holds iff $\|\Gamma\|_{\mathfrak{M}} \subseteq \|\phi\|_{\mathfrak{M}}$ and $@ \subseteq \|\phi\|_{\mathfrak{M}}$ hold separately, which are equivalent to, respectively, $\Gamma \models_{\mathfrak{M}} \phi$ (Γ strictly implies ϕ in model \mathfrak{M}) and $@ \models \phi$.

¹²Compare with the embedding of intuitionistic logic in the modal logic **S4**.

with what is true at $@$.¹³ This is at best an *ad hoc* solution, as it does not yet explain why some of a 's actual information might, despite being true, fail to qualify as information at s . Another, more radical way out specifies that at non-actual worlds one is not informed at all. Despite its formal elegance,¹⁴ this alternative is equally hard to motivate. Nevertheless, this is the best we can achieve within the present setting, and so it seems that we have once more reached the limits of what can be achieved within a particular framework.



In section 2 I emphasised the problems that arise from a purely syntactical characterisation of the logic for being informed. Here I take a closer look at the limitations of the sentential interpretation of the model-theory described in section 3. One thing that has become clear from the attempts to characterise an agent's basic information as a set of sentences is that once we extend an otherwise natural sentential proposal to informationally accessible worlds, we quickly reach the limits of what can be explained in such a setting. The problem faced by the sentential approach is structurally similar to the previous one: some distinctions we need to make cannot easily be accommodated by our formalism, and when we can accommodate these differences there is no guarantee that we capture the intended difference.

The problem can, at least to a first approximation, be traced back to a lingering commitment to the generic notion of information as mere semantic content (what Floridi [2004a] refers to as *weakly semantic information*). Clearly, if an agent's basic data are identified with a body of semantic content, and if at each world a subset of that content (its truthful part) is used to determine the set of informationally accessible worlds (strictly speaking, the content *is* just the set of informationally accessible worlds) then one cannot be less informed at an informationally accessible world than one is at the actual world. Identifying an agent's basic data with a set of formulae does not alter this, for sentences (and sets of sentences) pick out the same semantic content at any world. Thus, in a sense, the problem we face here is that the sentential interpretation does not allow us to look beyond the informational content we ascribe to an agent when we model that agent's basic data and information.

The moral of this result is that, even if we identify an agent's basic data and information with a set of formulae (a data or information-base), we model it as a set of cognitive commodities which: (a) in the case of data, picks out the same set of accessible worlds at all worlds, and (b) in the case of information, picks out a subset of the actually accessible

¹³The formulation of the second condition does not refer to what is actually true, and thereby also works if s settles an issue that is indeterminate at $@$. In the remainder of this paper I shall nevertheless stick to the initial condition.

¹⁴This is just the well-known idea from non-normal modal logics that at non-normal worlds anything is possible; see Priest [2001: 58-9].

worlds at all informationally accessible worlds. The latter suggests that, as a means to individuate an agent's basic data and information more finely than by means of its semantical content, the sentential approach is not entirely satisfactory. More exactly, the sentential approach works well as long as it is confined to the actual world; where it effectively settles the stronger independence thesis. This is because, as opposed to contents, sets of sentences allow us to discriminate between the consequences of an agent's true data, and the true consequences of of an agent's data. At non-actual worlds, by contrast, this approach enforces too much, and at least the most obvious ways to weaken it seem largely *ad hoc*.

5 TWO LEVELS OF INFORMATION AND THE KTB-ANALYSIS

Characterising an agent's basic data or information as a set of commodities is in no way confined to the sentential interpretation used in the previous section. As explained in Floridi [2005b] and exploited as a means to characterise an agent's information more finely in Allo [2009], data can be understood at different levels. One natural level is given by the *de dicto* reading of data, and sees data as a lack of uniformity at the level of symbolic representation. The sentential approach described in the previous section operates at this level. A less obvious level is given by the *de re* reading, which sees data as a lack of uniformity in the world. This is by far the most encompassing view of data, and gives us more control over how we ascribe data and information to agents.

To exploit this more refined view we need a different procedure to evaluate what counts as an informationally accessible world. We start as we did before by ascribing a set of basic data to each agent, but do not longer require that all these data are effectively in a sentential or otherwise symbolic form. If one thinks that at least an agent's basic beliefs are stored in some or other symbolic form (as presumed by the belief box metaphor), this choice implies that the totality of an agent's basic data does not coincide with its basic beliefs. The ultimate nature of belief is, however, irrelevant to the present discussion. We only need to be attentive to the fact that data and beliefs are not only different things, but that they also do not have to be stored in similar ways. As a working definition, we could stipulate that holding a particular datum implies access to that datum, but does not imply access to the content of that datum.

More important than the nature of *de re* data is the way we can formally exploit its properties to give an interpretation to the pure semantics of section 3. Let Ξ be a 's set of basic data. Now, let Γ be the adequate sentential expression of Ξ at the actual world. This coincides with a 's basic data as it can be represented by a set of formulae. As before, a 's basic information at the actual world can be identified as that part of Γ that is actually true, call this Δ .¹⁵ Using the above characterisation of a 's basic information at

¹⁵Remark that the intermediate step where a sentential reformulation of a 's basic data is

the actual world, the set of informationally accessible worlds can still be identified with $\|\Delta\|$. This account of basic information does not make a significant difference for what is informationally accessible from the actual world. But how does it behave at non-actual informationally accessible worlds? A natural way to understand an informationally accessible world s , is as a world where Δ is true, and a 's basic data still correspond to Ξ . The first question which then arises is how the sentential expression of Ξ at s , call this Γ' , compares to Γ . The second question is which subset of Γ' should qualify as a 's basic information at s . Of course, these questions cannot be answered independently of each other, as our main aim is still to vindicate the pure semantics set out in section 3.



Let us first look at how the *de re* reading of data affects the reflectivity of being informed. The main drawback of the sentential approach is that when $\Gamma = \Gamma'$, we have $\Delta \subseteq \Delta'$, and hence $\|\Delta'\| \subseteq \|\Delta\|$ which is just the transitivity of informational accessibility. As a result, the failure of transitivity is equivalent to $\|\Delta'\| \not\subseteq \|\Delta\|$, and implies $\Delta \not\subseteq \Delta'$. If, at least as a temporary assumption, we let Δ and Δ' be those subsets of Γ and Γ' that are true at, respectively, $@$ and s , then the latter also implies $\Gamma \neq \Gamma'$. As a result, to characterise being informed as a veridical, but non-reflective state, it is necessary that Ξ could have different sentential expressions at $@$ and at s , and sufficient if it could have a logically weaker expression. This summarises the first desideratum for an applied semantics for being informed.

The symmetry of informational accessibility can, by contrast, be shown to coincide with the requirement that $@$ be in $\|\Delta'\|$, and thus that Δ' be that subset of Γ' which is not only true at s , but also at $@$. This requirement can be satisfied in several ways, either by tweaking the relation between Δ' and Γ' (what qualifies as information at s), or between Ξ and Γ' (the sentential expression of Ξ at s), or maybe even by revising the relation between Γ and Δ (what qualifies as information at $@$). Independently of how this is actually achieved, this summarises the second desideratum for an applied semantics for being informed. An applied semantics should, however, not be reduced to the sheer compliance with these two desiderata; it should also be explanatory.



Elsewhere I already argued for the relevance of cognitive commodities (pieces of information), and the need to model these commodities in such a way that they do not uniformly pick out the same propositional contents at all points of a modal space (Allo [2009: sect. 4, 8]). I won't rehearse those reasons here. All that needs to be emphasised is that the sentential model of section 4 isn't adequate when data are merely seen as constraining affordances (Floridi [2005b: 3.2.1]). The general idea is this: Data, especially given is indispensable, for mere *de re* data cannot themselves be qualified as true or false.

those which are only specified as lacks of uniformity in the world, put an agent who has access to these data in a position where she can be (or become) informed of what is actually true. To see data as constraining affordances means that they facilitate knowledge (or at least being informed). The latter implies that seeing data in this open-ended way does not lead a generalised relativism. In the terminology of Floridi [2008] “they limit the possible models” (p. 325); in our terminology this means that they exclude some worlds. Additionally, when we emphasise that they are merely affordances, we mean that data do not have to fully determine what the agent can be informed of. Even if they exclude worlds in a non-arbitrary way, the set of worlds that are actually excluded depend on the context. Of course, this is already true for one of the sentential models laid out in the previous section: when one’s basic information is the truthful part of one’s basic data, one’s basic information cannot but vary across the informationally accessible worlds.

Treating data as constraining affordances presumably requires a further dimension of variation: one that accounts for the relation between basic data and their sentential expression. The only thing we need to note is that there is no need for a full characterisation of the relation between *de re* data and their sentential rendition. As long as it allows that the sentential rendition of a set of basic *de re* data may vary in such a way that, when moving from the actual world to an informationally accessible world, (i) an agent’s basic information may change in such a way that it at least involves an information loss, but (ii) could also not change in such a way that it would lead to the exclusion of the actual world, any such relation will do.



By assembling the different pieces laid out above we can now define the relation of being informed as follows:

1. Given an agent a who holds the basic data Ξ which at the actual world can be expressed as Γ , we say that a is informed that ϕ whenever, in virtue of its truth, Δ qualifies as a ’s basic information at the actual world, and ϕ is true at all informationally accessible worlds (i.e. $\|\Delta\|$).
2. For every informationally accessible world $s \in \|\Delta\|$, the set Γ' is the sentential expression of Ξ at s only if $@ \in \|\Delta'\|$ and Δ' is the truthful part of Γ' at s . Thus, we say that a is informed that ϕ at s whenever ϕ is true at all $s' \in \|\Delta'\|$.

To show that the present proposal fits the bill two things have to be argued for. First, we have to show that the definition is such that the relation of informational accessibility is the kind of reflexive, symmetric, but non-transitive relation that is also used in the pure semantics. Second, we have to show that the present proposal is no longer a pure semantics, but that (in particular) the failure of transitivity and the symmetry of informational accessibility correspond to a real feature of (basic) data and information.

To comply with the pure semantics the above definition has to enforce that $@|_a@$ and $@I_a s \rightarrow sI_a@$ come out valid for all $a \in \mathcal{A}$, and that $(sI_a s' \wedge s'I_a s'') \rightarrow sI_a s''$ has counterexamples for at least some $a \in \mathcal{A}$. The first one is straightforward, as it is implied by the demand that Δ is a 's basic information at the actual world only if $@$ is in $\|\Delta\|$. To prove the second one, assume *for reductio* that for some agent a , the world s is an informational alternative to $@$, but that $@$ is not an informational alternative to s . This means that where Δ is a 's actual basic information, s belongs to $\|\Delta\|$, and a 's basic information at s ; i.e. Δ' is such that $@$ does not belong to $\|\Delta'\|$. The latter implies that the sentential expression of a 's basic data would have been such that its truthful part at s would exclude the actual world, which violates the condition for basic information at a non-actual informationally accessible world.

Finally, to construct a counterexample for $(s_1 I_a s_2 \wedge s_2 I_a s_3) \rightarrow s_1 I_a s_3$, let Γ_i and Δ_i stand for, respectively, a 's basic data and a 's basic information at s_i . It has already been noted that any such counterexample would have to imply $\Delta_1 \not\subseteq \Delta_2$, with Δ_2 such that at least one ϕ implied by Δ_1 is not implied by Δ_2 . Furthermore, since validity has to be evaluated at the actual world, the only relevant counterexamples to the transitivity of informational accessibility are those where the base-world is the actual one (i.e. $s_1 = @$). Thus, we only need to verify that Γ_2 can simultaneously have a truthful part that is logically weaker than Δ_1 and not exclude the actual world. This condition is obviously satisfied since not excluding the actual world only puts an upper bound on Δ_1 , but no lower bound.



None of the above already hints at more concrete interpretations of failures of the transitivity of informational accessibility or the non-exclusion of the actual world. To a first approximation, my suggestion is that the failure of transitivity is due to the potential opacity of *de re* data, whereas the restriction on potential sentential expressions of *de re* data at informationally accessible worlds, which enforces the non-exclusion of the actual world, can be seen as a means to capture the idea that pieces of information encapsulate their truth.

Where Ξ is an agent's basic *de re* data, Γ its sentential expression at the actual world, and Δ its truthful part, we should keep the following issues apart: (a) the fact that Ξ potentially picks out a different set of sentences at each world—this is a key-property of *de re* data as mere constraining affordances; (b) the fact that informational accessibility is merely a function of the truthful part of Γ —this is the key-property of informational accessibility; and (c) the fact that further properties of the relation of being informed can only be determined by putting further constraints on what may be an admissible sentential expression of Ξ at the other informationally accessible worlds. Given that, let me suggest the following general principle which any such constraint should comply with. Where s is an informationally ac-

cessible world, (i) Γ' is an inadmissible sentential expression of Ξ at s only if (through the workings of Δ' ; most likely the truthful part of Γ' at s) it can be dismissed solely in virtue of what is actually true, and (ii) Γ' cannot be inadmissible solely in virtue of its divergence from the actual sentential expression of Ξ . What this principle suggests is that an agent's actual state of information does not have to be sensitive to how it is related to its basic data—there is no subjective or epistemic limit to how *de re* data receive their sentential expression—, but it should (or at least could) be sensitive to what is actually the case—which is a purely objective limit.

Even when presented in this abstract fashion, the following should be clear. First, where Γ_s is the sentential expression of Ξ at s , the lack of a subjective constraint on how Ξ relates to Γ_s is already sufficient to lead (at least potentially) to situations where an agent's basic information at s is indeed logically weaker than its basic information at the actual world. As we've seen, this opens the door to failures of transitivity of informational accessibility, and does so in a way that is quite different from more common such failures (in particular, the non-transitivity of the relation of being not discriminable by a non-perfect epistemic agent, see: Williamson [2000: chapt. 5]). This is for instance revealed by considering the non-factive relation of holding data for ϕ , which, if defined as the truth of ϕ at all worlds in $\|\Gamma\|$ with Γ the sentential expression of Ξ , can be shown to be non-reflective in virtue of the lack of a subjective constraint on admissible sentential expressions. The latter shows that being informed is a non-reflective state merely because it presupposes holding data that can be represented in some or other sentential form, and holding data is itself non-reflective. Put differently, a *de re* datum is by its nature relatively open-ended, and, because it might be opaque to the agent who has access to that datum, this agent does not necessarily have the cognitive resources to fully constrain this open-endedness.

Second, if there is an objective constraint or limit on the admissible sentential expressions of Ξ , then, if suitably conceived, that could lead to the symmetry of informational accessibility. The requirement that for any informationally accessible world s , Γ_s could not be such that its truthful part Δ_s would rule out the actual world is exactly such a constraint. Admittedly, this is all but a specific constraint on admissible sentential expressions. All I have done is specify the minimal outcome that should be respected. Since there is no need to be specific about what it means in general to be a sentential expression of a *de re* datum, there is also no need to be any more specific than what is required by the non-exclusion of the actual world. By only specifying the required outcome, the symmetry of informational accessibility is enforced in a minimal way. Rather than a lack of specificity, this is in fact a positive feature which precludes the symmetry of informational accessibility from collapsing into something that might look like an introspective feature. Unlike the non-reflectivity of being informed, the non-exclusion of the actual world cannot be retraced to a feature of hold-

ing data.¹⁶ Also, because no known doxastic or epistemic logic depends on an accessibility relation that is merely symmetric (and reflexive), the non-exclusion of the actual world cannot either be regarded as a familiar feature of knowledge that is merely presented in a less familiar guise.

6 CONCLUDING REMARKS

By understanding information as content that is not just true at the actual world, but can also never exclude that actual world, a model-theory which almost vindicates the *KTB*-analysis is obtained. Almost, that is, because the resulting logic does not have to be a normal modal logic. As strong necessitation was never explicitly defended, this does not pose a problem. The model-theory in question is valuable because it warrants the independence of the state of being informed. Information is not knowledge, because it does not come with all the modal properties that are usually associated with knowledge,¹⁷ and it is not belief, because it has to be true. More importantly, information also does not coincide with true belief; something is not information if it is derived from false data, and information privileges the actual world by never excluding it.

The design of models for being informed remains a tricky exercise, especially if the resulting logic is fixed in advance. Two reasons can be given for the difficulties we faced along the way. A first reason is that the suggestion that information is truth-constituted cannot easily be integrated within a standard Kripke-style semantics. This problem was solved by giving the actual world a special significance, but this is at best a partial rendition of the intimate connection between truth and information—it merely simulates the required behaviour within a Kripke-model. A second reason is that when information is understood as true well-formed and meaningful content such content is naturally perceived as a persistent commodity with a determinate or world-independent semantic content. The latter, however, implies a sort of monotonicity condition which immediately warrants the transitivity of the informational accessibility relation. Being more specific about the different kind of data an agent could have access to, leads to a sound explanation of why this otherwise natural assumption of persistence might fail.

Arguably, the main contribution of this paper is that it makes an information-based epistemology more plausible by showing that information can indeed be a cognitively more valuable, rather than merely a different kind of factive commodity than true belief. This is what the strong independence thesis

¹⁶This might look odd, because it is specified as a constraint on the data one holds. Yet, it does not have to be so, for the actual constraint might still only have an impact on that part of the data which actually constitute information; it simply wasn't defined that way.

¹⁷This does not entirely square with the analysis in Floridi [forthcoming], but the relevant issue here is that being informed does not coincide with the usual externalist accounts of knowledge.

aims at. Because it relates the properties of being informed to the nature of data and information, the proposed model-theory supports this thesis in, what I believe is, a worthwhile and interesting way. Still, even if one does not adhere to this broader aim, or if one believes that being informed should be conceived according to the Dretsian orthodoxy, this paper also offers a number of methodological insights that are of independent value. These insights mainly concern the nature and pitfalls of logical modelling, and is emphasised by showing that the original **KTB**-analysis could be seen as an exercise in reverse engineering.

REFERENCES

- ADAMS, F., forthcoming, Information and Knowledge à la Floridi, *Metaphilosophy*.
- ALLO, P., 2009, Reasoning about data and information, *Synthese*, **167**(2): 231-249.
- ARLO COSTA, H. AND R. PARIKH, 2006, Tracking truth: Knowledge and conditionals in the context of branching time, *Formal Epistemology Workshop*.
- BLACKBURN, P., M. DE RIJKE AND Y. VENEMA, 2001, *Modal Logic*, Cambridge University Press, Cambridge.
- COPELAND, J. B., 1983, Pure Semantics and Applied Semantics. A Response to Routley, Routley, Meyer, and Martin, *Topoi*, **2**(2):197-204.
- CRESSWELL, M., 1966, The Completeness of S0.5, *Logique & Analyse*, **9**(34): 263-266.
- DRETSKE, F., 1970, Epistemic Operators, *The Journal of Philosophy*, **76**(24): 1007-1023.
- DRETSKE, F., 1999, *Knowledge and The Flow of Information*, CSLI, Stanford.
- DUNN, J. M., 2008, Information in Computer Science, in: *Handbook on the Philosophy of Information*, Van Benthem, J. and P. Adriaans, eds., Elsevier, Amsterdam: 581-608.
- FELDMAN, R., 1974, An alleged defect in Gettier counter-examples, *Australasian Journal of Philosophy*, **52**(1): 68-69.
- FLORIDI, L., 2004a, Outline of a Theory of Strongly Semantic Information, *Minds & Machines*, **14**(2): 197-222.
- FLORIDI, L., 2004b, On the Logical Insolvability of the Gettier Problem, *Synthese*, **142**(2): 61-79.
- FLORIDI, L., 2005a, Is Information Meaningful Data?, *Philosophy and Phenomenological Research*, **70**(2): 351-370.
- FLORIDI, L., 2005b, Semantic Conceptions of Information, in: *Stanford Encyclopedia of Information*, Zalta, E. N., ed.
- FLORIDI, L., 2006, The Logic of 'Being Informed', *Logique & Analyse*, **49**(196): 433-460.
- FLORIDI, L., 2008, The Method of Levels of Abstraction, *Minds and Machines*, **18**(3): 303-329.
- FLORIDI, L., forthcoming, Semantic Information and the Network Theory of Account, *Synthese*.
- FUHRMANN, A., 1989, Reflective modalities and theory change, *Synthese*, **81**(1): 115-134.
- GIRLE, R., 2000, *Modal Logics and Philosophy*, Acumen, Teddington.

- HUGHES, G. E. AND M. J. CRESSWELL, 1996, *A New Introduction to Modal Logic*, Routledge, London / New-York.
- KRATZER, A., 2002, Facts: Particulars or Information Units?, *Linguistics and Philosophy*, **25(5)**: 655–670.
- LEMMON, E. J., 1959, Is There Only One Correct System of Modal Logic?, *Proceedings of the Aristotelian Society (Supplementary Volume)*, **33**: 23–40.
- LEWIS, D., 1996, Elusive Knowledge, *Australasian Journal of Philosophy*, **74(4)**: 549–67.
- PLANTINGA, A., 1974, *The Nature of Necessity*, Oxford University Press, Oxford.
- PRIEST, G., 2001, *An Introduction to Non-Classical Logic*, Cambridge University Press, Cambridge.
- SHIN, H. S., 1993, Logical Structure of Common Knowledge, *Journal of Economic Theory*, **60(1)**: 1–13.
- SHOPE, R. K., 1979, Knowledge and falsity, *Philosophical Studies*, **36(4)**: 389–405.
- WILLIAMSON, T., 2000, *Knowledge and Its Limits*, Oxford University Press, Oxford.

Patrick Allo
Postdoctoral Fellow of the Science Foundation (FWO)
CLWF (Vrije Universiteit Brussel)
IEG (Oxford) & GPI (Hertfordshire)
patrick.allo@vub.ac.be
homepages.vub.ac.be/~pallo/