

This is a preprint of a paper accepted for publication in

Synthese (Springer)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

It is a publisher's requirement to display the following notice:

The documents distributed by this server have been provided by the contributing authors as a means to ensure timely dissemination of scholarly and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

In the case of Springer, it is the publisher's requirement that the following note be added:

"An author may self-archive an author-created version of his/her article on his/her own website and his/her institution's repository, including his/her final version; however he/she may not use the publisher's PDF version which is posted on www.springerlink.com. Furthermore, the author may only post his/her version provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The original publication is available at www.springerlink.com."

Semantic Information and the Network Theory of Account

Luciano Floridi^{1,2}

¹Research Chair in Philosophy of Information and GPI, University of Hertfordshire; ²Faculty of Philosophy and IEG, University of Oxford.

Address for correspondence: Department of Philosophy, University of Hertfordshire, de Havilland Campus, Hatfield, Hertfordshire AL10 9AB, UK; l.floridi@herts.ac.uk

Abstract

The article addresses the problem of how semantic information can be upgraded to knowledge. The introductory section explains the technical terminology and the relevant background. Section two argues that, for semantic information to be upgraded to knowledge, it is necessary and sufficient to be embedded in a network of questions and answers that correctly accounts for it. Section three shows that an information flow network of type *A* fulfils such a requirement, by warranting that the erotetic deficit, characterising the target semantic information *t* by default, is correctly satisfied by the information flow of correct answers provided by an informational source *s*. Section four illustrates some of the major advantages of such a Network Theory of Account (NTA) and clears the ground of a few potential difficulties. Section five clarifies why NTA and an informational analysis of knowledge, according to which knowledge is accounted semantic information, is not subject to Gettier-type counterexamples. A concluding section briefly summarises the results obtained.

Keywords

Account; Epistemic Logic; Explanation; Gettier Problem; Information Theory; Network Theory; Network Theory of Account; Philosophy of Information; Semantic information.

1. Introduction

The problem addressed in this article may be phrased rather simply: how does semantic information upgrade to knowledge? The solution, articulated and supported in the following pages, is equally straightforward: semantic information upgrades to knowledge if and only if it is correctly accounted for. As the reader may readily suspect, the difficulty lies in the theoretical work required to understand the problem in detail, to explain and defend the solution successfully and to show how the two are properly related. These tasks will be undertaken in the next four sections. They are followed by some concluding remarks, which briefly summarise the results obtained in the article. Before entering into a thorough investigation of the subject matter, it might be helpful to clarify in this section some key concepts and the general framework in which they will be used, lest the lack of a shared vocabulary hinders the work that lies ahead.

In recent years, philosophical interest in the nature of information has been increasing steadily.¹ In particular, one of the current debates concerns the veridical nature of semantic information. As with many other interesting philosophical debates, this too is somewhat old,² but it has been re-ignited by the proposal to analyse semantic information in terms of *well-formed*, *meaningful* and *veridical data* (Floridi (2005), Floridi (2010b)). Admittedly, the analysis – according to which semantic information encapsulates truth, exactly as knowledge does – has attracted some criticisms for being too restrictive.³ Nevertheless, such criticisms have been proved unjustified⁴ and, as a result, there is now a growing consensus about the following approach.

Semantic information is primarily understood in terms of *content about a referent*, where content is analysed in terms of *well-formed* and *meaningful data*.⁵ Strings or patterns of data may constitute sentences in a natural language, but of course they can also generate formulae, maps, diagrams, videos or other semiotic constructs in a variety of physical codes, being further determined by their appropriate syntax (well-formedness) and semantics (meaningfulness). By “about a referent” one is to understand the ordinary and familiar way in which some well-formed and meaningful data, constituting semantic information, concern or address a topic. Following Dretske (1981) and Dretske (1988), one may easily recognise this

¹ For an updated overview and guide to the literature see Floridi (2004a), Floridi (2010a) and Floridi (2010b).

² For example, Bar-Hillel and Carnap (1953) and Devlin (1991) argued against the veridical nature of semantic information, whereas Dretske (1981) and Grice (1989) argued in its favour.

³ See for example the discussion in Fetzer (2004), who criticises Floridi (2004c), with a reply in Floridi (2005); or the objections moved by Colburn (2000a), Colburn (2000b) and Dodig-Crnkovic (2005).

⁴ Floridi (2007) and Sequoiah-Grayson (2007).

⁵ On the analysis of data see Floridi (2008a).

“aboutness” feature in propositional attitudes such as “Mary is informed that the water in the electric kettle is boiling”, where “being informed” is used in the *statal*⁶ sense, i.e., in the sense that Mary holds (rather than is receiving) that information. This is the condition into which *a* enters (and may remain, if *a* is not a memoryless agent) once *a* has *acquired* the information (*actional* state of being informed as becoming 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 the crime was committed. In the rest of this paper, we shall be concerned with only this standard, *statal* and epistemically-oriented or factual concept of semantic information. I shall also take the liberty of dropping the qualification “semantic” when it becomes redundant.

In Floridi (2005), I argued that a definition of semantic information in terms of alethically-neutral content – that is, strings of well-formed and meaningful data that can be additionally qualified as true or untrue (false, for the classicists among us), depending on supervening evaluations – provides only necessary but insufficient conditions: if some content is to qualify as semantic information, it must also be *true*. One speaks of false information in the same way as one qualifies someone as a false friend, i.e. not a friend at all. In other words:

“[...] 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).

All this leads to a refinement of the initial definition into:

[DEF] *p* qualifies as semantic information if and only if *p* is (constituted by) *well-formed, meaningful and veridical data*.

[DEF] captures the general consensus reached by the debate and mentioned at the outset of this section. According to it, semantic information is, strictly speaking, inherently *truth-constituted* and not a contingent *truth-bearer*, exactly like knowledge but unlike propositions or beliefs, for example, which are what they are independently of their truth values and then, because of their truth-aptness, may be further qualified alethically.

⁶ The distinction is standard in linguistics, where one speaks 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)”). In this paper, I deal only with the *statal* sense of “is informed”. This is related to cognitive issues and to the logical analysis of an agent’s “possession” of a belief or some knowledge.

[DEF] offers several advantages. For example, it plays a crucial role in the solution of the so-called Bar-Hillel-Carnap Paradox (Floridi (2004c)) and provides a necessary element for a subjectivist theory of epistemic relevance (Floridi (2008b)). Here, it is worth emphasising that it forges a robust and intuitive link between semantic information and knowledge. More specifically, the veridical thesis contained in [DEF] corresponds to the one characterising the definition of knowledge. Taking advantage of this parallelism, one may rely on the ordinary apparatus of modal logic (e.g. Chellas (1980)) to formalise “ a is informed that p ” as $I_a p$, and hence formulate the veridicality thesis of semantic information in terms of the so-called veridicality axiom $\Box\phi \rightarrow \phi$, also known as **T**, **M** or **K2**, thus:

$$[\text{VT}] \quad I_a p \rightarrow p$$

The intended interpretation of [VT] is that a is informed (i.e., holds the information) that p only if p is true. In Floridi (2006), I have shown that information logic (*IL*) can then be satisfactorily modelled in terms of an interpretation of the relation “ a is informed that p ” based on the axioms of normal modal logic **B**. [VT] associates *IL* to epistemic logics (*EL*) based on normal modal logics **KT**, **S4** or **S5**. And it differentiates both *IL* and *EL* from doxastic logics (*DL*) based on **KD**, **KD4** and **KD45**, since, of course, no *DL* satisfies the veridicality axiom. It follows that *IL* allows truth-encapsulation (i.e., it satisfies [VT]) without facing either epistemic or doxastic collapse (i.e., merely morphing into another epistemic or doxastic logic). So knowledge encapsulates truth because it encapsulates semantic information, which, in turn, encapsulates truth, as in a three dolls matryoshka.

Despite its advantages, any approach endorsing [DEF] raises two major questions (Floridi (2004b)). One is upstream:

- a) what does it mean for semantic information to be truthful?

The other is downstream:

- b) how does semantic information upgrade to knowledge?

Both questions are prompted by [DEF] but neither is specifically about [DEF] only, and each fails to provide a starting point for a *reductio ad absurdum*. They are rather information-theoretical versions of classic conundrums: (a) is a request for a theory of truth and (b) is a request for a substantive analysis of knowledge. Since the goal of this paper is to seek to answer only (b), let me brush (a) away by adding two final clarifications.

First, in the rest of the paper [DEF] and [VT] will be assumed as unproblematic. The reader who disagrees and finds them controversial may still be willing to concede that, in many cases, semantic information is true, and grant that these will be the only cases with which we are going to be concerned.

Second, in Floridi (forthcoming) I argue that several classic theories of truth may be suitable to deal with (a), but that a new correctness theory of truth might be preferable. In this case too, the reader need not agree with me. As long as we *stipulate* in this paper that we shall deal only with semantic information that is true, I shall leave it to the reader to opt for her favourite way to substantiate such a claim.

With all these clarifications in place, we are now ready to look more carefully into the nature of the upgrading problem.

2. The Nature of the Upgrading Problem: Mutual Independence

[DEF] nests semantic information into knowledge so tightly that one is naturally led to wonder whether anything else might be missing, in order to upgrade from the weaker to the stronger phenomenon, and hence between their corresponding concepts. Indeed, the threshold can be so fine that one may often overlook it, and thus fail to distinguish between the two propositional attitudes, treating “Mary is informed that the water in the electric kettle is boiling” and “Mary knows that the water in the electric kettle is boiling” as if they were always interchangeable without any conceptual loss. In everyday life, this might be the norm, and the conflation is usually harmless: it can hardly matter whether the bus driver is informed or knows that the traffic light is red, as long as he behaves accordingly. Philosophically, however, the distinction captures an important difference, and hence it is important to be more accurate. It takes only a moment of reflection to see that one may be informed (hold the information) that p without actually knowing that p . Not only because holding the information that p does not have to be a *reflective* state (although it is not necessarily the case that $I_a p \rightarrow II_a p$, one may also object that $K_a p \rightarrow KK_a p$ is notoriously controversial as well) but also because, even when it is, it might still arguably be *opaque* and certainly *aleatoric* (epistemic luck), whereas knowledge cannot.

Consider *opaqueness* first. It is open to reasonable debate whether a messenger carrying (in her memory, in her hand on in a pocket, it does not matter) an encrypted message p that she does not understand – even if she is informed that she carries p – may be said to hold the information that p . On the one hand, one may argue that she is not genuinely

informed that p . On the other hand, one may retort that, if she can deliver the information that p (and we are assuming that she can) then she can legitimately be said to be informed that p or hold that information. The interesting point here is not to solve the dispute, but to note that the dispute itself is reasonable, whereas, if the same messenger knows that p , there can be no doubt that she must also understand the information carried by p . It might be open to debate whether holding the information that p is necessarily a non-opaque state, but such a dispute would be pointless in the case of knowing that p .⁷

Next, consider *epistemic luck*. When asking how semantic information may be upgraded to knowledge, we are not asking what further axioms may need to be satisfied by **K**. For even if we were to upgrade **K** all the way up to **S5**, as we are perfectly and indeed easily able to do, we would still be left with the problem of the non-aleatoric nature of knowledge. Now, raising the issue of epistemic luck serves two purposes. It further strengthens the conclusion that there is a clear difference between (holding) the semantic information that p and (having) the knowledge that p . And it points in the direction of what might be missing for semantic information to upgrade to knowledge.

Regarding the first purpose, epistemic luck affects negatively only knowledge but not semantic information. To see why, one may use a classic Russellian example: if one checks a watch at time t and the watch is broken but stopped working exactly twelve hours before ($t - 12$) and therefore happens to indicate the right time $t - 12$ at t , one is still informed that the time is t , although one can no longer be said to know the time. The same applies to a more Platonic example in which a student memorises, but fails to grasp, the proof of a geometrical theorem: she is informed (holds the information) that the proof is so and so, but does not really know that the proof is so and so. Generalising, Russell- Plato- or Gettier-type counterexamples may succeed in degrading “knowing” to merely “being informed” (“holding the information that”), but then “being informed” is exactly what is left after the application of such counterexamples and what remains resilient to further subjunctive conditionalization.

Regarding the second purpose, epistemic luck, if properly diagnosed, should be understood as a symptom of the disease to be cured, rather than the disease itself, and therefore as providing an indication of the sort of possible treatment that might be required. To explain how, let me introduce the following thought experiment.

⁷ Note that, as far as opaqueness is concerned, knowing and being informed that p might collapse from a first-person perspective, but not from a second-person perspective, but that the collapse is not necessary if one first-personally, reflectively realises that one does not really have a clue about p , apart from holding the information that p .

Imagine a memoryless Oracle, who can toss a magic coin to answer Boolean questions.⁸ The coin is magic because it unfailingly lands heads whenever the correct answer to the Boolean question is yes, and tails whenever it is no. The Oracle has two alternatives. Either she remains silent and does not answer the Boolean question at all. This happens whenever the question cannot be answered uncontroversially or unambiguously either yes or no. Examples include “is the answer to this question ‘no’?”, “do colourless green ideas sleep furiously?”, or “will there be a naval battle tomorrow?”. Or she can toss the coin and thereby give the correct answer by reading the result aloud. Let us assume that there is no significant time lag between question and answer: if no answer is provided within a few seconds, it means that the Oracle will provide no answer at all (recall that she has no memory). It seems clear that the Oracle is the ultimate reliable source of information, but that she has no propositional knowledge. Imagine now a Scribe. He knows that heads means yes and tails means no. He asks answerable Boolean questions of the Oracle and methodically records her correct answers in his scroll, thus acting as an external memory. The entries in the scroll are ordered pairs that look like this:

[...]

<Q: “Is Berlin the capital of France?” A: “no”>

<Q: “Is Berlin in Germany?” A: “yes”>

<Q: “Is Berlin the capital of Germany?” A: “yes”>

<Q: “Has Berlin always been the capital of Germany?” A: “no”>

<Q: “Did Berlin become the capital of reunified Germany in 1990?” A: “yes”>

<Q: “Is Berlin the largest city in Europe?” A: “no”>

<Q: “Is Germany in Europe?” A: “yes”>

[...]

The scroll will soon resemble a universal Book of Facts, with each entry (each ordered pair) as an information packet. Now, it has been customary, at least since Plato, to argue that the scroll contains at most information but not knowledge, and that the Scribe may at best be informed (even counterfactually so: if p were not the case, the Oracle would not have given the answer she has given), but does not know, that e.g., “Germany is in Europe”, because knowledge cannot be aleatoric. This much seems uncontroversial. What is less clear is the

⁸ Note that the example of the Oracle, the magic coin and the Scribe should not be confused with BonJour’s example of Norman, who is assumed to be a perfectly reliable clairvoyant (BonJour (1985). This because the latter is supposed to have no evidence at all indicating that he is a clairvoyant, and has no way of realising that his beliefs are caused by his clairvoyance, therefore having no justification for them in terms of his clairvoyant capacities.

exact nature of the problem. By seeking to uncover it, we enter into the second half of this section: understanding what the difference is between semantic information and knowledge.

It might be tempting to argue that epistemic luck is the actual problem because, if we were to depend on it for our knowledge of reality, sooner or later we would run into trouble. We cannot be lucky in all circumstances, and, even in the same circumstances, we might have been unlucky, so other epistemic agents might easily disagree with us, for they might enjoy different degrees of epistemic luck, which means that further coin-tossing would hardly help and that interactions with the world and other agents embedded in it might be utterly haphazard. Yet giving in to this temptation would be short-sighted. Semantic information is impervious to epistemic luck whereas knowledge is not, but epistemic luck is only a criterion that helps us to differentiate between the two, a device used to cast light on the real difficulty. This is why the Oracle-Scribe example ensures that we see that the erratic and unreliable nature of epistemic luck plays no role. By definition, the Oracle is infallible in the sense that she always provides the correct answer, and the Scribe is fully reliable, in the sense that he is perfectly able to record and later access the right piece of information. Moreover, if a second Scribe were to consult the Oracle, he would obtain the same piece of information (ordered pairs). Indeed, the Oracle would be the ultimate Salomonic judge of any Boolean dispute. Nevertheless, we are facing a case of information at most, not of knowledge. If the problem with epistemic luck were that we may never have it, or that we might not have had it, or that we may never have enough of it, or that different epistemic agents may have different degrees of it, then surely the argument should be that hoping or trusting to be always (by oneself) and consistently (with respect to others) lucky cannot be a successful epistemic strategy even in the short term, rather than, when one is actually lucky, that one still fails to win the epistemic game. But this is exactly what we are asserting above, and rightly so. There is indeed something epistemically unsatisfactory with answering questions by tossing a coin, even when the coin is magic, yet the aleatoric nature of the process is not the fundamental difficulty, it is only the superficial symptom, and that is why taking care of the features that are most obviously problematic by using a magic coin clarifies that we are still failing to tackle the real issue.

At this point, one may concede that, yes, epistemic luck is only evidence of a more profound failure, but then conclude that this failure might be related to truth-conductivity, subjective justification or a combination of both. Yet this too would be a mistake. By hypothesis, the procedure of asking Boolean questions of the Oracle and recording her answers is as truth-conducive as anyone may wish it to be. Likewise, the Scribe holding the

information contained in the scroll is perfectly justified in doing so, and his attitude is indeed very rational: given the circumstances and the availability of the Oracle, he ought to consult her, and rely on her answers both in order to obtain information and in order to justify and manage (increase, refine, upgrade etc.) his own information states (set of beliefs, in the doxastic vocabulary). He is not prey to some wishful thinking, but sensibly constrained by his source of information. So, epistemic luck is indeed a warning sign but neither of some alethic ineffectiveness on the side of the epistemic process nor of some rational laxity on the side of the knowing subject.

The problem lies elsewhere: the aleatorization of information (i.e., the randomization of the ordered pairs or scroll entries in the lucky sense seen above) dissolves the bonds that hold it together coherently (its consilience), like salt in water. If one analyses each entry in the scroll, there is clearly nothing epistemically wrong either with it or with the subject holding it. What the aleatoric procedure achieves is the transformation of each piece of information into a standalone, mutually independent item, entirely and only dependent on an external and unrelated event, namely, the tossing of the magic coin. The problem is therefore systemic: aleatorization tears information items away from the fabric of their inter-relations, thus depriving each resulting information packet of its potential role as evidence and of its potential value for prediction or retrodiction, inferential processes and explanation.

Consider our thought experiment once again. This time, in order to explain mutual independence, let us assume that the Oracle uses an ordinary coin and that we have no reassurance about the truth or falsity of each ordered pair so obtained. Each $\langle Q_x, A_x \rangle$ will now have a probability value P independent of any other⁹ ordered pair $\langle Q_y, A_y \rangle$ (for $x \neq y$), that is, $P(\langle Q_x, A_x \rangle \cap \langle Q_y, A_y \rangle) = P(\langle Q_x, A_x \rangle)P(\langle Q_y, A_y \rangle)$. More generally, the scroll will contain only mutually independent entries, in the precise sense that any finite subset S_1, \dots, S_n of ordered pairs listed in the scroll will satisfy the multiplication rule: $P(\bigcap_{i=1}^n S_i) = \prod_{i=1}^n P(S_i)$. This feature is somewhat hidden when the coin is magic, since, in that case, each ordered pair and any finite subset of them in the scroll has probability 1. But consider what happens in the process of making an ordinary coin increasingly better at providing the correct answer (i.e. more and more “magic”): all the difficulties concerning chance and unreliability, truth-conductivity and subjective justification gradually disappear, until, with a perfectly magic coin, total epistemic luck indicates no other problem but the semantic lack (if

⁹ This is to ensure that no confusion is caused by self-referential independence (strictly speaking x is independent of itself if and only if its probability is one or zero).

we are trying to upgrade semantic information to knowledge) or removal (if we are trying to downgrade knowledge to semantic information) of any structural pattern stitching the various pieces of information together.

Such mutual independence is not yet a difficulty *per se* yet, but it finally points towards the problem that we need to solve. As Dummett (2004) nicely puts it

We do not merely react piecemeal to what other people say to us: we use the information we acquire, by our own observation and inferences and by what we are told, in speech and writing, by others, to build an integrated picture of the world” (p. 29).

Yet, by definition, mutually independent pieces of information cannot yield this integrated picture of the world because they cannot *account* for each other, that is, they cannot answer the question *how come* that $\langle Q_x, A_x \rangle$. Both italicised expressions require clarification.

Plato¹⁰ famously discusses the importance of embedding truths (our packets of semantic information) into the right network of conceptual interrelations that can “provide reason” (*logon didonai*) for them in order to gain knowledge of them. Plato seems to have meant several different things with “provide reason”, as this could refer to giving a definition, a logical proof, some reasonable support (e.g. dialectically), an explanation (e.g., causally) or some clarification (e.g., through an analogy), depending on the context. We shall see that this range of meanings is worth preserving. It is roughly retained in English by “giving a reasoned account” or simply “accounting”, hence the use of the term above.¹¹

Aristotle, not less famously, discusses the range of questions that an account may be expected to answer. For our purposes, we may organise them into teleological (future-oriented why, or what for, or for what goal or purpose), genealogical (past-oriented why, or where from, or through which process or steps) and functional questions (present-oriented why, or in what way, or according to which mechanism). Again, in English “how come” captures these different meanings without too much semantic stretching. This point was famously stressed in the philosophy of biology by Mayr:

The functional biologist is vitally concerned with the operation and interaction of structural elements, from molecules up to organs and whole individuals. His ever-repeated question is ‘How?’ How does something operate, how does it function? [...] The chief technique of the functional biologist is the experiment, and his approach is essentially the same as that of the physicist and the chemist. [...] The evolutionary biologist differs in his method and in the problems in which he is interested. His basic question is ‘Why?’ When we say ‘why’ we must

¹⁰ See for example Taylor (1967) and more recently Taylor (2008), esp. pp. 185-187.

¹¹ This epistemic sense should not be confused with the sense in which “account” is technically used in communication studies, where it refers to the common practice of justifying one’s own behaviour, see chapter 13 in Whaley and Samter (2006).

always be aware of the ambiguity of this term. It may mean ‘How come?’ but it may also mean the finalistic ‘What for?’ When the evolutionist asks ‘Why?’ he or she always has in mind the historical ‘How come?’ [...] We can use the language of information theory to attempt still another characterization of these two fields of biology [i.e. functional biology and evolutionary biology]”.¹²

If we apply this clarification to our examples, when someone asks today “how come that Berlin is the capital of Germany?” one may be asking what future purposes this might serve (teleological question), or which events in the nineties led to the transfer of the capital from Bonn to Berlin (genealogical question), or (admittedly less obviously in this example) how Berlin works as the re-established capital of a re-unified Germany (functional question). “How come” questions (henceforth HC-questions) may therefore receive different answers. “How come that the water in the electric kettle is boiling?” may receive as an answer “because Mary would like some tea” (teleological account), or “because Mary filled it with water and turned it on” (genealogical account), or “because electricity is still flowing through the element inside the kettle, resistance to the electric flow is causing heat, and the steam has not yet heated up the bimetallic strip that breaks the circuit” (functional account).

In the next section, we shall see that the wide semantic scope of both expressions (“account” and “HC-questions”) is an important feature essential to develop a sufficiently abstract theory that can show how information can be upgraded to knowledge. At the moment, the previous clarifications suffice to formulate more precisely our problem (P) and working hypothesis (H) to solve it, thus:

P) (a packet of) semantic information does not qualify yet as (an instance of) knowledge because it raises HC-questions that it cannot answer;

H) (a packet of) semantic information can be upgraded to (become an instance of) knowledge by having the HC-questions it raises answered by an account.

What is an account then, and how does it work?

3. Solving the Upgrading Problem: the Network Theory of Account

Each piece of semantic information is an answer to a question. As a whole, it poses further questions about itself that require the right sort of information flow in order to be answered correctly, through an appropriate network of relations with some informational source. Until recently, it would have been difficult to transform this general intuition about the nature of epistemic account into a detailed model, which could then be carefully examined and

¹² Mayr (1961), the citation is from the text reproduced in his *Toward a New Philosophy of Biology: Observations of an Evolutionist*, pp. 25-26.

assessed. Fortunately, new developments in an area of applied mathematics and computational algorithms known as *network theory* (Ahuja et al. (1993), Newman et al. (2006)) has provided all the technical and conceptual resources needed for our task.

The task is fairly simple: we need to construct a network through which the right sort of information flows from a source s , to a sink target t . In this network, t poses the relevant questions and s accounts *for* t if and only if s provides the correct answers. If the biconditional holds, we shall say that the whole network yields an account *of* t . Let us see the details.

We start by modelling the network as a finite directed graph G , representing the pattern of relations (a set E of *edges*) linking s and t . The edges work like communication channels: they have a set capacity c (e.g., how much information they can possibly convey) and implement an actual flow f (e.g., the amount of information they actually convey), which can be, at most, as high as their capacity. The path from s to t is usually mediated, so we shall assume the presence of a set (V) of other nodes (called *vertices*) between s and t that relay the information. Figure 1 provides an illustration.

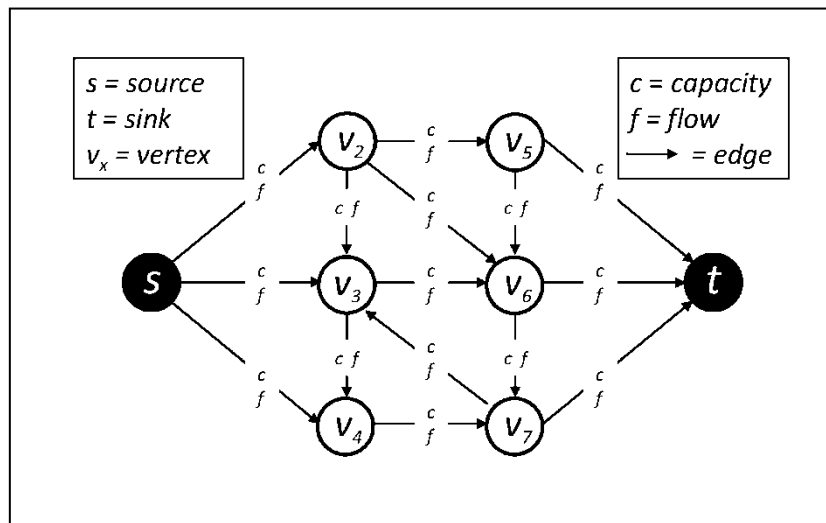


Figure 1 Example of a Flow Network

More precisely, the system just sketched qualifies as a flow network if and only if it satisfies the following conditions (where u and v are any two vertices generating an edge):

1. $G = (V, E)$ is a finite directed graph in which each edge $(u, v) \in E$ has a capacity $c(u, v) \geq 0$. Although we shall assume that c could be real-valued, for our purposes we may deal only with non-negative, natural values;
2. in G there are two special vertices: a source s , and a sink t ;

3. every vertex lies on some path from s to t ;
4. any (u, v) that is not an edge is disregarded by setting its capacity to zero;
5. a flow is a real-valued function on pairs of vertices $f: V \times V \rightarrow \mathbb{R}$, which satisfies the following three properties:
 - i. Capacity Constraint: $\forall v, u \in V, f(u, v) \leq c(u, v)$, that is, the flow along an edge can be at most as high as the capacity of that edge;
 - ii. Skew Symmetry: $\forall v, u \in V, f(u, v) = -f(v, u)$, that is, the net flow forward is the opposite of the net flow backwards;
 - iii. Flow Conservation: $\forall v, u \in V$ and $u \neq s$ and $u \neq t, \sum_{w \in V} f(u, w) = 0$, that is, the net flow to a vertex is zero, except for s , which generates flow, and t , which consumes flow. Given (b), this is equivalent to flow-in = flow-out.

The next step is to transform the flow network into an *information* flow network A , which can successfully model the process through which some semantic information is accounted for.

Since A is a flow network, it satisfies all the previous five conditions. In order to obtain a rather idealised but still realistic model of informational accounting, A needs to satisfy the following additional conditions:

a) *Single commodity*. This is a standard simplification in network theory. In A there is only one good (information as answers to questions) that flows through the network, with no constraint on which vertex gets which part of the flow. In real life, it usually matters how the flow is distributed through which vertices, but this feature would only increase the complexity of A with no heuristic added-value for our present purpose. Multi-commodity flow problems turn out to be NP-complete even for natural-valued flows and only two commodities. This is a good reminder that A is meant to be an abstract, conceptual model of the process of accounting, not a blueprint for some algorithmic application.

b) *Single source*. This is another standard assumption in network theory since, even if there were more than one source s_n , we could always add a special supersource S of information linking it to all the other s_n with edges with infinite capacity. By pushing the maximum flow from S to t we would actually produce the maximum flow from all the s_n to t .

c) *Redundancy = 0*. Intuitively, we assume that each packet of information is sent only once. More precisely, the vertices between s and t are not real secondary sources but rather ideal witnesses, constructed by means of a partition of the set of Boolean answers possible within the network (capacity) and actually required by t (flow). This because, contrary to ordinary physical flow networks (e.g. water through pipes, automobile traffic through a road network

and so forth), in A , s could send the same piece of information repeatedly through different channels without any loss, and this would make it more difficult to quantify it. It is possible to revise (c) by applying linear logic constraints to network theory that safeguard a realistic deployment of packets of information (interpreted as truths) as resources, but it would not be philosophically useful here.¹³

d) *Fidelity* = 1. Following information theory, we assume a completely accurate transfer of information from s to t . This means no noise and no equivocation, as specified in (e) and (f);

e) *Noise* = 0. Recall that, in information theory, noise is defined as any information received but not sent. In A , this means that any vertex different from s adds nothing to the information provided by s . Again, it is possible, but not philosophically useful, to model more complex scenarios, known as networks with gains, in which at least some vertices have a real-valued gain $g \geq 0$ such that, if an amount of information x flows into v then an amount gx flows out of v .

f) *Equivocation* = 0. In information theory, this is information sent but never received.

g) *Processing* = 0. This follows from conditions 5.i-iii and (e)-(f): every vertex between s and t merely retransmits the information it receives, without elaborating it, coding it or even reinforcing it (as repeaters do). Recent research on network information flow (Ahlsweide et al. (2000), Yeung (2008)) has proved that, in real circumstances, information can be multicast at a higher rate by using network coding, in which a receiver obtains as input a mix of information packets and elaborates which of them are meant to reach a sink. Yet this refinement is not essential for our purposes.

h) *Memory* = 0. As in (g), every vertex between s and t does not register the information flow, it merely multicasts it (see (j) below).

i) *Cost* = 0. Again, following information theory, we shall disregard any cost involved in the transmission of information from one vertex to another. Network theory does provide the technical tools to handle this problem, by assigning to each edge $(u, v) \in E$ a given cost $k(u, v)$ and then obtaining the overall cost of sending some flow $f(u, v)$ across an edge as $f(u, v) \times k(u, v)$. This would be crucial in any logistic context in which transmission costs need to be minimised, but it can be disregarded here.

j) *Routing scheme*: multicast. Realism requires that s may deliver its information to many vertices simultaneously.

¹³ More specifically, and as one of the anonymous reviewers rightly suggested, it would be interesting to explore the effects of dropping contraction and weakening.

The information flow network A that we obtain from conditions (1)-(5) and (a)-(j) is a standard idealization, which contains all the elements required for our theoretical purposes, but does not abstract from any feature that would be relevant. It merely simplifies our task, which is now that of showing how A models the process of accounting for some semantic information.

We have seen that epistemic luck dismantles the machinery of knowledge into its constitutive components, leaving them in perfect epistemic condition but piled up in a heap, unable to account properly for each other. This mutual independence is the semantic loss that needs to be tackled in order to upgrade semantic information to knowledge. We need to restore the epistemic fabric within which each piece of information is a thread. This is what (an implementation of) the information flow network A achieves, in the following way.

The semantic information to be accounted for is the sink t . Using our toy example, let us set $t =$ “the water in the electric kettle is boiling”. The sink t poses a number of HC-questions. For the sake of simplicity, we shall disregard the important fact that such questions will be formulated for a particular purpose, within a context and at some level of abstraction. Further simplifying, we transform each HC-question into a Boolean question. For example: “how come that the water in the electric kettle is boiling?” may become “Is the water in the electric kettle boiling because Mary wants some tea?”. So, t comes with an information deficit, which is quantifiable by the number of Boolean answers required to satisfy it. In our example, let us assume that t requires 10 Boolean answers. Accounting for t means answering t 's HC-questions correctly, that is, providing the necessary flow of information that can satisfy t 's Boolean deficit satisfactorily. The required answers come from the source s , but the connection between s and t is usually indirect, being mediated by some relay systems: a document, a witness, a database, an experiment, some news from the mass media, may all be vertices in the information flow, with the proviso that they are constituted by their capacity and flow values according to condition (c) above. Following standard practice, and again for the sake of illustration only, let us assume the presence of six intermediate vertices. Each vertex v_x and the source s can provide a maximum number of Boolean answers. This is the capacity c . An edge is now a vector with direction, indicating where the answers come from, and magnitude, indicating how many answers the starting point could provide in theory. In Figure 2, the edge (v_5, t) , for example, can convey up to 4 Boolean answers, while the total capacity of the selected area (known as a *cut*) is $20 + 5 + 4 = 29$.




Figure 2 An Information Flow Network with Capacities and Cut

The next task is to identify the flow of information, that is, the set of Boolean answers actually required by t . Information percolates through the network but, ultimately, it is assumed to come from a single source s . In most cases, the source and the layers of vertices have a much higher informational capacity c . This because s and any v_x may be a very rich source of information, like a complex experiment, a perceptual experience, an encyclopaedia, a large database, a universal law, a whole possible world, the universe or indeed our Oracle with a magic coin (recall the Supersource in (b) above). Clearly, s and the vertices between s and t can answer many more questions than the ones posed by t . Figure 3 shows a possible flow of information, given our example. The vectors (edges) now have a magnitude constituted by the numeric values for c (first value) and f (second value, in bold).




Figure 3 An Information Flow Network with Capacities and Flow

If all the HC-questions posed by t are correctly answered by s through A , then s accounts for t and A is an account of t . If the answers are incorrect or insufficient in number, then t remains unaccounted for, A is insufficient and may need to be improved or replaced. Before exploring some of the features of the model just proposed, let us take stock of the results obtained so far.

There is a difference between semantic information and knowledge, which can be highlighted by epistemic luck. The difference is that semantic information lacks the necessary structure of relations that allow different packets of information to account for each other. It follows that, for semantic information to be upgraded to knowledge, it is necessary to embed it in a network of relevant questions and corresponding correct answers. An information flow network of type A fulfils such a requirement, by making sure that the erotetic deficit, which the target semantic information t has by default, is satisfied by the flow of correct answers, provided by an informational source s .

At this point, we can become better acquainted with what I shall refer to as a network theory of account (henceforth NTA) by appreciating some of its major advantages and clearing the ground of a few potential difficulties. This is the task of the next section.

4. Advantages of a Network Theory of Account

Let me first dissipate two concerns regarding NTA: that it might be too general or, in another sense, too specific.

Regarding its generality, I suggested above that it should be appreciated as an advantage. In short, it is the same generality that we find in the supply and demand model in economics. More specifically, at least two main valuable features of NTA (the interested reader is referred to Chen (2003) for further details) make such generality welcome:

- *robust scalability*: the description of NTA outlined above highlights only the features that are of main interest here – in order to make sense of what an account for some semantic information amounts to, and hence how semantic information may be upgraded to knowledge – but it should be clear that more complex scenarios can be easily and elegantly handled by NTA, by relying on the wealth of theoretical results and successful applications of network theory; and
- *flexible semantics*: network theory deals as abstractly as possible with flow structures (this generality may remind the reader of category theory). Following network theory, NTA inherits a similarly high degree of generality. It thus provides a unifying

notation, terminology and modelling approach, which make it suitable for a variety of specific applications, depending on how the sink, the source, the layers of vertices, the edges, the information flow and the complexity of the network (see conditions (a)-(j) in the previous section) are interpreted.

Because of its robust scalability and flexible semantics, NTA has a very wide scope of applicability. It can make sense of simple and mundane forms of accounting, for example by modelling the account that might be given of why the water in the electric kettle is boiling, or a police account of an accident (Stering (2008)). But suitably interpreted, NTA may also model *scholarly* and *scientific explanations* of the semantic information under investigation, following a variety of deductive-nomological, inductive-statistical, teleological, functional, analogical, historical or psychological approaches. For example, in the deductive-nomological and the inductive-statistical cases, the source contains initial conditions and law-like generalizations, which enable the network to answer the HC-questions posed by the semantic information t concerning the event to be explained. Indeed, since the eighties some theories of explanation (Van Fraassen (1980) and Achinstein (1983), but see more recently Walton (2007)) have developed and defended erotetic approaches to explanation that are close to the analysis of account supported in this article. Furthermore, following Plato, NTA could also give rise to a *logical deduction* of t , which would then represent the theorem T to be proved starting from the premises contained in the source. In this case, NTA allows both constructive (proving T by answering the question “how come that T is the case?”) and non-constructive approaches (proving T by answering the question “how come that $\neg T$ is not the case?”). Alternatively, the network of vertices and edges and the probability of received answers through such paths could be given a Bayesian interpretation, as in Van Fraassen’s erotetic model of explanation (Van Fraassen (1980)). The list could be expanded. The interesting point to be stressed here is not how versatile NTA is, but that NTA’s plasticity is a valuable feature, obtained by raising the level of abstraction at which one may analyse what it means for some information (supply) to make sense of some other information by answering (satisfying) its HC-questions (demand).

Regarding its specificity, someone convinced by the previous defence of NTA’s abstract generality may still harbour doubts about its quantitative approach. Could one really attach numbers to both c and f in A ? The positive answer is twofold.

First, it is far from being implausible that one may be numerically precise about how many questions need to be answered (value of f) in order to account for t , and how many answers any v_x and s may provide in theory (value of c). For example, one may assume v_x and

s to be finite databases, from which a finite set of correct, Boolean answers might be extracted, only some of which will be relevant to answer t 's HC-questions correctly. Equally possible, and indeed feasible, is to transform any HC-question into a Boolean question. In this case, one only has to think of the twenty questions game. A handful of Boolean questions is often sufficient to account for some given information.

Second, even if the quantitative approach were utterly implausible and unfeasible, it would still remain logically possible, and this would be sufficient for our purpose, which is to clarify, through a simplified model, what accounting for some semantic information consists in. A related, but different, problem concerns what Boolean questions to ask and how to answer them correctly. About this, NTA remains silent, and rightly so. It is not up to a model of informational accounting to determine the nature of the specific information that will be conveyed through the network, in the same way that the engineer is not concerned about the sort of messages that will be exchanged through a digital system of communication. This does not make the formalism pointless. On the contrary, NTA can take full advantage of the theorems and algorithms available in network theory (Jungnickel (1999) and Cormen et al. (2001)). For example, in the previous section we have already encountered a limiting result with respect to the NP-completeness of two-commodity, natural-valued flow problems and we shall presently see that a classic result, known as the max-flow min-cut theorem (Elias et al. (1956), Ford and Fulkerson (1956)), can also be very enlightening.

With the previous concerns at least mitigated, if not entirely dispelled, it is useful to highlight now some positive and interesting features of NTA that will finally introduce the next section. I shall deal with them rather schematically because the goal is to give a general flavour of NTA, and only the last feature will be crucial for our current purposes.

NTA supports an epistemic or semantic interpretation of “accounting”. That is, both t and s are informational in nature: they may refer to, but are not in themselves, phenomena or events. So are the questions asked and the answers provided. However, nothing prevents the realist from expanding NTA into a more ontologically-committed approach, if she so wishes, depending on the theory of truth that she selects to ground the truthfulness of the information flowing through the network.

NTA lends itself to fast procedures of revision. It is clear that, if an account of t is unsatisfactory and there is a problem in the required flow of information, NTA helps to detect this more quickly, to check where the difficulty might be and to improve the configuration of A accordingly. In particular, inferences to the best account may require improved or different sources of information, if s does not answer correctly and in full all the questions posed by t .

It is true that, in ordinary life, our social network of interlocking *As* is often resilient to radical changes but, even in such stable scenarios, very significant experiences may lead to equally radical revisions. Othello or Hamlet are just idealised cases. This holds true in science as well. A Copernican revolution is any dramatic alteration in the accounting networks that cement our information.

The previous point is related to the possibility that not all semantic information might be upgradeable to knowledge. Sometimes HC-questions cannot be properly asked (recall Russell’s example of the broken watch) or are unanswerable anyway because the source that could account for *p* is inaccessible or cannot be established, as for example when trying to account for some archaeological information concerning a prehistorical civilization. Knowledge is a much rarer phenomenon than information.

With regard to ways of improving an interpreted *A*, the max-flow min-cut theorem is a classic cornerstone in optimization theory and flow network, which turns out to be very useful for NTA as well. The theorem states that, given a single-source, single-sink, flow network *N*, the maximum amount of flow *f* in *N* is equal to the capacity *c* of a minimum cut in *N*. The idea is simple: the maximum flow in *N* is determined by its bottleneck. Applied to *A*, this can be intuitively appreciated by looking at Figure 4, where it is easy to see that the maximum information flow reaching *t* is equivalent to the cut indicated by the selected area.

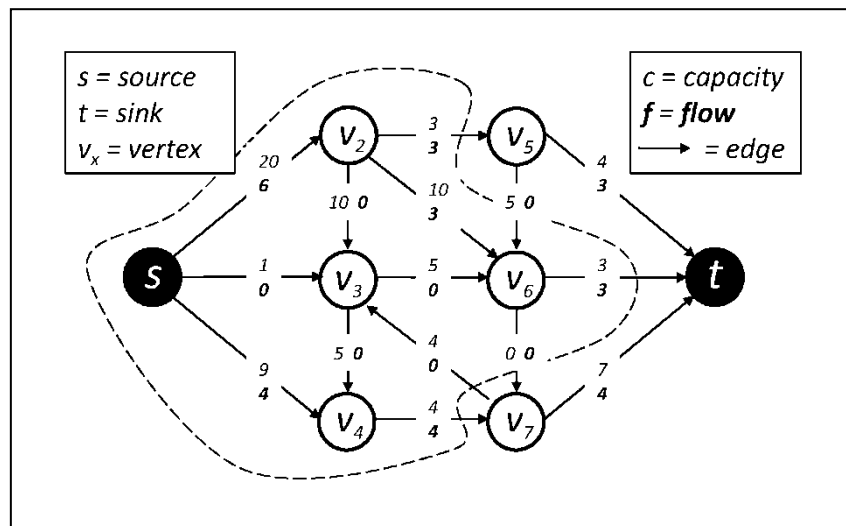


Figure 4 min-cut max-flow theorem applied to an Information Flow Network

In this case too, there are many technical results that could be useful for specific applications, such as the classic Ford-Fulkerson algorithm (Ford and Fulkerson (1956)), which calculates the maximum flow in a network. But the philosophically interesting point here is that we

have an elegant way of stating the equivalent of Ockham's razor for NTA: vertices and edges *non sunt multiplicanda praeter necessitatem*. In Figure 4, it is clear that vertices v_5 and v_7 are redundant in terms of capacity. The size and shape of the informational network A required to account for t (answer t 's HC-questions) can then be optimised by searching for the minimal set of vertices whose informational capacity equates the needed flow. Put simply, we should not look for more potential answers than we need.

Once the flow in A is optimised using (Ockham's) max-flow min-cut theorem, two other interesting features become obvious. First, networks can be interlocked like Lego bricks: what for one analysis is a sink can easily be a target for a different analysis, and this holds true for vertices as well. The building of larger networks is modular and perfectly feasible, like the scaling up of the features of the obtained complex network, in which conditions (a)-(j) may be variously satisfied. Depending on which theory of truth one adopts to analyse semantic information, the successful interlocking of a boundless set of accounts could actually be all that is required to make sense and support a coherentist approach.

Second, different accounts of the same semantic information posing the same HC-questions are perfectly possible, because NTA allows the constructions of different networks A_x that guarantee the same flow of information. Yet this does not commit NTA to some relativism: not every information network works and some informational networks work better than others. This is obvious in network theory, where the max-flow min-cut theorem can be used to prove that there may be more than one flow which attains the maximum amount, and more than one cut that attains the minimal weight. The same result seems to apply uncontroversially to NTA as well.

As for the actual nature of the elements constituting a simple network, the application of the max-flow min-cut theorem can also help us to understand the phenomenon of testimony in the following way. Imagine we wish to account for some semantic information t which actually refers to s ; for example, we wish to account for $t =$ "Julius Caesar was assassinated in Rome on March 15, 44 BC" through an informational network A whose source s is ultimately some first-hand information about the set of events that occurred in Rome on March 15, 44 BC. All the vertices v_x between s and t are testimonies. We can now define the concept of *perfect testimony* very precisely: an ideal scenario for a historian would be to have all the answers that she might wish to be provided by s (that is, s 's capacity c) wholly preserved through the network, but this can happen only if there is no bottleneck narrower than s ' c . In others words, the informational flow reaching t should be equal to the informational capacity of s , so that the min-cut should include only s . In order to evaluate a

group of testimonies, we should set the capacity of s and then check whether the min-cut includes any v_x . For any v_x that gets caught in the min-cut is telling less than the truth, the whole truth and nothing but the truth.

The previous features should have given the general flavour of how NTA works and of its scope of applicability. The last aspect to which I wish to call the reader's attention in this section is the special relation that NTA establishes between s and t in A , as this will play a key role in the next section.

It is plausible to assume, although this assumption will not be further developed here, that the logic of HC-questions may be easily rendered by adapting that of why-questions (Bromberger (1966), Koura (1988), Burgin and Kuznetsov (1994)). What is of the utmost importance here is that, within such logical analysis, one feature is both uncontroversial and crucial. It can be introduced by referring to some familiar instances of accounting. Suppose a theory explains a particular set of phenomena: for example, it explain why metals expand when heated, by indicating that, as temperature increases, the kinetic energy of the atoms in the metal increases as well, this leads to a greater average separation, hence to an expansion of the internal structure of the metal and so of its volume. Although this explanation might be expressed in more ("crystalline lattice") or less ("wobbling") technical terms, thus generating different networks with the same sink t , there is one aspect that remains invariant under different formulations and in each issuing network: *explanans* and *explanandum* cannot be de-coupled without making the explanation incorrect. Or put more simply: an explanation of p is correct if and only if it applies correctly to p . Hopefully the reader will find this uncontroversial to the point of triviality, but note the biconditional and let me stress the importance of the sufficiency condition. Consider now a second example, equally unproblematic: suppose a conclusion is validly deduced from a set of premises. We know that classical formal validity is monotonic: if the deduction is valid it remains valid, and the conclusion remains validly deduced from the premises, no matter how many other premises are added. The only way of de-coupling the conclusion from the premises is by showing that the deduction itself was not valid in the first place. Finally, take any ordinary way of accounting for some semantic information, such as the example of "the water in the electric kettle is boiling": it seems clear that, in all the cases we have encountered so far, either an account works (is correct) or it does not (is not correct), but that it makes no sense to talk of an account that is correct and yet fails to apply for some unknown reasons. Explanations, deductions, teleological or functional "giving reasons" and other forms of accounting all appear to share this essential feature: accounting is rigid, so A is a correct account of t if and

only if in A s correctly accounts for t and they cannot be de-coupled without revising the initial assumption that s did account for t correctly and hence that A was a correct account of it. Visually, t and s may be seen as the two end, bounding points of a line segment represented by the information flow: they are part of it as they define it. I shall refer to this characteristic as the *rigid coordination* between the accounting s and the accounted t in A .

One might object that we are not entirely sure that such rigid coordination does extend to all forms of accounting covered by NTA. After all, the list of types of accounting was left open, so how could one be so confident? This reluctance is understandable and might be overcome not so much by pointing out that the burden of showing that some forms of accounting are not rigid is on the side of the sceptic, but, more easily, by agreeing that, should indeed some forms of accounting turn out to be non-rigid, it will be sufficient to accept that, nevertheless, some major, important forms of accounting, such as the ones we have discussed in this article, are rigid, and that these are the only ones to which we are referring.

5. The Network Theory of Account and the Gettier Problem

So far, I have argued that (an interpreted) A provides the *necessary* condition to upgrade semantic information to knowledge. The time has come to deal with a difficult question: does (an interpreted) A also provide the *sufficient* condition to upgrade semantic information to knowledge? The alert reader may spot here the ghost of a Gettier-type problem, with which we may as well deal openly by rephrasing the problem thus: is the analysis of knowledge as accounted semantic information Gettierisable? The short answer is no, the long answer requires some toil.

To begin with, it is important to clarify that Gettier-type problems are logically unsolvable because they are a sub-class of the more general “coordinated attack” problem, which is demonstrably unsolvable in epistemic logic (Floridi (2006)). The difficulty at the root of this mess is that the tripartite definition presupposes the possibility of coordinating two resources, the objective truth of p and the subjective justification of the knowing subject S , which, by hypothesis, can always be de-coupled. There is a *potential lack of successful coordination*, between the truth of p and the reasons that justify S in holding that p , that is inherently ineradicable. So a Gettier-type counterexample can always arise because the truth and the justification of p happen to be not only mutually independent (as they should be, since we are dealing with fallibilist knowledge) but may also be opaquely unrelated, that is, they may happen to fail to converge or to agree on the same propositional content p in a

relevant and significant way, without S realising it (Gettierization). All this entails that the tripartite definition of knowledge is not merely inadequate as it stands, as proved by Gettier-type counterexamples, but demonstrably irreparable in principle because of the constraints it sets up, so that efforts to improve it can never succeed. With an analogy, the problem is not that one cannot square the circle, but that squaring the circle with algebraic means (straight-edge and compass) is impossible.

Given such a disheartening conclusion, one is entitled to interpret Gettier-type problems as symptoms of a bankrupt approach. The assumption remains, however, that in many cases we do enjoy epistemic propositional states: Mary knows that the kettle is boiling. So the constructive strategy consists in breaking away from the constraints that make the problem unsolvable: we no longer try to define knowledge doxastically and by relying on a logic of subjective justification, but informationally, and by using a logic of objective accounting. Of course, the new tools require shaping and sharpening, but that was exactly the task of the previous sections. So we are now ready to reap the fruit of our labour: some semantic information t (which is constitutively true), if correctly accounted by an information flow network A , is rigidly coordinated to the source s that correctly accounts for it, and cannot be de-coupled from it without making A an incorrect account, so it follows that Gettier-type counterexamples cannot arise. In epistemic logic, this is equivalent to saying that the Byzantine Generals (in our case the two resources s and t) do not try to coordinate their attack infallibly, which is impossible (Fagin et al. (1995)), but rather join forces first, and then attack, which is perfectly feasible.

Let us now consider what happens to our Scribe. So far we have employed an extensional approach: packets of semantic information have been treated as conceptual artefacts or, more figuratively, as items in the Scribe's scroll. We can now translate them intentionally, in the following way: a knowing subject (e.g., the Scribe) S knows that t if and only if:

- i) t qualifies as semantic information;
- ii) A accounts for t , that is, $A(A, t)$;
- iii) S is informed that t ; and
- iv) S is informed that $A(A, t)$.

This informational definition of knowledge faces at least one major objection, but, before discussing it, a few essential clarifications are in order.

Clauses (i) and (ii) seem to require no further comments, but clause (iii) is meant to satisfy at least the information logic based on **B**, if not some higher epistemic logic, and this leads to a first clarification. Depending on whether we assume *S*'s informational states in (iii) and (iv) to be non-opaque – that is, depending on whether *S* not only holds the information, but also understands that *t* and that $A(A, t)$ – we may be able (or fail) to include current artificial agents among the class of epistemic subjects. Since at least 2005 (First International Symposium on Explanation-aware Computing), there has been increasing interest in so-called explanation-aware computing (ExaCt) and more results have become available in recent years. However, it is important to stress that the sort of explanatory processes in question in ExaCt are not the ones that have been discussed here. The goal is to develop ways in which artificial expert systems may interact more profitably with human users, and hence increase their rate of success at solving problems collaboratively by “explaining” their operations and making their procedures and results more accessible. So we should be rather cautious: extensionally, knowledge is accounted information, and this is why we say that a scientific textbook or a website like Wikipedia, for example, contains knowledge, not just information. However, intentionally it seems that knowing requires understanding, or at least that the two are mutually related, and hence that current artificial agents cannot qualify as knowing subjects. They may hold knowledge extensionally, but they cannot know intentionally. This of course says nothing about futuristic artefacts that, should AI ever become possible, would be welcome to join us.

A second, apparent restriction comes with the more or less explicit holding not just of an informational content *t*, but also of a satisfactory account for it. It seems clear that animals do not hold explicit accounts for their information, so it follows that even the smartest dog can at most be informed that the neighbour's cat is a nasty beast, and yet not know it. However, animals do not hold justifications for their beliefs either, but when we acknowledge the old, doxastic, tripartite definition to be more inclusive, we mean that, as observers, it allows us to attribute to animals justificatory grounds supporting their beliefs implicitly. But if this is the case, then the same stance can be adopted in the case of holding an account. The dog knows that the neighbour's cat is a nasty beast because we may attribute to it the (at least implicit) memory of the historical account, for example, of the events that led to such belief. Animals do not hold *explicit accounts* for their information but it seems unproblematic to attribute to them both reasonable levels of understanding (contrary to engineered artefacts) and *implicit accounts* of their information, and therefore knowledge.

A third restriction concerns human knowing subjects. It is an advantage of the informational analysis of knowledge over the doxastic one that the former but not the latter allows for a graded evaluation of epistemic states. This is an important requirement. The doxastic approach is binary: either the Scribe knows that t or he does not, and if he does, his knowledge would be as good as that of an omniscient subject. This is simplistic and the informational approach redresses the situation by making the acknowledgement of expertise possible: the Scribe might know that p better than his dog does because he can provide an account for it, not just hold an implicit account of it. However, a scientist or an historian, for example, might know that p better than the Scribe. This because it is possible to agree on a minimal range of HC-questions that need to be answered correctly in order to qualify as a knowing subject – this is what we ordinarily do in educational and training contexts – but of course there is a vast number of further HC-questions that only an expert will be able to answer. Mary may know that her TV is not working properly because she is well informed about it and what accounts for it, but only the expert will have the right level of advanced knowledge to answer further HC-questions. Knowledge comes in degrees, and insipience as well as omniscience are not only a matter of scope – as we have seen above when discussing the possibility of not upgradeable information – but also of depth.

The profile of a knowing subject that emerges from the informational analysis of knowledge is, unsurprisingly, rather Greek. One important difference, however, is that the analysis links propositional knowledge to practical knowledge (know-that to know-how) in a way that Plato and Aristotle might have found less congenial, but might have pleased Bacon and Kant for being closer to their constructionist approach to knowledge. For it seems clear that knowing that t relies on knowing how to build, articulate and defend a correct account for t . Yet this is often acknowledged in Greek epistemology only partly and somewhat reluctantly, not in terms of ability to manufacture the required conceptual artefact, but merely in terms of ability to convey its properties. In Plato, it is the user that is considered to know something better than the artisan that has produced it. The informational analysis of knowledge is more engineer-friendly: according to it, the production of knowledge that t relies, ultimately, on the intelligent mastery of the practical expertise (including modelling or, more mundanely, story-telling) required to produce not only t but also its correct account A .

The last comment concerns the potential objection anticipated above, to which we can now finally turn. One may contend that the informational analysis of knowledge merely shifts the de-coupling problem. In the doxastic analysis, this affects the relation between the truth of t and S ' justification for believing in it. In the informational analysis – the objection continues

– the problem merely resurfaces by affecting the relation between the correct account of t and the possibility that S may hold it.

This objection deserves to be taken seriously, not because it represents anything close to a refutation, but because it does highlight a significant difficulty, which is different from what the objection seems to imply, and that can be turned into an advantage. Let me explain.

The objection suggests that we did not really get rid of Gettier-type counterexamples but only moved them out of sight. This is mistaken. The task was to show how semantic information can be upgraded to knowledge and the previous analysis provides the necessary and sufficient conditions to achieve this. The problem left unsolved is not the potential Gettierisation of the informational analysis because – once the logic of accounting replaces that of justification – the condition of possibility of Gettier-type counterexamples (i.e., decoupling) is removed. Nonetheless, the objection is correct in raising the more generic suspicion that something has been left unsolved. For the trouble is that the informational analysis converts Gettier-type problems into sceptical ones. How can S be certain that A is the correct account of t ? This is not among the questions answered by any account of t . Indeed, it must be acknowledged that nothing has been said in this article that goes towards tackling this sceptical question. But then, nothing should, because this is not the challenge we had set out to address. Of course, one may find this unsatisfactory: we are jumping out of Gettier’s frying pan only to land into the sceptic’s fire. Yet such dissatisfaction would be ungenerous. The sceptical challenge concerns the truth of t and, broadly speaking, the correctness of an account A of t (or of the answers offered with respect to the HC-questions posed by t) and S ’ possibility of not being mistaken about holding A . But such a challenge was always going to affect any analysis of knowledge, including the doxastic one. So, by converting Gettier problems into sceptical problems we have made progress, because the latter problems are not made any more serious by such conversion and we now need to take care of only one set of difficulties instead of two. Fighting on only one front is always preferable and it is an improvement. Von Clausewitz *docet*.

6. Conclusion

Knowledge and information are members of the same conceptual family. What the former enjoys and the latter lacks, over and above their family resemblance, is the web of mutual relations that allow one part of it to account for another. Shatter that, and you are left with a pile of truths or a random list of bits of information that cannot help to make sense of the reality they seek to address. Reconstruct that network of relations, and information starts

providing that overall view of the world which we associate with the best of our epistemic efforts. In this paper, I have tried to show how this upgrading of semantic information to knowledge is possible, welcome and resilient, by relying on some results from network theory. The proposal, if successful, is not without its costs. Exchanging Gettier-type counterexamples, which affect the justification of *S*, for doubts about whether *S* might hold the correct account of his information, means reverting to the challenge posed by the sceptic. It is a serious challenge that will have to met, but, fortunately, not in this context.¹⁴

¹⁴ I would like to acknowledge the useful comments and criticisms by Patrick Allo and Sebastian Sequoiah-Grayson on a first draft of this paper. Penny Driscoll kindly copyedited the last version. The two anonymous reviewers of the journal provided many useful comments. All the aforementioned people helped me to improve the paper substantially but they are not responsible for any remaining mistakes.

References

- Achinstein, P. 1983, *The Nature of Explanation* (New York; London: Oxford University Press).
- Ahlsvede, R., Cai, N., Li, S.-Y. R., and Wai-Ho Yeung, R. 2000, "Network Information Flow", *IEEE Transactions on Information Theory*, 46(4), 1204 - 1216.
- Ahuja, R. K., Magnanti, T. L., and Orlin, J. B. 1993, *Network Flows : Theory, Algorithms, and Applications* (Englewood Cliffs, N.J: Prentice Hall).
- Bar-Hillel, Y., and Carnap, R. 1953, "An Outline of a Theory of Semantic Information" repr. in Bar-Hillel [1964], pp. 221-74.
- BonJour, L. 1985, *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press).
- Bromberger, S. 1966, "Why-Questions" in *Readings in the Philosophy of Science*, edited by Baruch A. Brody (Englewood Cliffs: Prentice Hall, Inc.). 66-84.
- Burgin, M., and Kuznetsov, V. 1994, "Scientific Problems and Questions from a Logical Point of View", *Synthese*, 100(1), 1-28.
- Chellas, B. F. 1980, *Modal Logic : An Introduction* (Cambridge: Cambridge University Press).
- Chen, W.-K. 2003, *Net Theory and Its Applications: Flows in Networks* (London: Imperial College Press).
- Colburn, T. R. 2000a, "Information, Thought, and Knowledge", *Proceedings of the World Multiconference on Systemics, Cybernetics and Informatics*, 467-471.
- Colburn, T. R. 2000b, *Philosophy and Computer Science* (Armonk, N.Y.: M.E. Sharpe).
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. 2001, *Introduction to Algorithms* 2nd ed. (Cambridge, Mass.; London: MIT Press).
- Devlin, K. J. 1991, *Logic and Information* (Cambridge: Cambridge University Press).
- Dodig-Crnkovic, G. 2005, "System Modeling and Information Semantics", *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).
- Dretske, F. I. 1988, *Explaining Behavior: Reasons in a World of Causes* (Cambridge, Mass.: MIT Press).
- Dummett, M. A. E. 2004, *Truth and the Past* (New York; Chichester: Columbia University Press).
- Elias, P., Feinstein, A., and Shannon, C. E. 1956, "Note on Maximum Flow through a Network", *IRE Transactions on Information Theory*, IT-2, 117-119.
- Fagin, R., Halpern, J. Y., Moses, Y., and Vardi, M. Y. 1995, *Reasoning About Knowledge* (Cambridge, Mass ; London: MIT Press).
- Fetzer, J. H. 2004, "Information: Does It Have to Be True?", *Minds and Machines*, 14(2), 223-229.
- 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, "Open Problems in the Philosophy of Information", *Metaphilosophy*, 35(4), 554-582.
- Floridi, L. 2004c, "Outline of a Theory of Strongly Semantic Information", *Minds and Machines*, 14(2), 197-222.
- Floridi, L. 2005, "Is Information Meaningful Data?", *Philosophy and Phenomenological Research*, 70(2), 351-370.
- Floridi, L. 2006, "The Logic of Being Informed", *Logique et Analyse*, 49(196), 433-460.

- Floridi, L. 2007, "In Defence of the Veridical Nature of Semantic Information", *The European Journal of Analytic Philosophy*, 3(1), 1-18.
- Floridi, L. 2008a, "Data" in *International Encyclopedia of the Social Sciences*, edited by William A. Darity (Detroit: Macmillan).
- Floridi, L. 2008b, "Understanding Epistemic Relevance", *Erkenntnis*, 69(1), 69-92.
- Floridi, L. 2010a, *Information - a Very Short Introduction* (Oxford: Oxford University Press).
- Floridi, L. 2010b, *The Philosophy of Information* (Oxford: Oxford University Press).
- Floridi, L. forthcoming, "Semantic Information and the Correctness Theory of Truth".
- Ford, L. R., and Fulkerson, D. R. 1956, "Maximal Flow through a Network", *Canadian Journal of Mathematics*, 8, 399-404.
- Grice, H. P. 1989, *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press).
- Jungnickel, D. 1999, *Graphs, Networks, and Algorithms* (Berlin; New York: Springer).
- Koura, A. 1988, "An Approach to Why-Questions", *Synthese*, 74(2), 191-206.
- Mayr, E. 1961, "Cause and Effect in Biology", *Science*, 134 1501-1506. Also in E. Mayr, *Toward a New Philosophy of Biology: Observations of an Evolutionist* (Cambridge, MA: Harvard University Press, 1988) and in E. Mayr (ed.), *Evolution and the Diversity of Life* (Cambridge, MA: Harvard University Press, 1997).
- Newman, M. E. J., Barabási, A.-L., and Watts, D. J. (ed.) 2006, *The Structure and Dynamics of Networks* (Princeton, N.J.: Princeton University Press).
- Sequoiah-Grayson, S. 2007, "The Metaphilosophy of Information", *Minds and Machines*, 17(3), 331-344.
- Stering, R. 2008, *Police Officer's Handbook : An Analytical and Administrative Guide* (Sudbury, Mass.: Jones and Bartlett Publishers).
- Taylor, C. C. W. 1967, "Plato and the Mathematicians: An Examination of Professor Hare's Views", *The Philosophical Quarterly*, 17(68), 193-203.
- Taylor, C. C. W. 2008, "Plato's Epistemology" in *The Oxford Handbook of Plato*, edited by Gail Fine (New York; Oxford: Oxford University Press). 165-190.
- Van Fraassen, B. C. 1980, *The Scientific Image* (Oxford: Clarendon Press).
- Walton, D. 2007, "Dialogical Models of Explanation" in *Explanation-Aware Computing: Papers from the 2007 Aaai Workshop, Association for the Advancement of Artificial Intelligence*, (AAAI Press). 1-9.
- Whaley, B. B., and Samter, W. (ed.) 2006, *Explaining Communication: Contemporary Theories and Exemplars* (London: Routledge).
- Yeung, R. W. 2008, *Information Theory and Network Coding* (New York: Springer).