

Semantic Information and the Correctness Theory of Truth

Luciano Floridi

Received: 3 January 2009 / Accepted: 20 August 2010 / Published online: 29 September 2010
© Springer Science+Business Media B.V. 2010

Abstract Semantic information is usually supposed to satisfy the veridicality thesis: p qualifies as semantic information only if p is true. However, what it means for semantic information to be true is often left implicit, with correspondentist interpretations representing the most popular, default option. The article develops an alternative approach, namely a correctness theory of truth (CTT) for semantic information. This is meant as a contribution not only to the philosophy of information but also to the philosophical debate on the nature of truth. After the introduction, in Sect. 2, semantic information is shown to be translatable into propositional semantic information (i). In Sect. 3, i is polarised into a query (Q) and a result (R), qualified by a specific context, a level of abstraction and a purpose. This polarization is normalised in Sect. 4, where $[Q + R]$ is transformed into a Boolean question and its relative yes/no answer $[Q + A]$. This completes the reduction of the truth of i to the correctness of A . In Sects. 5 and 6, it is argued that (1) A is the correct answer to Q if and only if (2) A correctly *saturates* Q by *verifying* and *validating* it (in the computer science’s sense of “verification” and “validation”); that (2) is the case if and only if (3) $[Q + A]$ generates an adequate model (m) of the relevant system (s) identified by Q ; that (3) is the case if and only if (4) m is a *proxy* of s (in the computer science’s sense of “proxy”) and (5) proximal access to m commutes with the distal access to s (in the category theory’s sense of “commutation”); and that (5) is the case if and only if (6) reading/writing (*accessing*, in the computer science’s technical sense of the term) m enables one to read/write (access) s . Sect. 7 provides some further clarifications about CTT, in the light of semantic paradoxes. Section 8 draws a general conclusion about the nature of CTT

L. Floridi (✉)

Department of Philosophy, University of Hertfordshire, de Havilland Campus, Hatfield,
Hertfordshire AL10 9AB, UK
e-mail: l.floridi@herts.ac.uk

L. Floridi

Faculty of Philosophy and IEG, University of Oxford, Oxford, UK

as a theory for systems designers not just systems users. In the course of the article all technical expressions from computer science are explained.

1 Introduction

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. The debate is somewhat old,² but has been re-ignited by the proposal to analyse semantic information in terms of *well-formed*, *meaningful* and *veridical data* (Floridi 2004b). Admittedly, the analysis—according to which semantic information encapsulates truth, exactly as knowledge does—has attracted some criticisms for being too restrictive.³ Such criticisms, however, have been proved unjustified.⁴ 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*. I shall discuss the formal nature of content in the following pages but, at the moment, suffice it to say that it is analysable 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 and 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, 1988), one may easily recognise this “aboutness” feature in propositional attitudes such as “Mary is informed that there is some beer in the fridge”, where “being informed” is used in the *statal*⁶ sense, i.e., in the sense that Mary holds that information. This is the condition into which *a* enters once *a* has *acquired* the information (*actional* state of being informed) that *p*. It is the sense in which a witness, for example, is informed (holds the information) that the suspect was with her at the time when the crime was

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

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), with a reply in Floridi (2005); or the objections moved by Colburn (2000a, b) and Dodig-Crnkovic (2005).

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

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

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.

committed. In the rest of this paper, we shall be concerned with only this standard, statal and epistemically oriented concept of semantic information.

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. 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.

[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 2008c). Here, it is worth emphasising that it forges a robust and intuitive link between semantic information and knowledge. More specifically, the veridical thesis—the condition in [DEF] that the data need to be veridical—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 $\Box \varphi \rightarrow \varphi$, also known as **T**, **M** or **K2**, thus:

[VT] $I_a p \rightarrow p$

The intended interpretation of [VT] is that a is informed 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, 2010b). One is upstream:

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

The other is downstream:

(b) how can semantic information turn into knowledge?

Both questions are prompted by [DEF] but neither is specifically about [DEF] only, so 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 philosophical 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 (a), let me brush (b) away by adding a final clarification.⁷

[DEF] nests semantic information into knowledge so tightly that one is naturally led to wonder whether anything else might be missing, in order to escalate 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 there is some beer in the fridge” and “Mary knows that there is some beer in the fridge” as if they were always losslessly interchangeable. 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. Philosophically, however, the distinction captures an important difference, and hence it is important to be more accurate. For 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 $I_{ap} \rightarrow II_{ap}$, one may also object that $K_{ap} \rightarrow KK_{ap}$ is notoriously controversial as well) but also because, even when it is, it might still arguably be *opaque* and certainly *unjustified*.

Consider opacity first. It is open to discussion whether a messenger carrying (whether in her memory or in her hand 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 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 comprehend 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 .

Next, consider the degree of justification. Epistemic luck does not affect informativeness negatively. 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 at $t - 12$ h and therefore happens to indicate the right time $t - 12$ at t , one still holds the information that the time is t , although one can no longer be said to know the

⁷ The interested reader is referred to Floridi (forthcoming), the “twin article” where I develop and defend a full answer to question (b).

time. The same applies to a more Platonic example in which a student memorises, but fails to understand, 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. The additional difficulty is that the counterexamples show both that some justificatory variable might have a key role to play in full epistemic states, besides reflectivity and transparency, and that this variable too is still insufficient to guarantee the delivery of knowledge every time. Sometimes, one may be (reflectively and transparently) informed that p and fully justified in holding the information that p and yet still fail to know that p .

Rotten as all this may be, it is not all, for there is further bad news. One can also prove that Gettier-type problems are logically unsolvable by showing that they are a sub-class of the more general “coordinated attack” problem, which is demonstrably insolvable in epistemic logic (Floridi 2004b). This entails that the tripartite account is not merely inadequate as it stands, as proved by Gettier-type counterexamples, but demonstrably irreparable in principle, so that efforts to improve it can never succeed. Although it is useful to know that we should stop trying to fix this approach and start looking for a different one, the disappointing conclusion is that, as far as question (b) above is concerned, we lack even a promising strategy to upgrade I_{ap} to K_{ap} .

So much regarding (b), which we can now leave on one side. Prospects are much brighter when it comes to question (a). In this case, the challenge is not a shortage, but rather an overabundance of viable answers, since we are spoiled for choice by a variety of theories of truth.⁸ Admittedly, in the literature on semantic information there appears to be at least an implicit predilection for some version of a Tarskian and/or correspondentist approach.⁹ And yet, at least in principle, nothing prevents each of the major theories of truth from answering (a). They simply would have been refuted a long time ago if they couldn't. It follows that some initial tolerance towards a pluralistic approach to (a) might be unavoidable, if not methodologically welcome. Of course, if this were all that one could sensibly recommend about (a), there would be little reason to pursue any further investigation. There is, however, another way of approaching (a), which opens up an interesting line of enquiry, that further expands the menu of viable philosophical theories of truth.

Consider the strategy sketched above. It consists in selecting the best available theory of truth and testing how well it might be applied and adapted in order to explain the truthfulness of semantic information. With some negligible adjustments, such a top-down approach is comparable to the so-called “design pattern” technique (Gamma et al. 1995) in software engineering (Sommerville 2007). This consists in

⁸ In this paper, I have relied especially on Lynch (2001), Engel (2002) and Künne (2003), among the many introductions and anthologies available on the major theories of truth, as particularly helpful.

⁹ See for example Popper (1935), Dretske (1981), Fox (1983), Israel and Perry (1990), Barwise and Seligman (1997) and Bremer and Cohnitz (2004).

identifying and specifying the abstract features of a design structure (e.g. how to build a paying system for a candy vending machine), which are then generally reusable solutions to commonly occurring problems in the construction of an artefact (e.g. the paying system for a drinks vending system). In our case, we have several design patterns for the concept of truth. We know that they are robust, because they have been tested and refined since Ramsey, if not Aristotle. We also know that they are reusable: although they have been developed to deal primarily with propositional or sentential truths, one may reasonably expect them to be effectively adaptable to truthful semantic data (e.g. the map of the London underground) as well. So, when our artefact, i.e. semantic information, is proved to require the particular feature of being truthful, a sensible alternative is to consider such design patterns and try to identify the ones that best satisfy the constraints and requirements imposed by the development of the artefact itself. Oversimplifying, one may answer question (a) above by choosing whichever pre-packaged theory of truth turns out to be most suitable. This strategy may be classic, is certainly viable but it is hardly innovative. I shall not pursue it in the following pages, although I shall return to more standard theories of truth in Sect. 7. The reason for this choice will be clear in a moment.

The other approach is bottom-up and suggests the sort of strategy that will guide the rest of this investigation. It consists in assuming the artefact itself as given, and then trying to discover the principles governing its properties and workings by analysing its structure, function and operations. In software engineering, this technique is known as “reverse engineering”. This is “the process of extracting the knowledge or design blueprints from anything man-made” (Eilam 2005, p. 3). It consists in examining an existing artefact in order to identify its components and their interrelationships and hence create representations of it in other forms or at a higher level of generalization. Imagine reverse engineering the candy vending machine in order to understand how it works and then re-use what you have discovered in order to engineer a new vending machine. Following this strategy, one may answer question (a) by assuming the occurrence of some semantic information and then disassembling it in order to reveal what its components are and how they interact with each other in order to deliver information. We have the artefact and we seek to understand its mechanism by taking it apart, hopefully in the right way and places. Note that this second strategy is perfectly compatible with the first, once it is realised that there is a virtuous cycle of feedback between design patterns and reverse engineering results. Contrary to the first strategy, however, reverse engineering promises to deliver a more innovative analysis, as it avoids approaching the problem of truth from pre-established theories and explores it from a new perspective. After all, the first strategy merely retrofits some already existing theory of truth to semantic information, instead of trying to develop a customised solution which may then be generalisable. By reverse engineering semantic information, the goal is to articulate and support a theory of truth that explains what it means for some semantic information p to be true and hence expand the number of viable options at our disposal when considering which philosophical theories of truth are available. The cost to be paid for this innovation is that our bottom-up strategy will also be uphill, if I may be allowed to combine the two metaphors: it is much more

economical to choose from a pre-established menu than to develop a new approach. I can only hope that the reader will find the effort rewarding and the result enlightening. And now it is time to start climbing.

2 First Step: Translation

A large variety of kinds of semantic information, from traffic lights to train timetables, from road signs to fire alarms, falls within the scope of [DEF]. This is how it should be but it is awfully inconvenient for our purposes. For in order to reverse-engineer semantic information in such a way that its components might easily be identified, disassembled and explained, it would be far easier and more fruitful to concentrate on just one kind, the propositional one, which lends itself to such a treatment straightforwardly. So, our first step will be to ensure that all kinds of semantic information that satisfy the *definiens* in [DEF] are indeed translatable into *propositional* semantic information, thus guaranteeing that what will be concluded about the latter may be extendable to the former. At this point, the reader who finds such “translatability” uncontroversial, or indeed trivial, may wish to skip the rest of this section. The one who finds it impossible may concede the restriction of scope as a matter of convenient stipulation, although the rest of this section purports to show that the burden of proof is on her shoulders. As for the rest of us, what follows should be sufficiently convincing to make our second step unproblematic.

Syntactically (or in terms of information theory), the propositional translatability of any kind of semantic information is unquestionable and a matter of daily experience. After all, analogue information is reproducible digitally to any chosen degree of accuracy, its digital version is equivalent to finite lists of zeros and ones, and these can be further encoded into as many answers to questions asked in a suitably chosen language, and hence ultimately translated into statements of that language. That doing any of this would be sheer madness is irrelevant here. For the question is not how difficult or costly this process would be, e.g. in terms of accuracy, time and memory resources, but that it might be possible at all. More to the point is whether some non-propositional, semantic information—the sort of information provided by the map of the London Underground, for example—may always be translatable *semantically* into propositional semantic information, at least in principle. Mind, not all of it at once, and not even part of it at every level and for every kind and degree of detail [henceforth Level of Abstraction or simply LoA, Floridi (2008b)], but any of it at the right LoA, depending on needs and requirements. Since the difference between a syntactic and a semantic translation may not be very familiar, let me first introduce it with an example.

Consider being able to reproduce the map of the London Underground on graph paper by being told, say over the phone, the position and colour of each square on the paper: the communication over the phone would provide a syntactic translation, with the end result (the coloured graph paper representing the map) constituting a test about whether the translation worked. Contrast it now to being able to travel from one station to another on the London Underground, by receiving verbal

instructions from someone who is navigating using the visual indications provided by the map. This is a semantic translation, and your trip is a test of its accuracy.

Suppose now that a semantic translation from non-propositional into propositional information, of the kind just illustrated, were sometimes impossible, even in principle. Then there would be some residual semantic information, conveyed non-propositionally (e.g., by the map), that one would necessarily be unable to convey propositionally, independently of the resources available. We would then have reached the limits of the informational powers of any natural language, even natural languages formally expanded, e.g. mathematically. Allegedly, we should still be able to point to the information in question (in the previous example, suppose we are both looking at the same map), but we would be unable to generate the right sort of propositional content that could adequately convey it. This is a *reductio ad absurdum*. For here we are not engaging with some Wittgensteinian limits of the “sayable”, with Kantian *noumena*, with some linguistically-ungraspable sensations, or some mystical experience enjoyed while looking at the map of the London Underground. We are talking about what the map of the London Underground can encode, in terms of information about travelling through the network, positions of the stations, interconnections, available routes etc., which, allegedly, would be at least partly beyond the expressive power of any natural language to convey. But since natural languages have been acknowledged to be “semantically omnipotent” at least since Leibniz (Formigari 2004, pp. 91–92), one can arguably assume that the translation is always possible, even if it is likely to be onerous at times and hence often unfeasible in terms of resources. So, in the rest of the paper, we shall treat semantic information as possibly semiotic-dependent (it may always require a code) but not as semiotically bounded (codes are translatable propositionally, if expensively resource-wise). The same point can be expressed formally and succinctly thus:

$$[\text{Tr}] \quad \forall x (\text{DEF}(x) \wedge \text{Non-prop}(x)) \rightarrow \exists y (\text{Prop-}t(y, x) \wedge \text{DEF}(y))$$

The intended interpretation of [Tr] is that, if any data (the domain on which the quantifiers range) satisfy [DEF] but are not propositional, then there is a propositional translation of those data which also satisfies [DEF]. Note that we do not need to assume the stronger principle of translational equivalence: pictures may be worth thousands of words, but there might be thousands of words that are priceless. All that [Tr] needs to guarantee is that the conclusions reached about the alethic nature of propositional semantic information will be exportable to the truthful nature of non-propositional semantic information as well. In other words, that what can be concluded about the truth of “there is some beer in the fridge” is equally applicable to the truthfulness of a picture conveying the same information visually.

3 Second Step: Polarization

Once some information *i* is formulated propositionally, the second step is to follow a standard approach, in information theory, to the quantification of information, and

disassemble i into a combination of a query Q and a result R .¹⁰ In short, we have (the asterisk is a reminder that the formula is provisional and will have to be refined):

$$[\text{POL}^*] \quad i = Q + R$$

That $[\text{POL}^*]$ is always achievable is warranted by the fact that any propositional i is equivalent to a message, and that any message is a combination of querying and resulting data encoded in the same set of symbols of the chosen language.¹¹ The polarization of i into $Q + R$ offers several advantages. We shall exploit four of them.

First, $[\text{POL}^*]$ highlights the need to specify the *context* (C) in which, the *level of abstraction* (LoA) at which, and the *purpose* (P) for which the query is formulated and hence it is expected to be satisfied by the result. For the sake of simplicity, below I shall refer to the combination of these three parameters by means of the acronym CLP. The first two requirements were stressed by Austin (1950). “Where is the beer?” is asked by someone in some specific circumstance (the context), by relying on a specific granularity of discourse or detail, what I have called LoA. In our example, there might be no beer (if no beer has been purchased) or, if the sender of the query knows that some beer has been purchased, answering that “the beer is somewhere” would amount to a joke or a mistake in the choice of LoA, if the sender wishes to know the precise location of the beer, e.g. left in the car or carried inside the house or placed in the fridge. The third requirement was stressed by Strawson (1964). LoAs are always teleological and queries are formulated (results are offered) for a purpose, even if the purpose might be implicit. In the example, one may wish to make sure that the beer has been placed in the fridge and not left in the car, for example. Queries cannot acquire their specific meaning in isolation or independently of their CLP parameters. It is a bit of a pain, but we need to keep these variables in mind, lest the conceptual mess becomes unmanageable. So, as a memory aid, let me revise $[\text{POL}^*]$ by adding a combined index, thus:

$$[\text{POL}] \quad i^{\text{CLP}} = [Q + R]^{\text{CLP}}$$


A second advantage of the polarization of i into $Q + R$ is that it makes evident the role of R , which is to *saturate* Q . Although it is trivial to apply $[\text{POL}]$ to any piece of information, p , like “the beer is in the fridge”, in order to obtain:

[Ex. 1]	Query	“Where is the beer?” +
	Result	“In the fridge” =
	Information	“The beer is in the fridge”

¹⁰ A query is to be understood as a request for data sent (e.g., an illocutionary act performed) by a sender to a receiver, in the form of a message. Thus, it might have the format of a question (“where is the beer?”) as well as of an imperative (“tell me where the beer is”), or a string of symbols in a search engine. A result is also to be understood as a message, the requested data, sent by the receiver to the querying sender.

¹¹ Alternatively: every p can be transformed into a request of whether p plus a result, but more on this in the next section.

Fig. 1 The relation “is correctly saturated by” assigns to each query Q in A at least one result R in B



it is important to keep in mind that the correct interpretation of $i^{\text{CLP}} = [Q + R]^{\text{CLP}}$ in [POL] is not as (i) a request for *confirmation* or (ii) a *test*, but as (iii) a genuine request to erase a data deficit through *saturation*. The difference is that, in (i) and (ii), the sender of the query already holds the information that p , but wishes to double-check it, or to check whether the receiver also holds that information; whereas in (iii), the sender lacks the information that p and wishes to acquire it from the receiver by obtaining the missing data. Having said this, let me hasten to clarify a point that might be a source of potential confusion. The polarization of i does *not* really involve two agents. I shall speak sometimes as if the querying sender and the saturating receiver were two different entities, but this is only for heuristic purposes and ease of treatment. It is i that is being polarised, so sender and receiver are really the same entity. If you need an intuitive representation, imagine a language in which Mary can make statements not by uttering declarative sentences, but only by formulating questions followed by the appropriate answers. Her language does not enable her to say: “The beer is in the fridge” but only “Where is the beer? In the fridge”.

The third advantage is set-theoretic. Adopting a standard extensional theory of questions,¹² it is easy to see that [POL] allows us to treat “is correctly saturated by” as a relation r from a countable set of queries $A = \{Q \mid Q \in A\}$ to a countable set of results $B = \{R \mid R \in B\}$. Note that r is not yet a function because two or more propositional i , e.g. “the beer is in the fridge” and “the beer is in the kitchen” are analysed as “where is the beer?” + “in the fridge” and “where is the beer?” + “in the kitchen”, thus mapping the same Q_1 both to R_1 and to R_2 (see Fig. 1). In Sect. 6, we shall see that the real crux is to provide an analysis of correctness that does not beg the question.

The fourth advantage is that [POL] can be normalized. This is our next step.

4 Third Step: Normalization

In real life, queries and results share, in variable proportions, the amount of semantic content that is to be found in the corresponding semantic information. In [Ex. 1], the full semantic content to be found in “the beer is in the fridge” is allocated partly to Q , which contains a request for location and a reference to the

¹² This is a rather standard approach, see Groenendijk and Stokhof (1994) and Szabolcsi (1997).

object to be located, and partly to R , which contains a reference to the requested location of the object to be located. Although a step forward in the disassembling process, this is still unsatisfactory because it makes it very hard to quantify—precisely, consistently and uniformly across the whole class of $Qs + Rs$ —how much content is allocated to which side of the polarised information. In order to uncover what lies under the thick layer of content, it would be useful to shovel it all on one side, by shifting all the content, still embedded in R , to the left, until R is completely streamlined. At the same time, however, weakening R should not lead to an over-strengthening of Q into a rhetorical question, since a question that requires no answer would be a mere transliteration of i itself and would only defy the purpose. Luckily, a little trick from information theory comes to our rescue: we can reach the right balance, in shifting all the content on the side of the queries, by normalising them into yes/no questions, that is (again the asterisk reminds us that the formula is only a first approximation):

$$[\text{NORM}^*] \quad [Q + R]^{\text{CLP norm}} \Rightarrow [Q_{0/1} + A_{0/1}]^{\text{CLP}}$$

The intended interpretation of $[\text{NORM}^*]$ is that a query Q and a result R , both CLP-parameterised, can be normalised into a Boolean Question Q and a Boolean Answer A (the 0/1 subscripts are there to remind us of their Boolean nature), equally CLP-parameterised. This is very much easier done than said, so let us look at our example again. By applying $[\text{NORM}^*]$ to $[\text{Ex. 1}]$, we obtain:

[Ex. 2]	Question	“Is the beer in the fridge?” +
	Answer	“Yes” =
	Information	“The beer is in the fridge”

Of course, this is not what happens in the real world, where one cannot expect a querying sender to be able always to maximise the content of her questions, for she often lacks much more than just a positive or negative saturation. However, recall that we are disassembling semantic information as a given artefact: all the content is already provided, and hence some idealization, typical of controlled experiments, is perfectly reasonable. Recall also that $[\text{NORM}^*]$ does not really involve two agents. This time, imagine Mary being able to state that the beer is in the fridge only by uttering “is the beer in the fridge? Yes”.

Once again, $[\text{NORM}^*]$ offers several nice advantages for our analysis, three of which will be immediately useful for our next step.

The first advantage is semantic: it is now easy to see that it is really Q and not A that sets the scope of the CLP parameters. A Boolean answer can only endorse the *context* (C) in which, the *level of abstraction* (LoA) at which, and the *purpose* (P) for which the Boolean question is formulated; it can neither change nor challenge them. So we can revise $[\text{NORM}^*]$ thus:

$$[\text{NORM}] \quad [Q + R]^{\text{CLP norm}} \Rightarrow Q_{0/1}^{\text{CLP}} + A_{0/1}$$

The second advantage is set-theoretic: the normalization transforms the relation r “is correctly saturated by” into a function f from a still countable domain of




Fig. 2 The function f (= is correctly saturated by) assigns to each Boolean question Q in A exactly one Boolean answer (either *Yes* or *No*) in B . Note that Q_3 , for example, corresponds to a negative truth, e.g. “the red wine is not in the fridge” in the case in which the fridge does not contain any red wine

Boolean questions $A \{Q \mid Q \in A\}$ to a codomain of only two possible Boolean answers $\{Yes, No\}$. Figure 2 provides a graphical illustration.

Correctness is now a functional concept, but it is still premature to investigate it. At this stage, what matters is that the downsizing of the codomain of the function represents the extensional counterpart of a third, informational advantage: [NORM] shifts all the content in i to Q . We have seen that this re-location of content is what motivates the normalization in the first place. To understand how it works and why it is useful, we need to recall a few other elementary facts in information theory.¹³

As is well-known, given a set of N equiprobable symbols, information theory quantifies the amount of information in a symbol thus:

$$\log_2(N) = \text{bits of information per symbol}$$

It follows that a coin ($N = 2$), by producing a head (h) or tail (t), delivers at most (if it is fair) 1 bit of information, whereas two coins ($N = 4$), deliver at most (again, if they are both fair) 2 bits of information (e.g. $\langle h, t \rangle$), and so forth.

Imagine now a biased coin, which makes obtaining h more likely. The more biased the coin is, the more likely h is, the less information is provided by the answer, the smaller the information deficit becomes, up to the point when, if both sides of the coin are heads, the bias is total, the probability of h is 1, the information conveyed by h is 0 bit and so is the receiver’s information deficit. All this means that, since [NORM] transforms queries into yes/no questions [that can be answered by tossing a coin A with different degrees of bias, the worst scenario is one in which Q corresponds to an information deficit that requires at most 1 bit of information from A to be saturated. However, even a $A_{0/1}$ worth a full bit of information fails to add anything, in terms of semantic content, to what is already contained in Q . It follows that, whatever the specific semantic content in i is, [NORM] shifts it entirely to Q , exactly as we wished.

As a consequence, we now have an intuitive way of defining semantic content as unsaturated information or, more formally:

¹³ See Floridi (2010a) for a more detailed but still introductory presentation.

$$\begin{aligned}
 [\text{CONT}] \quad \text{Content in } i^{\text{CLP}} &= \text{Content in } Q_{0/1}^{\text{CLP}} \\
 &= i^{\text{CLP}} - A_{0/1} \\
 &= i^{\text{CLP}} - n \text{ bit of information, for } n = 0 \text{ or } 1
 \end{aligned}$$

[CONT] is not just interesting in itself but provides a reassuring test, since it is perfectly consistent with a theory of strongly semantic information (Floridi 2004c). In particular, it shows that tautologies and contradictions are pure semantic contents, equally uninformative or, to phrase it differently, that they provide no semantic information about their referents, over and above their contents (in both cases the coin we are tossing has two identical sides, as it were). This is as it should be, so our reverse engineering seems to be proceeding in the right direction.

5 Fourth Step: Verification and Validation

We have now disassembled semantic information into two components. By combining [POL] and [NORM], we obtain:

$$[\text{PN}] \quad i^{\text{CLP}} = Q_{0/1}^{\text{CLP}} + A_{0/1}$$

Let us now scrutinize each component separately.

On the one hand, we have seen that $Q_{0/1}$ sets the CLP parameters. Since it provides all the content in i , $Q_{0/1}$ also identifies its referent, that is, what i is about. We can express all this more precisely by saying that $Q_{0/1}^{\text{CLP}}$ identifies a *system* s (the referent of i) and provides all the *semantic content* (the content in i) for a *model* of s (namely, $Q_{0/1}^{\text{CLP}} + A_{0/1}$) within a given context, at a particular LoA and for a purpose.

On the other hand, although $Q_{0/1}^{\text{CLP}}$ in [PN] is still neither a test nor a request for confirmation but a request for saturation, clearly the sort of saturation in question can no longer be a matter of content, as it was in [POL]. $A_{0/1}$ acts only as a Boolean key, that either fails to apply at all (see $\neg A_{0/1}$ in Fig. 3) or that applies and then either locks or unlocks the content provided by $Q_{0/1}^{\text{CLP}}$, thus generating a partial model (henceforth just model) of the targeted system. Once again, a conceptual distinction and some technical vocabulary from software engineering (Fox 2007) can help to clarify this crucial point.

Software Verification and Validation (V&V) is the overall process of checking the “fitness for purpose” of an artefact, by ensuring that the software being developed or modified:

- (a) complies with some given *specifications*, regulations or pre-conditions imposed at the start of the development process; and
- (b) accomplishes its intended purpose, meeting its *requirements*.

The two phases are complementary.

In phase (a), called *verification* (no relation with the philosophical concept), we check whether the artefact is being developed in the right way, that is, whether we

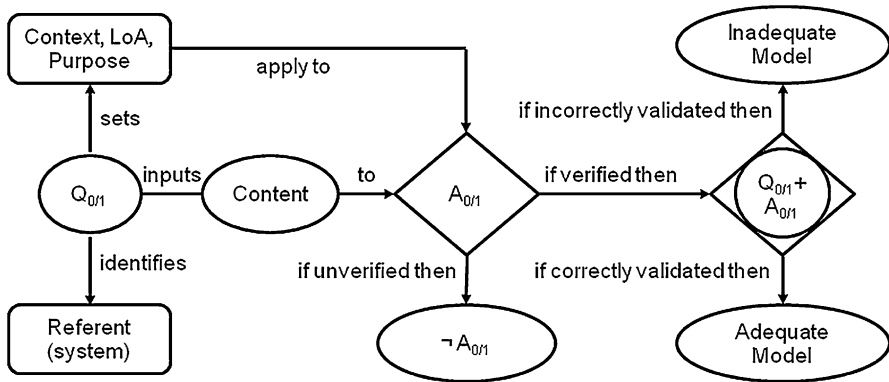


Fig. 3 Summary of the first four steps in the analysis of semantic information. The process starts with $Q_{0/1}$ on the left

are constructing (or have constructed) what we have (or had) planned to construct. In phase (b), known as *validation* (again, no relation with the logical concept either), we check whether the right artefact is being developed, that is, we check whether we are constructing what is required. The V&V process applies to a variety of artefacts and products and helps to clarify the twofold role played by $A_{0/1}$ in [PN]. Let me first show how by relying on our example [Ex. 2].

Given the question “is the beer in the fridge?”, any Boolean answer— independently of whether it is “yes” or “no”—implicitly verifies (in the V&V sense) that the question complies with the pre-conditions (i.e., the specifications) regulating its proper formulation, including its context, LoA and purpose. A question like “Is the fridge in the beer?” fails to qualify as something that can receive either a “yes” or a “no” answer because it fails the verification check, since it blatantly fails to develop the semantic artefact in the right way. Once the question is verified—once it is shown to have been formulated properly—the specific answer, either “yes” or “no”, validates (gives a green or a red light to) its content. If this process seems to be prone to error recall that we started by assuming p in order to obtain Q and A , so the possibility of re-obtaining p by re-combining Q and A is a priori guaranteed by hypothesis and sceptical suggestions would merely be out of place here. All this can be formulated more precisely by saying that $A_{0/1}$ saturates $Q_{0/1}^{CLP}$ by

1. implicitly *verifying* its CLP parameters: both “yes” and “no” implicitly signal that the question is being asked in the right context, at the right LoA and for the right purpose and
2. explicitly *validating* its content, as a model of the system (roughly: “yes” and “no” provide a green or a red light for the question, respectively).

Figure 3 summarises how far we have progressed in reverse engineering semantic information.

Clearly, a correct saturation consists in a *correct verification* and a *correct validation*. It has taken several clarifications and distinctions and a bit of technical vocabulary, but we have finally reached the heart of our problem.

6 Fifth Step: Correctness

Let us quickly review our progress. Simplifying, we now know that p qualifies as semantic information (for an epistemic agent a , first parameter) about a system s (second parameter) if and only if p is true; that p is true if and only if A correctly saturates the Boolean question Q corresponding to p ; and that A correctly saturates Q if and only if it correctly verifies and validates it, thus generating an adequate model m of s . Having reduced truth (of semantic information) to adequacy (of the corresponding model m) via correctness (of A with respect to Q), our next challenge is the analysis of the correctness of A .

The challenge consists in negotiating two crossroads in a row. The first is represented by the twofold correctness of the saturation. I shall return to the issue of what it means for A to verify correctly Q in Sect. 7.5. Here, let me just highlight the fact that the correct verification of Q by A is a formal precondition for the development of an adequate model m of the targeted system s : it is necessary for, but does not contribute to, the truthfulness of i . In other words, the analysis of the correctness of the verification cannot help us in understanding what it means for semantic information to be truthful. At this crossroad, the really interesting path is represented by the correct validation of Q by A . By following it, we encounter the second crossroad, represented by two further alternatives. For now we can either analyse correctness of the validation in terms of some concept of truth, thus showing consistency but also failing to provide a non-circular analysis of what it means for semantic information to be true. Or we can move forward, and check whether a further reduction of the correctness of the validation—and hence of the adequacy of the issuing model in terms that are truth-poietic but not truth-dependent—is possible. Let us quickly review the circular path first.

A useful way to test whether our reverse engineering process is still on the right track is by showing that we have not lost touch with our starting point. Statistics provides the standard analysis of what it means for a model to be adequate (Freedman et al. 2007). A model is adequate with respect to its target system if it is *valid*. This is now the *statistical* (not the software engineering or the logical) concept of *validity*, which is to be understood as the result of a combination of *accuracy* and *precision*, two other technical concepts borrowed from statistics. Although one might have the impression that we are actually gaining some new ground, it is easy to see that this road only leads back to our starting point. For statistical *accuracy* is the degree of conformity of a measure or calculated parameter (belonging to the model) to its actual, that is, *true*, value (belonging to the system). And statistical *precision* is the degree to which further measurements or calculations show the same or similar results. So it turns out that the statistical concepts of validity, accuracy and precision—even assuming that we could adapt them to our less quantitative needs, and hence exploit them to clarify what we mean by an adequate model—ultimately presuppose a truth-dependent relation of *conformity* and hence cannot provide a philosophically foundational analysis of truth itself without begging the question. The silver lining in all this is that such internal coherence is reassuring: we have not got lost in some conceptual wilderness, while searching for the mechanism that generates semantic information. Encouraged by

the knowledge that we could still go back to square one should we wish to do so, let us not press the panic button but push forward.

The second path should lead us away from semantics and epistemology, if we want to avoid ending up back where we started, and take us into the realm of pragmatics, that is, the realm of actual and hopefully successful *interactions*—between an agent a holding the information that p , the model m generated by p , and the system s modelled by m —that can provide some *exogenous* grounding for the evaluation of the quality of the model itself. In order to achieve this, I shall ask the reader to bear with me a bit longer, as I need to introduce two technical concepts to make sense of such interactions.

One is that of *proxy*, and is borrowed from Information and Communication Technology (Luotonen 1998). Technically, it refers to a computer agent (e.g., a network service) authorized to act on behalf of another agent (the client), e.g., by allowing another computer to make indirect network connections to other network services (the server). In this sense, a proxy can be an interface for services that are remote, resource-intensive, or otherwise difficult to use directly. Note that the “proxy-ing” system need not be a copy, an image, a representation or a reproduction of the “proxy-ed” system (the client).

The other concept is that of *commutative diagram*, and is borrowed from category theory (Barr and Wells 1999). Technically, it refers to a diagram of objects (vertices) and morphisms (arrows) such that, when selecting two vertices, one can follow any directed path through the diagram and obtain the same result by composition.

Adapting these two concepts to our needs, we can now reverse engineer the correctness of the validation, and hence the adequacy of the ensuing model, in terms of the commutativity of the accessibility relation, thus (see Fig. 4 for a more intuitive presentation, all Greek letters in [COR] refer to paths in the diagram in Fig. 4):

[COR] $A_{0/1}$ correctly validates $Q_{0/1}^{CLP}$ about a target system s identified by $Q_{0/1}^{CLP}$ if and only if $Q_{0/1}^{CLP} + A_{0/1}$ generates (β) an adequate model m of s ; and m is an adequate model of s if and only if m is a proxy (δ) of s such that, if a holds (α) $Q_{0/1}^{CLP} + A_{0/1}$, then a 's proximal access (γ) to m commutes with a 's distal access (ϵ) to s .

[COR] offers two advantages and raises a problem. The first advantage is that it finally introduces an explicit reference to an informee a . This is crucial, since




Fig. 4 The meaning of [COR]. $Q + A$ is a simplification for $Q_{0/1}^{CLP} + A_{0/1}$

semantic information is an *objective* (i.e., not subjective) but also *liminal* (that is, neither internal nor external) and *relational* (that is, neither absolute nor relative) concept, like food. It makes little sense to talk about the presence and nature of food without any reference to the specific type of feeder. Likewise, something counts as semantic information only with respect to a specific type of informee. Note that this does not mean that we are going to relativise truth to some specific agent, translating “ p is true” into “ p is true for a ”. Relational does not mean relative. The move here is rather Kantian: we are looking into the conditions of possibility of the truthfulness of some semantic information p about s given the fact that p objectively counts as semantic information about s for an epistemic agent a .

The second advantage is that [COR] explains the well-known fact that semantic information provides *distal* access to its target. If the agent in the bedroom upstairs asks whether the beer is located in the fridge, and the agent in the kitchen downstairs answers positively, then the agent upstairs, by having proximal access to this overall piece of information, gains distal access to the presence of the beer in the fridge, as long as the answer is *correct*. [COR] merely combines this into a single agent’s informative state.

The problem concerns the interpretation of the relation of distal and proximal *accessibility*. If we were to interpret it alethically or epistemically this would obviously fail to take us off the semantic merry-go-round and, sooner rather than later, we would be sent back where we came from. The good news is that we do not need to go down that modal road. On the contrary, the sort of accessibility at stake here is a matter of pragmatic or factual interaction, which provides an exogenous grounding of correctness. It is the one that we find specified in computer science, where accessibility refers to the actual permission to *read* (technically, *sense* and *retrieve*) and/or *write* (again, technically *modify* and *record*) data as a *physical* process. The result is that a ’s proximal access to m commutes with a ’s distal access to s if and only if a can read/write s by reading/writing m .

The writing of s through the writing of m is admittedly rare, but it is useful to illustrate it in order to convey the sense of concrete interaction with the targeted system that is involved. Thus, we have left behind a magic culture that considered it an ordinary phenomenon (cf. the practice of sticking pins in a doll as a method of cursing an individual). Nevertheless, self-fulfilling prophecies (Bill Gates confessing that “Microsoft’s shares are overvalued”), performative sentences (the baptising priest declaring that “the name of this girl is Mary”), magic-placebo formulae (the guru concluding that “you are now healed”), authoritative-fictional descriptions (“Sherlock Holmes never visited the Bodleian Library” written by Conan Doyle), God’s *intellectual intuition* that p , according to Kant, and other ways of “doing things with words” (“this train is not leaving the station” uttered by a dictator) are a good reminder that it is far from impossible to modify/record a system by accessing only its model. Of course, access to m is most commonly used in order to read (i.e., sense and retrieve) s by reading (ditto) m . One gains distal access to (part of) the actual, physical system represented by the fridge in the kitchen and its contents (one senses and retrieves the data in question at a distance) by gaining proximal access to its (partial) model represented by the semantic information “the beer is in the fridge”. A way of conveying the same point is by relying on a subjunctive

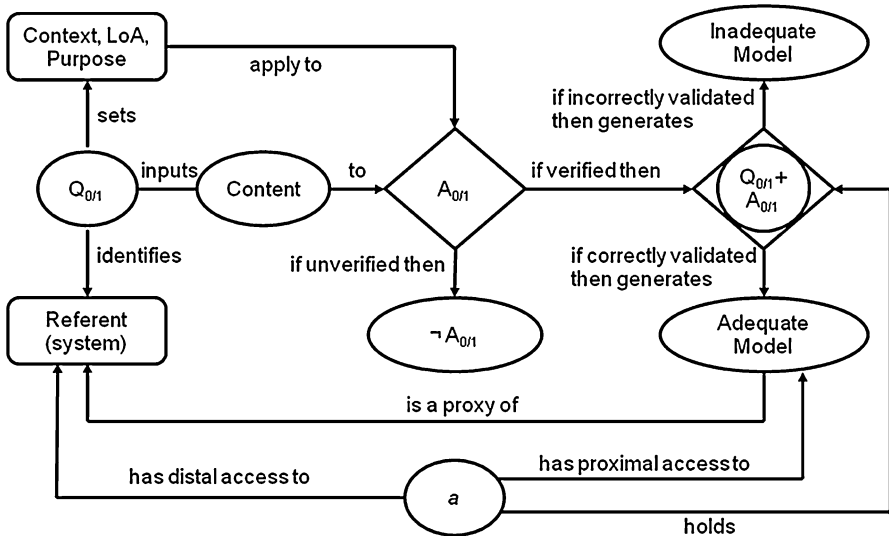


Fig. 5 The correctness theory of truth

formulation: the proximal read/write access to m as a proxy of s commutes with the distal read/write access to s if and only by having read/write access to m one were having read/write access to s . This happens in space as well as time: imagine the question being “Will the train leave from platform one?” and the answer being “yes”. Semantic information may be seen as a way of being telepresent (Floridi and Sanders 2005).

We needed actual interaction with the system being modelled in order to ground exogenously the correctness of the (validation provided by the) answer to the question pertaining to it, and we have now obtained it. Our toiling is almost over. Putting together this last piece of our jigsaw puzzle, we obtain Fig. 5 (the reader may check that this is simply the result of merging Figs. 3 and 4, even if this may not be immediately obvious visually).

Figure 5 represents the blueprint of the mechanism that underlies the truthful nature of semantic information. If we apply it to our example, we obtain:

1. “the beer is in the fridge” qualifies as semantic information if and only if
2. “the beer is in the fridge” is true; this is the case if and only if
3. “yes” is the correct answer to (i.e., correctly saturates by correctly verifying and validating) the question “is the beer in the fridge?”; this is the case if and only if
4. “is the beer in the fridge?” + “yes” generate an adequate model m of the relevant system s ; this is the case if and only if
5. m is a proxy of s and proximal access to m provides distal access to s ; and finally this is the case if and only if
6. reading/writing m enables one to read/write s .

That is, if “the beer is in the fridge” qualifies as semantic information, then holding that semantic information is tantamount to accessing the particular feature of the system addressed by the model which, in our example, is the location of the beer inside the fridge.

7 Some Implications and Advantages of the Correctness Theory of Truth

A good way to explore some of the features of what I shall label the *correctness theory of truth* (henceforth CTT) is by making explicit some of its implications and advantages. They will also help us to grasp the similarities and differences between CTT and other standard approaches to truth.

7.1 Truthmakers and Coherentism

At the beginning of this article, we saw that semantic information is, strictly speaking, not truth-bearing but truth-constituted. What is truth-bearing is rather content, which gets upgraded to semantic information only if it is truthful. It follows that, since a truth-maker is that in virtue of which a truth-bearer is true, CTT is compatible with a variety of theories of truth-makers, insofar as these are successfully adaptable and applicable to content. This is fine but probably less interesting than the fact that, since CTT seeks to reduce truth to correctness, it also translates the question about truth-makers into a question about correctness-makers: what is it that in virtue of which a correctness-bearer (i.e., $A_{0/1}$ as a correct answer) is indeed correct? A quick, Aristotelian reply may point to the system s as the most plausible candidate. For a look at Fig. 5 may suffice to convince one that it is because of what s is that the model m may qualify as adequate, and hence that the answer $A_{0/1}$ that generates m may be correct. The story, however, is slightly more complicated. To show why let me first sketch two analogies.

Consider the case in which a non-atomic formula F in propositional logic (e.g. $(\neg P \vee Q) \rightarrow S$) is declared to be well-formed. One may ask what it is that in virtue of which F is indeed well-formed, that is, what its well-formedness-maker is. Pointing to the right sub-set of formation rules as the relevant well-formedness-makers would definitely be a good answer but, like pointing to the system s before, it would also be only partial. For what would be missing is the implicit fact that we are talking about a dynamic system: F is well-formed also because it has been (or may be) constructed according to the relevant formation rules *recursively*. So the formation rules in question are only the most salient, necessary source of the well-formedness of F . Strictly speaking, they would be insufficient by themselves.

Consider next the more complex case of a bottle of Chianti wine. If we call that wine the Chianti-taste-bearer and ask what its Chianti-taste-maker is, pointing to Sangiovese grapes is fine, indeed necessary, but also insufficient. Clearly, the whole process, through which the wine is produced, makes a very significant difference and is as much a part of that in virtue of which the wine bears the taste of Chianti

(the Chianti-taste-maker) as the original grapes. The Sangiovese grapes are only the necessary source of the Chianti-taste of the wine in question.

If we now go back to what a correctness-maker is according to CTT, given the engineering approach adopted by the theory, and hence the treatment of semantic information as an artefact, it should be clear that the dynamic “making” plays an essential role. The essential element that makes $A_{0/1}$ correct is indeed the system s , which is the necessary correctness-maker that is the source of the correctness of $A_{0/1}$. But s becomes sufficient only if it is embedded in the right sort of network of dynamic relations, as shown in the previous pages. For CTT, the system s is the source of correctness but the correctness-maker only in a loose way of speaking. The necessary and sufficient correctness-maker is the whole complex construct, represented by the configuration of the entire distributed system¹⁴ in which s is embedded as a source of correctness, including the commuting relation. Two final clarifications are now in order.

First, the previous point might be looked upon with some sympathy by followers of coherentist approaches to truth. At the risk of losing some allies, let me clarify that this would be a mistake. CTT may indeed be compatible with a coherentist theory of truth, but it is very far from being (committed to) one. Depending on the CLP parameters, some models play the role of systems against which other models are evaluated. In our example, we take a propositional model (“the beer is in the fridge”) as less fundamental than a perceptual model (e.g., the observable presence of beer in the fridge, or the grasping of some beer once the fridge is opened). We could have used a memory-based system (the recollection that the beer had been placed in the fridge) or a testimony-based alternative (the reassurance by someone who knows where the beer is, that is, in the fridge). This is where CTT is closer to a coherentist theory of truth. However, the network of dynamic relations specified above (the blueprint illustrated in Fig. 5) has very little to do with either coherence or consistency within an information system, which should be analysed independently of a correctness theory of truth, in terms of information *integrity*. Thanks to the commuting relation, correctness is not an internal property of the system, but the external feature of $A_{0/1}$ that guarantees the successful, pragmatic interaction with s through m .

Second, note that, as I have just specified, the system s , which is the source of the correctness of $A_{0/1}$ and hence of the adequacy of the model m , may be another model n . In this case, supporters of some forms of relativism and internalism may rejoice, but this reaction too would be premature. CTT offers an *intra-model* not an *infra-model* theory of truth. The reader may recognise here a Kantian point: CTT analyses truth as a relation between models and never shifts from talking about semantic information to talking about systems in themselves, yet this is not a form of antirealism. In CTT, truth is ultimately a matter of assessment of what is claimed to be the case against what is taken to be the case, within the (often very rigid) constraints and (often rather limited) affordances offered by the targeted system.

¹⁴ The occurrence of the term “system” here is unfortunate but inevitable (it is dictated by standard terminology in model analysis). Luckily, it should not generate any confusion, since it clearly refers to the whole blueprint described by Fig. 3.

Ultimately, it is the way the system is (the beer being in the fridge) that interactively constrains (recall the relation of commutation) the value of its models (“the beer is in the fridge”, “there is beer in the fridge” “the fridge contains some beer”, “if you are looking for some beer look inside the fridge” and so forth) and determine their truth, even if the only way to deal with the system epistemically is through its models. The relation between model and system is not one of pictorial representation or resemblance or similarity (no metaphorical mechanism of mirroring or photocopying is at stake here) but one of fit, in the way in which a key corresponds to a lock. To think that CTT supports an “any key works” sort of policy would be a mistake. The philosopher sharing some (naïve, direct, scientific, commonsensical, etc.) form of realism may be willing to accept CTT but then graft to it some form of (possibly privileged) epistemic access to s , or what is taken to be the case, in terms of what is actually the case. This is indeed an option. But it is not the one that I or any Kantian favour, because the alternative advocated by CTT points in the direction of a more modest epistemology, a safer commitment and a less ontologically-overloaded conception of truth. As in statistics, in CTT we never talk about the ultimate, real nature of the world in itself, but compare data sets (the model) to data sets (the system), or phenomena (in the Kantian sense) to other phenomena (still in the Kantian sense). According to CTT, truth is a successful transduction of models among possible worlds. Truth is commutation. Further discussion of this crucial point leads to the next topic.

7.2 Accessibility and Correspondence

The reader might have noticed that, in CTT, the system s is first identified by $Q_{0/1}^{CLP}$ and then modelled by $Q_{0/1}^{CLP} + A_{0/1}$ and that this twofold manoeuvre is paralleled by the double-access that a enjoys (at least in principle) to s as posed by $Q_{0/1}^{CLP}$ and to s as modelled by $Q_{0/1}^{CLP} + A_{0/1}$. We are able to check whether an answer is correct, and hence the issuing model is adequate, with respect to a posited s (whether some content about a targeted s is true and hence qualifies as semantic information) only if we have both s and (at least in principle) an *alternative* way of reading/writing s . When it comes to empirical knowledge, nobody could check the truthfulness of a newspaper by buying a second copy of the same issue, to paraphrase Wittgenstein. So, contrary to what the correspondence theory of truth sometimes seems to suggest, in CTT truth is about positing *and* modelling a system and therefore having double-not single-access to it. Metaphorically, we capture the world and its features by using pincers, not harpoons. In scientific research, this is well-known and common practice. For in order to understand whether a model is correct, a scientist looks at the data set and considers whether the model can successfully predict the behaviour of some aspect of the system being modelled (Davison 2003). The better the model, the smaller the disagreement between its forecast and what happens in the *observed* system. In the technical vocabulary of testing theory, this means that, for a given model under test (MUT, our m), there is a reference model (RM, also known as an oracle, because it is assumed to be an infallible source of truth, our s) that serves as

the basis for the construction of MUT and can then be used to establish whether MUT is behaving as required, through some alternative access to RM itself. This scientifically realistic, if not philosophically realist, feature allows CTT not to solve but to bypass two classic problems threatening theories of truth as correspondence according to which only one relation connects truth-bearers and truth-makers:

- (a) how systems (facts, in some correspondentists' terminology) may be understood independently of models (true sentences, still following the same terminology) without merely making them their "tautological accusatives", to paraphrase Armstrong; and
- (b) a version of the slingshot argument to the effect that all models (true sentences) correspond to the same "Great System" (the "Great Fact").

CTT avoids (a) because it argues that specific systems are posited and accessed independently of how they are modelled in the first place, and there are only specific models of specific systems, developed by fulfilling specific CLP parameters. So CTT's twofold access approach further explains why it is not an internalist theory of truth and avoids (b): although all truth evaluations occur between models, there is no ultimate "Great System" which all models adequately describe.¹⁵

A significant advantage of CTT is that it avoids overloading truth with a double-task. Some theories of truth are *double-tasking* in that they require truth to work both semantically—in order to explain what it means for truth-bearers to have the truth-value they have, e.g., what it means for "the beer is in the fridge" to be true—and ontologically, in order to explain what the world is like if truth-bearers have the truth-value they have, e.g., what the world is like if "the beer is in the fridge" is true. Metaphorically, such theories identify only one road between truth-makers and truth-bearers, which they assume to be two-way. On the contrary, CTT decouples the semantic from the ontological task and requires truth to be only a semantic relation between models. In this, the similarity with Tarski's approach is obvious: according to CTT "snow is white" is true if and only if "yes" is the correct answer to "is snow white?". The difference lies in the pragmatic (as opposed to model-theoretic) and hence exogenous turn that CTT takes when it grounds the correctness of the answer: having read/write access to the model *m* that "is snow white? + yes" generates commutes with having read/write access to the substance in question and its whiteness (the system *s*).

What CTT also shows is that deflationist theories of truth, when applied to semantic information, may be right, but in a trivial and uninteresting way. Since semantic information encapsulates truth, it is not truth-bearing but truth-constituted, so qualifying it as true is worse than informationally redundant, it is pointlessly noisy. If "the beer is in the fridge" qualifies as information, to add that it is true fails to provide any further information and only messes up the communication, wasting resources. But to strip semantic information of such a uselessly¹⁶ redundant qualification leaves the problem of its truthfulness (or of the truthfulness of the

¹⁵ Young (2002) has shown that even in the case of a correspondence theory of truth it is at least controversial whether the slingshot argument undermines it.

¹⁶ Redundancy is often useful, but in this case it is pointless redundancy that is in question.

corresponding content) untouched, and hence unsolved. We still need to run our reverse engineering process in order to understand what it means for p to qualify as semantic information. And as soon as we transform p into a Boolean question, we know that the problem of the truth of p has been transformed into the problem of the correctness of the answer.

7.3 Types of Semantic Information and the Variety of Truths

We have already seen that CTT can account for the nature of tautologies and contradictions, but any acceptable theory of truth for semantic information should also be able to deal satisfactorily with a variety of genuine types of semantic information and hence with their truths. Happily, CTT proves to be sufficiently flexible. Here is a quick review.

We would like to be able to treat fictional truths, such as “Watson is Sherlock Holmes’ best friend”, future truths, such as “the flight will leave at 12.30 tomorrow”, negative truths, such as “whales are not fish” (see Fig. 4), ethical truths, such as “rape is morally wrong”, modal truths, such as “beer can be stored in a fridge”, dispositional truths, such as “sugar is soluble in water”, and metaphorical truths, such as “Achilles is a lion” (or even more complex cases such as “Mary is not a fox”) as genuine instances of semantic information. CTT allows this treatment rather easily. In each case, the system s in question, posed by Q (e.g., “is Watson Sherlock Holmes’ best friend?”), is distally accessed through the model generated by the correct answer (“yes”) because CTT is not ontologically committed to the empirical existence of s but rather treats it as the reference model (s could be a segment of any possible world). A major advantage, over standard theories of truth as correspondence, is that this allows CTT to avoid any reference to some existing fictional facts, negative facts, queer moral facts, parallel modal facts, dispositional fact or metaphorical facts, to which such truths would allegedly correspond. We never check semantic information (e.g., “whales are not fish”) against some fact (about their non-fishiness), we check it against other semantic constructs, which might be narrative (in Sherlock Holmes’ case), decisional (in the flight’s case), biological (in the whales’ case), ethical (in the rape case), modal (in the storability case), dispositional (in the solubility case) and so forth.

One may object that treating fictional, empirical, ethical, modal, dispositional and metaphorical instances of semantic information (independently of whether negative or positive, or past, present or future) as all bona fide true impoverishes our capacity to discriminate between reality, imagination and social conventions or stipulations. But this would be a fair criticism only if one were to forget that the whole analysis must be conducted by paying careful attention to the LoA, the context and the purpose of the corresponding questions. To simplify, “Achilles is a great warrior” is an instance of semantic information, and hence it is true, not only because “yes” is the correct answer to the corresponding question, but also because we take for granted Homer’s Iliad as the right CLP framework. Consider “snow is white”, “milk is white” and “teeth are white”. Comparing these instances of semantic information is enlightening because, from such truths taken separately, it does not follow necessarily, at least not in CTT, that therefore “milk, snow and teeth have the

same colour” is also true. This is because of the crucial role played by the CLP parameters. “Milk, snow and teeth have the same colour” is true if and only if “yes” is the correct answer to the corresponding Boolean question, but now one cannot determine whether that answer is indeed correct unless one specifies the context in which, the LoA at which, and the purpose for which that question is being asked. Change the available palette (different LoA) or the purpose (redecorating the living room, say, instead of having one tooth replaced), for example, and the question may receive different answers. This is not relativism, it is, for want of a better word, “preciseism”. It is a fallacy to fuse two or more instances of semantic information into a large instance without making their CLP parameters homogenous, at least implicitly. If this seems too easy and commonsensical, it is worth recalling that we are only reaping the fruits of the hard labour done in the previous pages.

Our opponent may still be unconvinced. He might retort that there is still a risk of causing an inflation of truths. Such concern is misplaced. “The earth is flat”, “Sherlock Holmes is happily married to Watson”, “in 2012 the Olympic Games will take place in Rome”, “horses are oviparous”, “the use of violence against women is always justified” fail to qualify as semantic information because they are false, this because the corresponding questions are correctly answered in the negative, and this because affirmative answers do not commute with the systems posed by the corresponding questions. The point is important and deserves a fuller treatment in the next section.

7.4 A Deflationist Interpretation of Falsehood as Failure

CTT treats untruth (falsehood) as commutation failure. The treatment comes as rather natural if one realizes that

- (i) in logic programming, negation as failure (NAP) is a non-monotonic inference rule used to derive $\neg P$ from the failure to derive P (Gabbay et al. 1993); that
- (ii) the so-called stable model semantics, which gives a semantics to logic programming with NAP, is a simplified form of autoepistemic logic (Nerode and Shore 1997), and that
- (iii) $\neg P$ may have not only the classic meaning but also the modal meanings, in autoepistemic logic, of “ P is not believed”, “ P is not known” or “ P cannot be shown” (Gelfond 1987).

The further but rather simple step taken by CTT consists in interpreting “ P is not true” (false for the classicist) as $\neg P$ and then analysing $\neg P$ as equivalent to commutation failure of the relevant diagram. The expanded autoepistemic semantics can then be given in terms of “ P is not information”. To illustrate more intuitively what all this amounts to, and see the advantage of such minimalism, consider the following example: “the earth has two moons”. Following CTT, the usual analysis requires a specification of the CLP parameters, posed by the corresponding question “does the earth have two moons?”. Once we have ascertained that we are talking about our planet considered astronomically and in light of our current knowledge (not, for example, of some twin earth in another

possible world; or some future earth whose moon has been split into two; or some other planet also called earth; or some earth described in a sci-fi novel as having two moons; or some ancient text in which the earth is described as having two moons, etc.), the answer “yes” provides a model (the earth with two moons) the proximal access to which fails to commute with the distal access to the astronomical system in question. There is a failure in the information flow, and this is what it means for “yes” to be incorrect, and hence for “the earth has two moons” to be untrue (false). The advantage of this minimalism is that there is no need to treat truth and untruth (falsehood) in the same way: untruth (falsehood) is best understood as the mere absence of truth, a lesson well-known to any non-Manichean philosopher.

7.5 The Information-Inaptness of Semantic Paradoxes

Semantic paradoxes are often seen as the ultimate benchmark of a theory of truth. The point of this section, however, is not to argue in favour of a CTT-based solution of them—an impossible task, given the nature and length of this article—but rather to see what semantic paradoxes may teach us about CTT.

Consider first the task of preventing the occurrence of semantic paradoxes. In this, CTT’s strategy is partly Russellian, partly Tarskian. This comes as no surprise if one realises that, technically speaking, CTT—with its emphasis on the importance of the CLP parameters and especially on the Method of Abstraction and its use of Levels of Abstraction—represents a late incarnation of Russell’s approach to semantic paradoxes in terms of type theory. The modern lineage, of some interest for the historian, is through the adoption and refinement, in programming language theory, of Russell’s and (later) Church’s theory of types in order, for example, to construct type-checking algorithms to analyse compilers for programming languages and avoid the disasters caused by unconstrained self-reference. CTT is simply reclaiming to philosophical analysis what was its own in the first place.

Consider next the task of treating semantic paradoxes once they have occurred. CTT can explain their occurrence in terms of failure to respect some constraints, e.g. about object language and metalanguage. It can then interpret their value, as alleged instances of semantic information, by relying on the reverse engineering procedure detailed in the previous pages, with the following results.

Semantic paradoxes are notoriously caused by self-referential mechanisms. *Internal* semantic paradoxes are those in which the self-referential relation occurs within the message itself (the semantic information *i*), independently of the sender. The classic example is of course “this sentence is false”. Following CTT, the verdict on similar paradoxes is that they fail to pass the *verification* stage, in the computer science sense introduced in Sect. 5. For consider the erotetic structure of “this sentence is false”. Once the CLP parameters are taken care of, if “this sentence is false” must count as semantic information, it must be true, and hence informationally equivalent to “is this sentence false?” + “yes”. But then it becomes easier to see that, before trying to understand the role of “yes”, one should acknowledge that “is this sentence false?” is a question, not a declarative sentence

at all, which is not truth-apt (it makes no sense to ask whether it can be correctly qualified as either true or untrue). So CTT can show this and other internal semantic paradoxes (e.g., “the next sentence is false. The previous sentence is true.”) to be badly engineered informational artefacts, comparable to any blueprint of a perpetual motion machine. Note that this applies to vicious as well as virtuous cases: “this sentence is true” is equally self-referential, it also fails to pass the verification stage (“is this sentence true?” is not truth-apt) and hence cannot count as semantic information, according to CTT.

The previous approach is ineffective towards *external* semantic paradoxes. In this case, the self-referential relation is between the message (the semantic information *i*) and its sender. The classic example is of course “Cretans always lie”, suitably refined. In this case, there is nothing wrong with the erotetic structure of the message (“do Cretans always lie?” + “yes”). The problem is with its relation to the sender, when the message comes from a Cretan like Epimenides. Recall the example in which Mary—now Epimenides—can make statements not by uttering declarative sentences but only through Boolean questions followed by the corresponding Boolean answer. If “Cretans always lie” counts as semantic information it should be true, and hence equivalent to “do Cretans always lie?” + “yes”, where both “do Cretans always lie?” and “yes” are messages sent by the same source. And this is where the problem arises. For imagine the case in which you wish to know whether Cretans always lie. Asking a Cretan whether they do would provide you with no information: you would not know whether Cretans lie all the time, no matter what the Cretan answers. This means that a self-certifying question cannot be informatively asked to the source that needs to be certified. But this holds true even when it is the source itself that asks and then answers the self-certifying question. Mary cannot convey any semantic information by saying “am I lying? Yes” because, by asking *Q*, she has ipso facto forfeited the possibility of answering it informatively. As in the previous case, the analysis treats vicious and virtuous cases in the same way. Informationally speaking, “Cretans never lie”, uttered by a Cretan, and “I always tell the truth”, run into the same problem faced by their paradoxical counterparts: they are equally disqualified by CTT as failing to pass the verification step to qualify as semantic information.

To summarise, both internal and external semantic paradoxes are faulty artefacts that fail to qualify as semantic information because they fail to pass the verification stage. This does not mean that they are useless informationally. Semantic paradoxes may help the flow of information by fulfilling a phatic function: they can perform the social task of establishing, prolonging or discontinuing communication, or simply confirming whether the receiver is still there, exactly like “how are you?” or the inarticulate sounds made by a listener during a telephone conversation are not meant to provide (or gain) any information.

The reader acquainted with the literature on semantic paradoxes may still be left with at least one further doubt: what happens when the semantic paradox has an erotetic format to begin with? Russell formulated his own paradox in terms of a question, but one may retort that, in his case, the problem is set-theoretical, not semantic. Nevertheless, there are other paradoxes that are both semantic and

erotetic, such as Smullyan's "is the answer to this question "no"?".¹⁷ How does CTT fare in this case? The answer is simple. If p is to count as semantic information, the relation between p and [$p?$ + answer] must be a biconditional. But this means that, independently of which answer one may later provide to "is the answer to this question "no"?", in order to count as the first half of the erotetic equivalent of some semantic information, that question must correspond to the message "the answer to this question is "no"" (or "the answer to this question is not "no""), but note that this is not a question, but a declarative sentence, hence it is malformed. It follows that this version of the semantic paradoxes too poses no problem for CTT, which diagnoses them as cases of verification failure.

8 Conclusion

We have come to the end of a rather long journey. The hope is that the effort might have been rewarding both in itself, if it has been clear enough, and in terms of the final result, if it has been sufficiently convincing. At this point, the reader will probably wish me to keep this conclusion as short as possible. I shall oblige, by adding only a final comment.

Theories of truth often seem to be developed with passive viewers of an outside world in mind, detached observers, whether inside or outside Plato's cave, TV watchers, radio listeners, movie goers, in short, *systems users*, according to the computer science terminology favoured in this article. The correctness theory of truth, proposed in the previous pages, should rather be seen as an attempt to cater for a different sort of customer, namely embodied and embedded, creative agents, who interact with reality, shape and build it, Plato's artisans, writers not just readers, players not audience, in short *systems designers*. To these customers, truth is about constructing and handling artefacts and interacting with them successfully, not merely experiencing them passively. Unfortunately, this is not very Greek, but it is still a very respectable tradition to which both Russell and Tarski belong, insofar as their groundwork in model theory concerned the design of systems.

Acknowledgments A first version of the paper was presented at the *I International Meeting of Experts in Information Theories* (Leon, Spain, 6–7 November, 2008) and I am very grateful to Francisco Salto for his kind invitation, and to the audience for the lively and valuable discussion. Several revised versions were then presented at the following meetings: *Formal Philosophy Seminar*, University of Leuven, February 2009 (thanks to Sebastian Sequoiah-Grayson); Department of Philosophy, Bilkent University, March 2009 (thanks to Hilmi Demir); Dipartimento di Filosofia, Università degli Studi di Padova, March 2009 (thanks to Massimiliano Carrara and Francesca Menegoni); *North-American Computing And Philosophy* (NACAP) Conference, Indiana University, June 2009 (thanks to Anthony Beavers); Akademie der Wissenschaften and Georg-August-Universität Göttingen, July 2009 (thanks to Robert Schaback); Dipartimento di Filosofia, Università degli Studi di Genova, October 2009 (thanks to Carlo Penco); The Moral Sciences Club, Cambridge University, October 2009 (thanks to Jane Heal and Steven Methven). I would also like to acknowledge the help, useful comments and criticisms by Patrick Allo, Mark Jago, Sebastian Sequoiah-Grayson, and Matteo Turilli. I remain deeply indebted to Michael Dummett and Susan Haack for their clarifications and feedback, which date to almost 20 years ago. I am afraid this paper has been a work in progress for quite some time. Joanna Gillies kindly copyedited the

¹⁷ For the attribution to Smullyan see Landini (2007).

last version. Finally, I wish to acknowledge the feedback provided by the two anonymous referees, who offered many constructive suggestions on how to improve the paper and saved me from more than a couple of blunders. All the aforementioned people helped me to develop the paper substantially, but they are not responsible for any remaining mistakes.

References

- Austin, J. L. (1950). Truth. In J. O. Urmson & G. J. Warnock (Eds.), *Philosophical papers* (pp. 117–133). Oxford: Oxford University Press.
- Bar-Hillel, Y., & Carnap, R. (1953). An outline of a theory of semantic information. Reprinted (1964) In Y. Bar-Hillel (Ed.), *Language and information: Selected essays on their theory and application* (pp. 221–274). Reading, MA; London: Addison-Wesley.
- Barr, M., & Wells, C. (1999). *Category theory for computing science* (3rd ed.). Montreal: CRM Press.
- Barwise, J., & Seligman, J. (1997). *Information flow: The logic of distributed systems*. Cambridge: Cambridge University Press.
- Bremer, M., & Cohnitz, D. (2004). *Information and information flow—An introduction*. Frankfurt: Ontos Verlag.
- Chellas, B. F. (1980). *Modal logic : An introduction*. Cambridge: Cambridge University Press.
- Colburn, T. R. (2000a). Information, thought, and knowledge. In *Proceedings of the world multiconference on systemics, cybernetics and informatics*, pp. 467–471.
- Colburn, T. R. (2000b). *Philosophy and computer science*. Armonk, NY: M.E. Sharpe.
- Davison, A. C. (2003). *Statistical models*. Cambridge: Cambridge University Press.
- Devlin, K. J. (1991). *Logic and information*. Cambridge: Cambridge University Press.
- Dodig-Crnkovic, G. (2005). System modeling and information semantics. In J. Bubenko, O. Eriksson, H. Fernlund, & M. Lind (Eds.), *Proceedings of the fifth promote IT conference, Borlänge, Sweden*. Lund: Studentlitteratur.
- Dretske, F. I. (1981). *Knowledge and the flow of information*. Oxford: Blackwell.
- Dretske, F. I. (1988). *Explaining behavior: Reasons in a world of causes*. Cambridge, MA: MIT Press.
- Eilam, E. (2005). *Reversing: Secrets of reverse engineering*. Indianapolis, IN: Wiley.
- Engel, P. (2002). *Truth*. Chesham: Acumen.
- Fetzer, J. H. (2004). Information, misinformation, and disinformation. *Minds and Machines*, 14(2), 223–229.
- Floridi, L. (2004a). Information. In L. Floridi (Ed.), *The Blackwell guide to the philosophy of computing and information* (pp. 40–61). Oxford: Blackwell.
- 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 W. A. Darity (Ed.), *International encyclopedia of the social sciences*. Detroit: Macmillan.
- Floridi, L. (2008b). The method of levels of abstraction. *Minds and Machines*, 18(3), 303–329.
- Floridi, L. (2008c). 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 network theory of account. *Synthese*.
- Floridi, L., & Sanders, J. W. (2005). Internet ethics: The constructionist values of homo poieticus. In R. Cavalier (Ed.), *The impact of the internet on our moral lives*. New York: SUNY.
- Formigari, L. (2004). *A history of language philosophies*. Amsterdam: John Benjamins Pub.
- Fox, C. J. (1983). *Information and misinformation : An investigation of the notions of information, misinformation, informing, and misinforming*. Westport, CT: Greenwood Press.
- Fox, C. J. (2007). *Introduction to software engineering design*. Boston, MA/London: Pearson/Addison Wesley.

- Freedman, D., Pisani, R., & Purves, R. (2007). *Statistics* (4th ed.). New York: W.W. Norton.
- Gabbay, D. M., Hogger, C. J., & Robinson, J. A. (1993). *Handbook of logic in artificial intelligence and logic programming*. Oxford: Clarendon Press.
- Gamma, E., Helm, R., Johnson, R., & Vlissides, J. (1995). *Design patterns: Elements of reusable object-oriented software*. Reading, MA: Addison-Wesley.
- Gelfond, M. (1987). On stratified autoepistemic theories. In *Proceedings of national conference on artificial intelligence (AAAI)*, pp. 207–211.
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Groenendijk, J., & Stokhof, M. (1994). Questions. In J. Van Benthem & A. Ter Meulen (Eds.), *Handbook of logic and language*. North-Holland: Elsevier Science.
- Israel, D., & Perry, J. (1990). What is information? In P. P. Hanson (Ed.), *Information, language, and cognition* (pp. 1–28). Vancouver: University of British Columbia Press.
- Künne, W. (2003). *Conceptions of truth*. Oxford: Clarendon Press.
- Landini, G. (2007). *Wittgenstein's apprenticeship with Russell*. Cambridge, UK: Cambridge University Press.
- Luotonen, A. (1998). *Web proxy servers*. Upper Saddle River, NJ: Prentice Hall PTR.
- Lynch, M. P. (2001). *The nature of truth: Classic and contemporary perspectives*. Cambridge, MA: MIT Press.
- Nerode, A., & Shore, R. A. (1997). *Logic for applications* (2nd ed.). New York: Springer.
- Popper, K. R. (1935). *Logik Der Forschung : Zur Erkenntnistheorie Der Modernen Naturwissenschaft*. Wien: Springer.
- Sequoiah-Grayson, S. (2007). The metaphilosophy of information. *Minds and Machines*, 17(3), 331–344.
- Sommerville, I. (2007). *Software engineering* (8th ed.). Harlow: Addison-Wesley.
- Strawson, P. (1964). Identifying reference and truth-value. *Theoria*, 30, 96–118.
- Szabolcsi, A. (1997). *Ways of scope taking*. Dordrecht: Kluwer Academic.
- Young, J. O. (2002). The slingshot argument and the correspondence theory of truth. *Acta Analytica*, 17(1), 121–132.