A Subjectivist Interpretation of Relevant Information

LUCIANO FLORIDI, HERTFORDSHIRE/OXFORD

1. Introduction

A frequent complaint about current theories of information[1] is that they are utterly useless when it comes to establishing the actual *relevance* of some specific piece of information. As a rule, agents assume that some content is by default an instance of information (Sperber and Wilson [1995]). What they often wonder is whether and how far that content may contribute to the formulation of their choices and purposes, the development of their decision processes and eventually the successful pursuit of their goals.

In light of this problem, this paper pursues two goals. The first is to provide a subjectivist interpretation of epistemic relevance (i.e. epistemically relevant semantic information, more on this presently), thus satisfying those critics who lament its absence and, because of it, may be sceptical about the utility of using information-theoretical concepts to tackle conceptual problems and cognitive issues in real life. The second goal is to show that such a subjectivist interpretation can (indeed must) be built on a veridical conception of semantic information, thus vindicating a strongly semantic theory of information (Floridi [2004b]) and proving wrong those critics who argue that misinformation can be relevant.

The two goals are achieved through a strategy of progressive refinements. In § 2, the distinction between system-based or causal and agent-oriented or epistemic relevance is introduced. In § 3, I discuss the most common and basic sense in which semantic information is said to be epistemically relevant. This has some serious shortcomings, so, in § 4, the basic case is refined probabilistically. The new version too can be shown to be only partly satisfactory, so in § 5 there will be a second, counterfactual revision. The limits of this version are finally overcome in § 6, where the analysis is completed by providing a conclusive, meta-informational refinement. In § 7, some of the advantages of the metatheoretical revision are

illustrated. In § 8, I briefly outline some important applications of what I shall label the subjectivist interpretation of epistemic relevance. In § 9, I return to the problem of the connection between a strongly semantic theory of information and the concept of epistemic relevance and explain why misinformation cannot be relevant. In § 10, I conclude by briefly summarising the results obtained and the possible work that lies ahead.

A final warning before starting: "information" can mean many things (Floridi [2004a]; Floridi [2005a]). In what follows, I concentrate only on information understood as semantic information about reality, i.e. factual information with an epistemic or cognitive value. A train timetable, a theory in a physics book, the map of the London underground, a police report about a road accident, the description of Peter's breakfast, the bell ringing when someone is at the door, are all typical illustrations that may be kept in mind.

2.   Epistemic vs. Causal Relevance

Following previous taxonomies by Cohen [1994] and Borlund [2003], approaches to the study of relevance can be divided into two groups, depending on whether they focus on a more system-based or a more agent-oriented concept of relevance. System-oriented theories (S-theories) usually analyse relevance in terms of *topicality, aboutness* or *matching* (how well some information matches a request), especially in the information retrieval (IR) literature, and various forms of *conditional in/dependence* (how some information can help to produce some outcome), especially in logic, probability theory, philosophy of science and AI.

Agent-oriented theories (A-theories), on the other hand, tend to analyse relevance in terms of *conversational implicature* and *cognitive pertinence*, especially in philosophy of language, pragmatics and psychology, and perceived *utility, informativeness, beneficiality* and other ways of "bearing on the matter at hand" in relation to an agent's informational needs, especially in IR literature and in epistemology. Adapting a distinction introduced by Hitchcock [1992], S-theories and A-theories may be seen to be interested mainly in *causal relevance* and *epistemic relevance* respectively.

S-theories clearly do not try to define, but rather presuppose, the fundamental concept of relevance understood as a relation between some information and an informee.[2] A similar conclusion may be reached regarding

the logical literature, which has concentrated mainly on S-theories, providing a variety of formalizations of logics for relevance-related notions such as conditional independence, subjunctive conditionals, novelty, causal change and co-variance (also known as perturbation models).[3] In this context, Weingartner and Schurz [1986] distinguish between two types of relevance, one *à la* Aristotle (a-relevance) and the other *à la* Körner (k-relevance). Their point is that "an inference (or the corresponding valid implication) is a-relevant if there is no propositional variable and no predicate which occurs in the conclusion but not in the premises. And an inference (or in general any valid formula) is k-relevant if it contains no single occurrence of a subformula which can be replaced by its negation *salva validitate*".[4] Clearly, neither a-relevance nor k-relevance addresses the problem of epistemic relevance. It is not surprising then that some years later, in a ground-breaking article on relevant properties and causal relevance, Delgrande and Pelletier [1998] could still conclude that "as mentioned at the outset, we feel that 'relevant' is a concept for which we have no deep understanding" (p. 166). They made no attempt to connect their analysis to an informee-oriented explanation of epistemic relevance. However, in an equally important work on relevance relations in propositional logic, published the year before, Lakemeyer [1997] had already tried to bridge the gap between the two kinds of relevance: "Perhaps the most distinctive feature that sets this work apart from other approaches to relevance is the subjective point of view. In particular, we try to capture relevance relations relative to the deductive capabilities of an agent. For example, two agents who are given the same information may very well differ in their opinion about whether p is relevant to q. Even the same agent may at first miss a connection between the two, which may be discovered upon further reflection. For instance, a student solving a geometry problem involving a right-angled rectangle may not see the connection to the Pythagorean Theorem." (p. 138) We shall see that this is a promising starting point.

The current situation can be summarised thus: some philosophical work has been done on several formal aspects of system-based or causal relevance, but the key question, namely what it means for some information to be relevant to some informee, still needs to be answered. We lack a foundational theory of agent-oriented or epistemic relevance. The warming up is over. The time has come to roll up our sleeves.

3.   The Basic Case

Strawson once remarked that "stating is not a gratuitous and random human activity. We do not, except in social desperation, direct isolated and unconnected pieces of information at each other." (Strawson [1964], p. 92). Rather, according to his Principle of Relevance, we "intend in general to give or add information about what is a matter of standing or current interest or concern." (p. 92). He was right, of course, and one may add that giving or adding information happens most commonly through interactions of questions and answers. So let us start from an abstract definition of the most basic case of relevant information and then a couple of examples.

It is common to assume that some information i is relevant (R) to an informee/agent a with reference to a domain d in a context c, at a given level of abstraction[5] (LoA) l, if and only if:

1) a asks (Q) a question q about d in c at l, i.e. $Q(a, q, d, c, l)$, and
2) i satisfies (S) q as an answer about d in c, at l, i.e. $S(i, q, d, c, l)$

In short:
$$R(i) \leftrightarrow (Q(a, q, d, c, l) \wedge S(i, q, d, c, l)) \qquad [1]$$

The basic idea expressed by [1] is simple: "the train to London leaves at 13.15" is relevant to Mary if and only if Mary has asked for that piece of information about train timetables in such and such circumstance and with the usual linguistic conventions, and "the train to London leaves at 13.15" satisfies her request.

Formula [1] is what we find applied by services like Amazon or eBay, when they suggest to a user a new item that might be relevant to her, given her past queries. It is also what lies behind the working of databases and Boolean searches, including Google queries.

3.1   Advantages of the basic case
The formulation provided in [1] has several advantages, which explain why it is so popular.

a) [1] explicitly identifies semantic information as the ultimate relevance-bearer. Other candidates in the literature on relevance comprise events, facts, documents, formulae, propositions, theories, beliefs, and messages, but Cohen [1994] has convincingly argued that relevance is proposi-

tional. He is largely correct, but while any proposition may be interpreted informationally, not all semantic information (e.g. a map) is propositional, so [1] simply brings to completion his reduction.

b) [1] takes into account the informee's interests by explicitly making the relevance of *i* depend on her queries. No semantic information is relevant *per se*, relevance being an informee-oriented concept, as anyone who has been listening to airport announcements knows only too well. This move is crucial, since it means that causal relevance can be better understood if the informee is considered part of (i.e., is embedded in) the mechanism that gives rise to it. More explicitly, this means grounding relations of causal relevance on relations of epistemic relevance.

c) [1] couples relevance and the domain *d* about which, the context *c* in which, and the LoA *l* at which the relevant information is sought. Relevance is *situational* (Borlund [2003]): the same informee can find the same information relevant or irrelevant depending on *d*, *c* and *l*.

d) [1] analyses relevance *erotetically*, in terms of logic of questions and answers (Groenendijk [2003]), and this is a strength, since it is a standard and robust way of treating semantic information in information theory (Shannon and Weaver [1949 rep. 1998]), in information algebra (Kohlas [2003]) and in the philosophy of information (Floridi [2004a]). Note that the class of questions discussed excludes those which are "loaded".[6]

e) [1] also seeks to provide an *objective* sense of relevance insofar as *i* is not any information, but only the information that actually satisfies *q* at some LoA *l*.

f) Finally, [1] constrains the amount of *subjectivity* involved in the analysis of relevance. This is achieved by assuming that the agent *a* in [1] is a *type of rational agent* which satisfies the so-called Harsanyi doctrine (Harsanyi [1968]). This point deserves some comments.

According to the Harsanyi doctrine, also known in game theory as the "common prior assumption", if two or more rational agents share a set of beliefs (the common prior assumption) about the possible state of the world, expressed by means of a probability distribution over all possible states, then – if they receive some new information about the world and if they update their set of beliefs by making them conditional (Bayesian learning) on the information received – they obtain the same revised probability (the posterior probability). So, if their new, updated beliefs differ, the conclusion is that this is because they have received different information. As Aumann

[1976] synthetically put it: "differences in subjective probabilities should be traced exclusively to differences in information".

The model is both famous and controversial. In our case, it can be used not as an abstract, if still phenomenologically reliable, description of agents' behaviour, but as a definition of what an idealised yet not unrealistic rational agent should be. The proposal is to define *a* as belonging to the class of (rational) agents who, if they share the same information about the probable realization of an event, should hold the same beliefs about it (they reach the same subjective probability assignments). This allows one to treat differences in beliefs among rational agents, and hence in their querying processes, as completely explainable in terms of differences in their information.[7] In game theory, this is called reaching consistent alignment of beliefs.

To conclude, the connection between the informee-oriented and the query-satisfaction-based features explains that [1] supports a *subjectivist* interpretation of epistemic relevance in terms of the degree of *a*'s interest in *i*. It is the sense in which one speaks of a subjectivist interpretation of probability, and should not be mistaken for any reference to the idiosyncratic inclinations of an empirical epistemic agent or their phenomenological analysis, as can be found e.g. in Schutz [1970].

3.2   Limits of the basic case

Common sense and scientific literature thus provide a good starting point, namely [1]. Despite its popularity and several advantages, however, the basic case is severely limited. One main shortcoming is that [1] is brittle, in that it is forced to declare *i* irrelevant when condition Q (*a, q, d, c, l*) is not satisfied. Obviously, even if *a* does not ask *q*, *i* (understood, following [1] as the answer to *q* about *d* in *c* at *l*) may still be highly relevant to *a*. This is what researchers and salesmen alike find distressing.

4.   A probabilistic revision of the basic case

The first step is to revise [1] by making more explicit the relation between *i* and *q*. We can then move from a rigid double-implication to a more flexible, functional relation between the degree of relevance and the degree of probability of the two conditions concerning the questioning and the answer.[8]

Call *A* the degree of *adequacy* of the answer, that is, the degree in which *i* satisfies *q* about *d* in *c* at *l*. One can define *A* as precisely as one wishes by adapting the statistical concept of *validity*. *Validity* is the combination of *accuracy* and *precision*, two other technical concepts also borrowed from statistics.[9] We shall say that *i* is an *adequate* answer to *q* insofar as it is a *valid* answer to *q*, that is, insofar as it is both an *accurate* and a *precise* answer to *q*.

We can now make [1] more resilient by considering the probability that *a* may ask *q* *and* the probability that *i* may answer *q* adequately. Unfortunately, the probability of asking a question is unrelated to the probability of receiving an adequate answer (or life would be much easier), so the two events are independent and their conjunction translates into a simple multiplication. By adopting this refinement we obtain:

$$R\,(i) = \mathrm{P}(Q\,(a, q, d, c, l)) \times \mathrm{P}(A\,(i, q, d, c, l)) \qquad [2]$$

4.1 Advantages of the probabilistic revision

[2] combines the advantages of [1] with the possibility of talking about degrees of epistemic relevance (not just Boolean quantities) and adequacy. This is coherent with a broader informational approach: in [2], the more likely *a* is to ask *q* and the more adequate *i* is as an answer to *q*, the more relevant *i* becomes to *a*.

4.2 Limits of the probabilistic revision

The main disadvantage of [2] is that the epistemic relevance of *i* decreases too rapidly in relation to the decrease in the probability of *Q*, and it becomes utterly counterintuitive in some cases. Realistically, the informee *a* cannot be considered omniscient, even if *a* is assumed to be so modal-logically (Floridi [2006]). The world is informationally opaque to *a*, at least empirically, so *a* may often fail to request the information that would actually be epistemically relevant to her, seen from a sort of God's-eye perspective. What happens when the probability that *a* may ask *q* is less than 1? As Figure 1 shows, in [2] there are four possible trends, since *R* tends towards 0 or 1 depending on whether both P(*Q*) and P(*A*) tend towards 0 or 1. Three out of four cases in [2] are realistic and unproblematic. But when P(*Q*) tends to 0 while P(*A*) tends to 1, we re-encounter the counterintuitive collapse of epistemic relevance already seen in § 1.2.c: *i* is increasingly irrelevant

epistemically because it is increasingly unlikely that *a* may ask *q*, even when the adequacy of *i* is made increasingly closer, or equal, to 1.
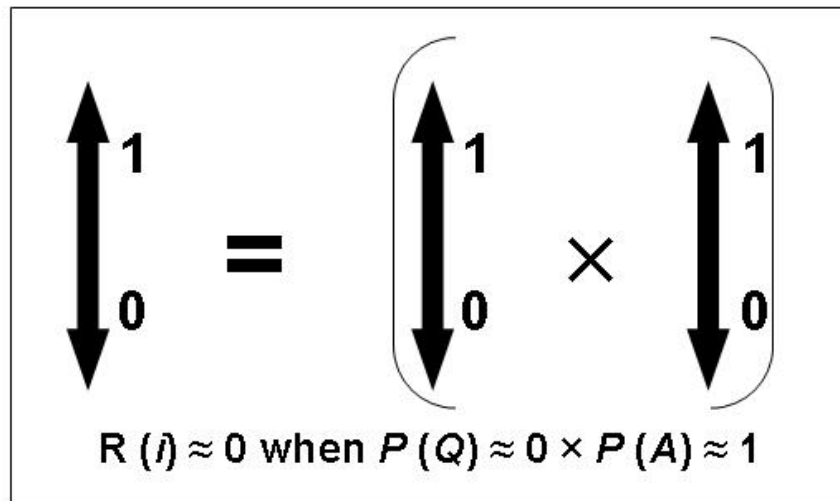

Figure 1: Four trends in formula [2]. The highlighted case is the problematic one.

5.   A counterfactual revision of the probabilistic analysis

The collapse can be avoided by revising [2] counterfactually. Instead of analysing the probability that *a* might ask *q*, one needs to consider two scenarios:

- the case in which *a* asks *q*, i.e. $P(Q) = 1$, and
- the case in which *a* does not but might ask *q*, i.e. $0 \leq P(Q) < 1$.

In the former case, the only variable that counts is the probability that *i* might be adequate. In the latter case, one can consider the probability that *a* would (have) ask(ed) *q* if *a* were (had been) sufficiently informed. Using the standard symbol "□→" for the counterfactual implication and simplifying a bit our notation by omitting (*q, d, c, l*), we obtain:

$$R(i) = \begin{cases} P(A(i)) & \text{if } P(Q(a)) = 1 \\ \\ P(Ia(i) \,\Box\!\!\to Q(a)) \times P(A(i)) & \text{if } 0 \leq P(Q(a)) < 1 \end{cases} \qquad [3]$$

The second line in [3] states that the epistemic relevance of *i* is a function of the probability that *i* might be an adequate answer to *q* times the probability that *a* would ask *q* if *a* were sufficiently informed about the availability of *i*.

5.1 Advantages of the counterfactual revision

The advantages of [3] are all the advantages of [1] and [2] plus the further advantage of solving the problem of the opacity of epistemic relevance, seen in § 4.2, and its corresponding collapse.

5.2 Limits of the counterfactual revision

The main limit of [3] may be labelled the counterfactual paradox of semantic information and is not avoidable without further revising the approach. According to [3], assuming, for the sake of simplicity, that $P(A(i, q, d, c, l))$ = 1, i would be maximally relevant epistemically only if the probability is also 1 that, if a had been informed that i was the answer, then a would have asked q to obtain i. But this conditional reminds one of *Meno's Paradox*.[10] For, if a had held i in the first place, *strictly speaking a* would not have needed to ask q to obtain i, so it is not true that a would have asked q had he held i. It follows that [3] *largely* fails to deliver a good analysis of epistemic relevance. "Strictly speaking" and "largely" are emphasised because, in practice, i would be epistemically relevant if a is assumed to be looking not for new information but for *confirmation*: a may ask q even if a already knows that i is the answer, if a wishes to be reassured that i is indeed the answer. Yet double-checking procedures are insufficient to rescue the analysis, for the complete reduction of relevance to confirmation would work as a *reductio ad absurdum*.

6. A metatheoretical revision of the counterfactual analysis

The solution is to bypass the paradox by revising [3] metatheoretically.[11] One can still rely on a's rationality to gauge the epistemic relevance of i to a herself without providing the actual content of i but only some information about its availability. For if a had been informed that new information (*ni*) about d was available, insofar as a would then have asked a question to retrieve i, it follows that i would have been correspondingly more or less epistemically relevant to a. Now, a simple way of constructing *ni* is by changing the LoA l. For example, if a had been informed that something had changed regarding the schedule of the meeting (higher LoA), a would probably have asked what had changed about it, and the information that

the meeting had been cancelled (lower LoA) would then be correctly analysed as highly epistemically relevant to *a*. In this way we obtain:

$$R(i) = \begin{cases} \mathrm{P}(A(i, q, d, c, l_m)) & \text{if } \mathrm{P}(Q(a, q, d, c, l_m)) = 1 \\[2em] \mathrm{P}(\mathrm{I}a(ni, d, l_n) \,\square\!\!\rightarrow Q(a, q, d, c, l_m)) \times \mathrm{P}(A(i, q, d, c, l_m)) & \\ & \text{if } 0 \leq \mathrm{P}(Q(a, q, d, c, l_m)) < 1 \end{cases} \qquad [4]$$

or, by simplifying our notation:

$$R(i) = \begin{cases} \mathrm{P}(A(i, l_m)) & \text{if } \mathrm{P}(Q(a, l_m)) = 1 \\[1.5em] \mathrm{P}(\mathrm{I}a(ni, l_n) \,\square\!\!\rightarrow Q(a, l_m)) \times \mathrm{P}(A(i, l_m)) & \text{if } 0 \leq \mathrm{P}(Q(a, l_m)) < 1 \end{cases} \qquad [5]$$

A final refinement can now complete the analysis. In most cases, *a* is not informed that *ni* is available. Rather, *a* may only be informed that *ni might* be available. So, instead of analysing the probability that *a* would ask *q* about *d* in *c* at l_m if *a* were informed that new information *ni* is available about *d* at l_n, one should consider, more realistically, the case in which *a* is informed that there is a probability P > 0 that there might be new information *ni* about *d* at l_n, that is, $\mathrm{P}(\mathrm{I}a\mathrm{P}(ni, l_n) \rightarrow Q(a, l_m))$. Note the scope of the two probabilities: the formula should not be interpreted as a problematic case of second order probability (Gaifman [1988]), as if the counterfactual depended on the probability of the probability of *a* being informed. It is actually *a* who is informed about the probability of *ni*. The revised formula, with the usual simplifications, is:

$$R(i) = \begin{cases} \mathrm{P}(A(i, l_m)) & \text{if } \mathrm{P}(Q(a, l_m)) = 1 \\[1.5em] \mathrm{P}(\mathrm{I}a\mathrm{P}(ni, l_n) \,\square\!\!\rightarrow Q(a, l_m)) \times \mathrm{P}(A(i, l_m)) & \text{if } 0 \leq \mathrm{P}(Q(a, l_m)) < 1 \end{cases} \qquad [6]$$

[6] synthesises the subjectivist interpretation of epistemic relevance.

7. Advantages of the metatheoretical revision

The availability of new information about d, retrievable at a higher LoA, is like a sealed envelope for a: a is informed that new information is available inside it, but does not hold the specific informational content (compare this to the message "you have mail" sent by an email client). In this way, no version of *Meno's paradox* arises and one can also account for the prima facie obligation that collaborative or informee-friendly informers may have towards a. The trite answer "I didn't tell you because you didn't ask", offered when someone fails to provide some epistemically relevant information, is now easily shown to be disingenuous. For either a should be assumed to be in a standing state of querying about (i.e., as being interested in) i, in which case the informer has a prima facie obligation to provide a with i even if a did not explicitly ask for it. Imagine the case in which Peter, a friend of Mary's, knows that she has lost her job, but that she has not yet been informed about this. It would be safe to assume Mary to be in a standing state of querying about such piece of information, so Peter, as a collaborative informer, has a prima facie obligation to inform her. Or a may simply be assumed to be reasonable enough to ask the appropriate question to obtain i, if provided with sufficient metainformation about the availability of i. In which case, the informer may have the prima facie obligation to provide at least enough metainformation about the availability of new information. Peter has at least the prima facie obligation to tell Mary that something might have happened regarding her job. Either way, not being explicitly asked by the informee fails to be a proper justification for the (informee-friendly) informer's silence.

A last, important advantage to be highlighted is that [6] is easily translatable into a Bayesian network, which then facilitates the computation of the various variables and subjective probabilities.

To summarise, [6] is easily implementable as a Bayesian Network. It explains why a collaborative informer has a prima facie epistemic obligation to inform a about i, or at least about its availability when the informer does not know what i amounts to, even if the informee does not ask for i. As we shall see in the next section, this is the fundamental assumption behind the juridical concept of relevant information. It is also what may generate conflicts in medical ethics, when epistemically relevant information may or may not be shared with all interested parties.

8.   Some illustrative cases

As anticipated, the previous analysis is compatible with a large variety of widespread usages of the concept of relevant information, to which it provides a unified, conceptual foundation. We have just seen the deontological and Bayesian contexts. Three other examples will suffice to illustrate the point and show how the conceptual ingredients found in [6] also occur in the literature on relevance, even if unsystematically.

The idea of interpreting relevant information erotetically was already exploited by Cohen [1994]. It is common in computer science and information science, where relevant information is broadly treated as "information whose subject matter matches that of a query" (Choo et al. [2000]).

The connection between relevance, probability and counterfactual inference is drawn, although not too clearly, in jurisprudence. For example, the *U.S. Federal Rules of Evidence 401. Article IV. Relevancy and its limits* states that "'Relevant evidence' means evidence having any tendency to make the existence of any fact that is of consequence to the determination of the action more probable or less probable than it would be without the evidence." Essentially, the law of evidence treats epistemic relevance as a relation between an informee *a* and two pieces of information *p* and *q*, such that it renders *p* (e.g. information about the involvement of an agent in a crime) more probable to *a* because of the occurrence of *q* (e.g. information about the time and location of an agent when the crime was perpetrated) either by itself, or in connection with other pieces of information (e.g. information about means of transportation).

Finally, in pragmatics, relevance theory (Sperber and Wilson [1995]) states that "In relevance-theoretic terms, an input is relevant to an individual when its processing in a context of available assumptions yields a positive cognitive effect. A positive cognitive effect is a worthwhile difference to the individual's representation of the world – a true conclusion, for example. *False conclusions are not worth having* (emphasis added). [...] Intuitively, relevance is not just an all-or-none matter but a matter of degree. [...] Thus, relevance may be assessed in terms of cognitive effects and processing effort:
Relevance of an input to an individual
a. other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time.

b. other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time." (Wilson and Sperber [2004], p. 608).

Although "relevance" is used in relevance theory as a technical term,[12] it is easy to see how several elements in the previous quotation can also be found included in [6], especially the informee-oriented, context-based, query-driven nature of relevance. The improvements encapsulated in [6] are threefold:

1.   semantic information (not just some linguistic item) is explicitly identified as the relevance-bearer;

2.   point (a) above is still assumed but it is now translated into a's (counterfactual) interest in asking q to obtain i, expressed by a's query. This translation no longer requires the problematic specification of what may count as "positive cognitive effects";

3.   point (b) above is replaced by degrees of probability of obtaining i, since [6] entirely decouples the degree of epistemic relevance of i from the degree of cognitive (or computational) obtainability of i. It seems counterintuitive to assume that "the greater the processing effort expended, the lower the relevance of the input to the individual at that time". Indeed, if it weren't for the technical use of "relevance" stressed above, one might argue exactly the opposite: *ceteris paribus*, some times it is precisely those bits of information more difficult to obtain (access, process etc.) that are the most epistemically relevant.[13]

A fundamental consequence of both the pragmatic approach (see the quotation above) and the subjectivist interpretation (see [6]) is that false semantic content fails to be relevant at all (for a different view see Dodig-Crnkovic [2006]). This is the next point to be discussed.

9.   Misinformation cannot be relevant

It is easy to be confused about both "relevance" and "misinformation". Regarding the former, we now have a clear analysis; regarding the latter, elsewhere (Floridi [2005b]) I have shown that misinformation is "well-formed and meaningful data (i.e. semantic content) that is false".[14] If we analyse epistemic relevance in terms of cognitive efforts, clearly misinfor-

mation makes no worthwhile difference to the informee/agent's representation of the world. On the contrary, it is actually deleterious. If the train leaves at 13.15, being told that it leaves at 14.25 is a nuisance to say the least. Likewise, if we endorse [6], clearly no rational informee/agent would be interested in receiving some misinformation as an answer to her query.[15] That one might not know whether the answer counts as information is an entirely different problem, one that involves trust, the reliability of both sources and methods of information processing and of course sceptical issues. That misinformation may turn out to be useful in some serendipitous way is also a red herring. False (counterfeit) banknotes may be used to buy some goods, but they would not, for this reason, qualify as legal tender. Likewise, astrological data may, accidentally, lead to a scientific discovery but they are not, for this reason, epistemically relevant information. Of course, there are many ways in which misinformation may be indirectly, inferentially or metatheoretically relevant, yet this is not what is in question here. The student who answers "Napoleon" to the question "who fought at Thermopylae?" has said something false and hence uninformative and *a fortiori* epistemically irrelevant to someone who asked the question in order to be informed about the battle, although his answer is informative about, and hence might be epistemically relevant to someone interested in assessing, the student's historical education. It is because of this distinction that the domain, context and the level of abstraction at which one is evaluating epistemic relevance need to be kept clear and fixed in the course of the analysis. If they are not, the outcome is conceptual carnage.

In the end, the previous discussion shows that we are on the right track. The pragmatic and the subjectivist interpretation of what may count as communicationally or epistemically relevant semantic information coherently converge on the same conclusion, even if they come from different perspectives: had *a* known that *i* was actually a piece of misinformation she would not have asked *q* in order to obtain *i* in the first place. Misinformation is not worth the effort, according to the pragmatic theory. It is unworthy of a rational agent's interest, according to the subjectivist interpretation. These are two sides of the same coin.

10. Conclusion

Agents require a constant flow and a high level of processing of relevant information in order to interact successfully among themselves and with the environment in which they are embedded. Standard theories of information are silent on the nature of relevant semantic information. In this paper, a subjectivist interpretation of relevance has been developed and defended. It is based on a counterfactual and metatheoretical analysis of the degree of relevance of (some semantic information) i to a rational informee/agent a as a function of the accuracy of i understood as an answer to a question q, given the probability that q might be asked by a. The interpretation, synthesised in [6], vindicates the strongly semantic theory of information, according to which semantic information encapsulates truth. It has been shown to be able to account satisfactorily for several important applications and interpretations of the concept of relevant information. Finally, the interpretation provides the missing foundation for a general theory of relevance. It constitutes the hub for several other theories of relevance already developed in the literature. And it is a hub that can be easily expanded by other modules. Two are worth stressing in this conclusion. First, [6] is easily combined with theories of belief upgrade. This is crucial, since the latter can explain how degrees of relevance may be dynamically upgraded following the evolution of a's background information and beliefs and feedback loops. Second, [6] is perfectly compatible with subjectivist interpretations of probability and Bayesian learning. Clearly these are implications and applications that will be worth developing.[16]

References

Bar-Hillel, Y. 1964, *Language and Information: Selected Essays on Their Theory and Application* (Reading, Mass – London: Addison-Wesley).

Bar-Hillel, Y., and Carnap, R. 1953, "An Outline of a Theory of Semantic Information" repr. in Bar-Hillel [1964], pp. 221-74.

Borlund, P. 2003, "The Concept of Relevance in IR", *Journal of the American Society for Information Science and Technology*, 54 (10), 913-925.

Bowles, G. 1990, "Propositional Relevance", *Informal Logic*, (90), 65-77.

Bremer, M., and Cohnitz, D. 2004, *Information and Information Flow – An Introduction* (Frankfurt – Lancaster: Ontos Verlag).

Choo, C. W., Detlor, B., and Turnbull, D. 2000, *Web Work: Information Seeking and Knowledge Work on the World Wide Web* (Dordrecht; Boston: Kluwer Academic Publishers).

Cohen, J. 1994, "Some Steps Towards a General Theory of Relevance", *Synthese*, 101, 171-185.

Colburn, T. R. 2000, *Philosophy and Computer Science* (Armonk, N.Y.: M.E. Sharpe).

Cooper, W. S. 1971, "A Definition of Relevance for Information Retrieval", *Information Storage and Retrieval*, 7, 19-37.

Cover, T. M., and Thomas, J. A. 1991, *Elements of Information Theory* (New York; Chichester: Wiley).

Crestani, F., Lalmas, M., Van Rijsbergen, C. J., and Campbell, I. 1998, "Is This Document Relevant?... Probably: A Survey of Probabilistic Models in Information Retrieval", *ACM Computing Surveys*, 30 (4), 528-552.

Delgrande, J. P., and Pelletier, J. 1998, "A Formal Analysis of Relevance", *Erkenntnis*, 49(2), 137-173.

Devlin, K. J. 1991, *Logic and Information* (Cambridge: Cambridge University Press).

Dodig-Crnkovic, G. 2005, "System Modeling and Information Semantics", *Proceedings of the Fifth Promote IT Conference*, Borlänge, Sweden, edited by Janis Bubenko, Owen Eriksson, Hans Fernlund, and Mikael Lind (Studentlitteratur: Lund).

Dodig-Crnkovic, G. 2006, "Model Validity and Semantics of Information" in *Model-Based Reasoning. Science and Engineering Abduction, Visualization, and Simulation*, edited by Lorenzo Magnani (London: King's College Publications).

Dongen, S. V. 2006, "Prior Specification in Bayesian Statistics: Three Cautionary Tales", *Journal of Theoretical Biology*, 242(1), 90-100.

Fetzer, J. H. 2004, "Information, Misinformation, and Disinformation", *Minds and Machines*, 14(2), 223-229.

Floridi, L. 2004a, "Information" in *The Blackwell Guide to the Philosophy of Computing and Information*, edited by L. Floridi (Oxford – New York: Blackwell), 40-61.

Floridi, L. 2004b, "Outline of a Theory of Strongly Semantic Information", *Minds and Machines*, 14(2), 197-222.

Floridi, L. 2005a, "Information, Semantic Conceptions Of", *Stanford Encyclopedia of Philosophy.* Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/information-semantic/>.

Floridi, L. 2005b, "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., and Sanders, J. W. 2004, "The Method of Abstraction" in *Yearbook of the Artificial – Nature, Culture and Technology, Models in Contemporary Sciences*, edited by Massimo Negrotti (Bern: Peter Lang), 177-220.

Gaifman, H. 1988, "A Theory of Higher-Order Probabilities" in *Causation, Chance and Credence*, edited by Brian Skyrms and William Harper (London, Ontario: University of Western Ontario Press),

Gärdenfors, P. 1978, "On the Logic of Relevance", *Synthese*, 37, 351-367. Reprinted in *Philosophy of Probability*, edited by J.-P. Dubucs, Philosophical Studies Series 56, Kluwer, Dordrecht 1993, pp. 35-54.

Greisdorf, H. 2000, "Relevance: An Interdisciplinary and Information Science Perspective", *Informing Science*, 3(2), 67-71.

Groenendijk, J. 2003, "Questions and Answers: Semantics and Logic" in *2nd Colognet-Elsnet Symposium – Questions and Answers: Theoretical and Applied Perspectives*, edited by R. Bernardi and M. Moortgat (Amsterdam: OTS).

Harsanyi, J. 1968, "Games with Incomplete Information Played By 'Bayesian' Players – Parts 1, 2, 3", *Management Science*, 14, 159-182, 320-134, 486-502.

Hintikka, J., and Suppes, P. (ed.) 1970, *Information and Inference* (Dordrecht: Reidel).

Hitchcock, D. 1992, "Relevance", *Argumentation*, 6(2), 251-270.

Israel, D., and Perry, J. 1990, "What Is Information?" in Hanson [1990], pp. 1-28.

Jones, D. S. 1979, *Elementary Information Theory* (Oxford: Clarendon Press).

Kohlas, J. 2003, *Information Algebras: Generic Structures for Inference* (London: Springer).

Lakemeyer, G. 1997, "Relevance from an Epistemic Perspective", *Artificial Intelligence*, 97 (1-2), 137-167.

LePore, E., and Smith, B. C. (ed.) 2006, *The Oxford Handbook of Philosophy of Language* (Oxford: Clarendon Press).

McCarthy, J., and Hayes, P. J. 1969, "Some Philosophical Problems from the Standpoint of Artificial Intelligence" in *Machine Intelligence*, edited by D. Michie and B. Meltzer (Edinburgh: Edinburgh University Press), 463-502.

Mizzaro, S. 1996, "Relevance: The Whole (Hi)Story", *Technical Report UDMI/12/96/RR (Dec.)* Dipartimento di Matematica e Informatica, Universita' di Udine, Italy).

Polani, D., Martinetz, T., and Kim, J. T. 2001, "An Information-Theoretic Approach for the Quantification of Relevance", *Proceedings of the 6th European Conference on Advances in Artificial Life*, edited by Jozef Kelemen and Petr Sosík (Springer), 704-713.

Polani, D., Nehaniv, C., Martinetz, T., and Kim, J. 2006, "Relevant Information in Optimized Persistence Vs. Progeny Strategies", *Artificial Life X : Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, edited by Luis Rocha, Larry Yaeger, Mark Bedau, Dario Floreano, Robert Goldstone, and Alessandro Vespignani (MIT), 337-343.

Saracevic, T. 1970, "The Concept Of "Relevance" In Information Science: A Historical Review" in *Introduction to Information Science*, edited by T. Saracevic (New York: Bower Company).

Schroder, J. 1992, "Korner's Criterion of Relevance and Analytic Tableaux", *Journal of Philosophical Logic*, 21(2), 183-192.

Schutz, A. 1970, *Reflections on the Problem of Relevance. Edited, Annotated, and with an Introduction by Richard M. Zaner* (New Haven & London: Yale University Press).

Sequoiah-Grayson, S. forthcoming, "The Metaphilosophy of Information".

Shannon, C. E., and Weaver, W. 1949 rep. 1998, *The Mathematical Theory of Communication* (Urbana: University of Illinois Press). Foreword by Richard E. Blahut and Bruce Hajek.

Sperber, D., and Wilson, D. 1995, *Relevance: Communication and Cognition* 2nd edition (Malden, MA ; Oxford: Basil Blackwell).

Tishby, N., Pereira, F., and Bialek, W. 1999, "The Information Bottleneck Method", *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, edited by B. Hajek and R. S. Sreenivas (The University of Illinois), 368–377.

Walton, D. N. 1991, "Critical Faults and Fallacies of Questioning", *Journal of Pragmatics*, 15, 337-366.

Weingartner, P., and Schurz, G. 1986, "Paradoxes Solved by Simple Relevance Criteria", *Logique et Analyse*, 29, 3-40.

Wilson, D., and Sperber, D. 2004, "Relevance Theory" in *The Handbook of Pragmatics*, edited by Laurence R. Horn and Gregory L. Ward (Malden, MA: Blackwell), 607-632.

Yus, F. 2006, "Relevance Theory Online Bibliographic Service". http://www.ua.es/personal/francisco.yus/rt.html

Ziv, Y. 1988, "On the Rationality of 'Relevance' and the Relevance of 'Rationality'", *Journal of Pragmatics*, 12, 535-545.

Notes

1. For an overview see Bremer and Cohnitz [2004] and Floridi [2004a].

2. The problem is accurately described in Crestani et al. [1998].

3. "A specific 'entity' (such as an action, training sample, attribute, background proposition, or inference step) is irrelevant to a task in some context if the appropriate response to the task does not change by an unacceptable [sic] amount if we change the entity in that context. Otherwise, we view that entity as (somewhat) relevant to the task. This view is explicitly stated in the paper by Galles and Pearl, which deals with causality and where a perturbation corresponds to a material change in the physical world." Subramanian et al. [1997], p. 2.

4. The adequacy of the Körner criterion of relevance for propositional logic has been proved by Schroder [1992].

5. The analysis of relevance also depends on the level of abstraction (Floridi and Sanders [2004]) at which the process of assessment is conducted. A level of abstraction may be seen as the precise specification of the way in which some information is being accessed and processed, cf. the analysis of "the point of view" according to which something is relevant in Cohen [1994].

6. A question Q is loaded if the respondent is committed to (some part of) the presupposition of Q (Walton [1991], 340) e.g. "how many times did you kiss Mary?" which presupposes that you did kiss Mary at least once.

7. Two further consequences are that (i) rational agents cannot possess exactly the same information and agree to disagree about the probability of some past or future events. In fact, they must independently come to the same conclusion, and (ii) they cannot surprise each other informationally.

8. Bowles [1990] follows a similar strategy to explain probabilistically the relation of relevance in propositional inferences.

9. *Accuracy* is the degree of conformity of a measure or calculated parameter to its actual (true) value. *Precision* (also called *reproducibility* or *repeatability*) is the degree to which further measurements or calculations show the same or similar results.

10. Plato, *Meno* 80d-81a.

11. This solution is partly adopted in information theory by Tishby et al. [1999], who "define the relevant information in a signal $x \in X$ as being the information that this signal provides about another signal $y \in Y$. Examples include the information that face images provide about the names of the people portrayed, or the information that speech sounds provide about the words spoken." Note that what they treat as "relevance" is really a quantitative relation of structural conjunction, which can be considered a necessary condition for semantic relevance, but should not be confused with it.

12. "Relevance here is a technical term (though clearly related to the natural language homonym), whereby an intepretation is relevant only in cases where the cognitive cost of processing the event which demands the attention of the agent is outweighed by the cognitive benefits of that processing (where benefits include deriving or strengthening new assumptions, and confirming or rejecting previous assumptions). 'Optimal relevance' states that the first interpretation which crosses the relevance threshold is the right one; that is, that the first relevant interpretation the addressee arrives at is the one the speaker intended to communicate." (Emma Borg, *Intention-Based Semantics*, in Lepore and Smith [2006], p. 255).

13. Ziv [1988] has argued that relevance theory needs to be supplemented by a theory of rationality of causal relations, in other words, what in this paper has been called causal relevance (following Hitchcock [1992]) and the assumption of a rational agent.

14. "Disinformation" is misinformation purposefully conveyed to mislead the receiver into believing that it is information.

15. This is consistent with the truth requirement established in Cohen [1994].

16. I am grateful to the organisers of the 30[th] Wittgenstein Symposium, especially Herbert Hrachovec and Alois Pichler, and to the participants to the meeting, especially Fred Dretske, for the opportunity to discuss the contents of this paper and for their feedback during the presentation. The reader interested in the topic may wish to know that a much longer, more detailed and slightly more technical version of this paper is forthcoming in *Erkenntnis*.