

Understanding Epistemic Relevance

Luciano Floridi

Abstract Agents require a constant flow, and a high level of processing, of *relevant* semantic information, in order to interact successfully among themselves and with the environment in which they are embedded. Standard theories of information, however, are silent on the nature of epistemic relevance. In this paper, a subjectivist interpretation of epistemic relevance is developed and defended. It is based on a counterfactual and metatheoretical analysis of the degree of relevance of some semantic information i to an informee/agent a , as a function of the accuracy of i understood as an answer to a query q , given the probability that q might be asked by a . This interpretation of epistemic relevance vindicates a strongly semantic theory of information, according to which semantic information encapsulates truth. It accounts satisfactorily for several important applications and interpretations of the concept of relevant information in a variety of philosophical areas. And it interfaces successfully with current philosophical interpretations of causal and logical relevance.

1 Introduction

A frequent complaint about current theories of information¹ is that they are utterly useless when it comes to establish 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

¹ For an overview see Bremer and Cohnitz (2004) and Floridi (2004a).

L. Floridi (✉)

Department of Philosophy, School of Humanities, University of Hertfordshire,
Hatfield AL10 9AB, UK

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

purposes, the development of their decision processes and eventually the successful pursuit of their goals.

The complaint is not to be underestimated. Questions of relevance affect many critical contexts, from the most mundane transactions to scientific experiments, from medical diagnoses to juridical procedures. And yet, the complaint may seem unfair, for no theory of information, from the most purely syntactical² to the most strongly semantical,³ was ever meant to cast any light on the phenomenon of relevance. This is true but, unfortunately, critics may still retort that they have at least a normative point. Information theories should care more about the relevance-related features of what they model as information. If they don't, this is not only their problem but also a good reason to disregard them when informational needs become increasingly pressing.

This “normative” objection easily morphs into a full-blooded dilemma. On the one hand, theories that formalise syntactical or structural properties of information⁴ rely on probability theory, they are statistical in nature and their pervasive applications are scientifically sound. Yet these theories abstract from any semantic feature, relevance included, and hence they seem inconsequential for the investigation of further epistemological and communication issues depending on it. On the other hand—the objection continues—there are philosophical theories that seek to capture the most salient semantic properties of information, through a variety of techniques, from situation semantics to the semantics of possible worlds or a modified calculus of probabilities.⁵ But if they end up making the concept of semantic information encapsulate that of true content (well-formed and meaningful data qualify as information only if they are also true⁶), then they are mistaken. For any theory that imposes a truth condition on the concept of semantic information cannot therefore explain how some misinformation (semantic content actually false) may still be relevant. The result is that current theories are either irrelevant or mistaken. The only way forward—the objection concludes—may be to analyse semantic information in terms of well-formed and meaningful data, without including any further truth constraint,⁷ and then trying to understand relevance in these terms. This is, however, inconsistent with the most accredited theories of relevance, according to which falsities are irrelevant (more on this in Sect. 9). Obviously something has to go, but it is unclear what.

In light of these problems, I shall pursue two goals in this paper. The first is to provide a subjectivist interpretation of epistemic relevance (i.e. epistemically relevant semantic information), thus satisfying those critics who lament its

² The classic reference is Shannon and Weaver (1949, rep. 1998), see Jones (1979) for an introduction.

³ The list includes: Bar-Hillel and Carnap (1953), Bar-Hillel (1964), Hintikka and Suppes (1970), Israel and Perry (1990), and Floridi (2004b).

⁴ In this sense, theoretical information theory is a branch of probability theory, and applied information theory a branch of engineering, see Cover and Thomas (1991).

⁵ For an overview see Floridi (2004a).

⁶ This is argued in Floridi (2005b), see Sequoiah-Grayson [forthcoming] for a recent defence.

⁷ Defenders of the alethically neutral nature of information include Devlin (1991); Colburn (2000); Fetzer (2004); Dodig-Crnkovic (2005); the latter two criticise Floridi (2004b).

absence and, because of it, may be skeptical about the utility of using information-theoretical concepts to tackle philosophical 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. This means showing that the second horn of the dilemma outlined above is actually blunt. That is what has to go.

The two goals are achieved through a strategy of progressive refinements. In Sect. 2, the distinction between system-based or causal and agent-oriented or epistemic relevance is introduced. In Sect. 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 Sect. 4, the basic case is refined probabilistically. The new version too can be shown to be only partly satisfactory, so in Sect. 5 there will be a second, counterfactual revision. The limits of this version are finally overcome in Sect. 6, where the analysis is completed by providing a conclusive, meta-informational refinement. In Sect. 7, some of the advantages of the metatheoretical revision are illustrated. In Sect. 8, I briefly outline some important applications of what I shall label the subjectivist interpretation of epistemic relevance. In Sect. 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 Sect. 10, two common objections are answered; their discussion helps to clarify further the proposed theory. In Sect. 11, 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, 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

Most of the literature on relevance⁸ does not so much interpret the nature of the phenomenon as actually use the corresponding concept for specific applications. For example, relevant information is essential in many epistemological analyses, especially in the so-called *relevant alternatives theory*, but the question about what exactly makes some information relevant is normally left unanswered

⁸ See for example Yus (2006), a bibliography online on relevance theory in pragmatics and related disciplines. For recent review articles on relevance in information science see Greisdorf (2000) and the very useful Borlund (2003). Philosophical accounts of relevance include Gärdenfors (1978); Cohen (1994); Lakemeyer (1997); and Delgrande and Pelletier (1998), all works that have influenced the research for this paper.

(Moser 2002). True, we encounter plenty of hints about what it might mean for some information p to be relevant, yet these normally amount to more or less implicit endorsements of a variety of commonsensical and pretheoretical understandings of the concept, which fail to provide a conceptual foundation and a shareable, explanatory frame. To make things worse, the theories of relevance currently available come from a variety of fields that often do not speak to each other: several branches of computer science and of information science, statistics and probability theory, AI, cognitive science, epistemology, logic, philosophy of language, linguistics and jurisprudence. The risk of gerrymandering is obvious. It was already stressed by Cohen (1994).

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. The problem is accurately described in Crestani et al. (1998): “The concept of relevance is arguably the fundamental concept of IR. In the above-presented model, we purposely avoid giving a formal definition of relevance. The reason behind our decision is that the notion of relevance has never been defined precisely in IR. Although there has been a large number of attempts towards a definition of the concept of relevance (Saracevic 1970; Cooper 1971; Mizzaro 1996), there has never been agreement about a unique and precise definition. A treatment of the concept of relevance is outside the scope of this paper and we will not attempt to formulate a new definition or even accept a particular already existing one. What is important for the purpose of our survey is to understand that relevance is a relationship that may or may not hold between a document and a user of the IR system who is searching for some information: if the user wants the document in question, then we say that the relationship holds.”

Similar conclusions 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).⁹ 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*”.¹⁰ 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

⁹ “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.

¹⁰ The adequacy of Körner criterion of relevance for propositional logic has been proved by Schroder (1992).

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 abstraction¹¹ (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:

$$R(i) \leftrightarrow (Q(a, q, d, c, l) \wedge S(i, q, d, c, l)) \quad (1)$$

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. Finally, understood as in [1], relevance is the semantic counterpart of the algebraic concept of marginalization, $(\varphi, x) \mapsto \varphi^{\downarrow x}$, in information algebra.¹²

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 propositional. 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,

¹¹ 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).

¹² Many thanks to Jürg Kohlas for having called my attention to this equivalence.

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”.¹³
- 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.¹⁴ In game theory, this is called reaching consistent alignment of beliefs.

¹³ 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.

¹⁴ 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.

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. Three of the main shortcomings are:

- a) [1] is insufficiently explanatory, since the relation between i and q is left untouched: how *adequate* must i be as an answer to q in order to count as *relevant* information?
- b) [1] is too coarse, for it fails to distinguish between degrees of relevance and hence of epistemic utility of the more or less relevant information. It might be relevant to a that the train has been delayed, but it is even more relevant to a that the train has been delayed by one hour instead of ten minutes, yet [1] can not capture this distinction.
- c) [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.¹⁵

Call A the degree of *adequacy* of the answer, that is, the degree to 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.¹⁶ We shall say

¹⁵ Bowles (1990) follows a similar strategy to explain probabilistically the relation of relevance in propositional inferences.

¹⁶ *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.

that i is an *adequate* answer to q insofar as it is a *valid* answer to q , that is, insofar as it is an answer to q both *accurate* and *precise*.

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:

$$R(i) = P(Q(a, q, d, c, l)) \times P(A(i, q, d, c, l)) \quad (2)$$


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 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 Fig. 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.

Fig. 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 “ $\Box \rightarrow$ ” for the counterfactual implication and simplifying a bit our notation by omitting (q, d, c, l), we obtain:

$$R(i) = \begin{cases} P(A(i)) & \text{if } P(Q(a)) = 1 \\ P(Ia(i) \Box \rightarrow Q(a)) \times P(A(i)) & \text{if } 0 \leq P(Q(a)) < 1 \end{cases} \quad (3)$$

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 Sect. 4.2, and its corresponding collapse.

5.2 Limits of the Counterfactual Revision

The first limit requires some fine-tuning: it concerns a potentially circular use of counterfactuals. The metalinguistic interpretation of counterfactuals *à la* Goodman requires a reference to a characterization of the *relevant* initial conditions that would make a counterfactual true. Thus, using Quine’s classic example, “if Julius Caesar had been in charge of U.N. Forces during the Korean War, then he would have used (a) nuclear weapons or (b) catapults”, one can make sense of the general scenario and hence of the two alternatives only if the *relevant* domain knowledge is available. This is obviously circular and would not do. However, the more standard Stalnaker-Lewis semantics allows an interpretation of counterfactuals in terms of possible worlds close to the actual world within a specific similarity-structure in a logical space. This is good news because, although references to agents’ epistemic interests might be brought to bear in this case as well, the move is unnecessary: the closeness or similarity function may be computed on the basis of a purely extensional analysis or probabilistic projection, starting from the given, actual world. The Stalnaker-Lewis approach is far from being uncontroversial or devoid of

problems, but it does allow one to avoid the circularity of having to establish what information is metatheoretically or contextually relevant to the agents in order to evaluate some further relevant information.

The second limit of [3] may be labelled the counterfactual paradox of semantic information and it 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*.¹⁷ For, if a had held i in the first place, *strictly speaking* a would not have been in any need 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.¹⁸ 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:

¹⁷ Plato, *Meno* 80d–81a:

“Meno: And how will you investigate, Socrates, that of which you know nothing at all? Where can you find a starting-point in the region of the unknown? Moreover, even if you happen to come full upon what you want, how will you ever know that this is the thing that you did not know?”

Socrates: I know, Meno, what you mean; but just see what a tiresome dispute you are introducing. You argue that man cannot enquire either about that which he knows, or about that which he does not know; for if he knows, he has no need to enquire; and if not, he cannot; for he does not know the very subject about which he is to enquire.”

¹⁸ 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.

$$R(i) = \begin{cases} P(A(i, q, d, c, l_m)) & \text{if } P(Q(a, q, d, c, l_m)) = 1 \\ P(IaP(ni, d, l_n) \square \rightarrow Q(a, q, d, c, l_m)) \times P(A(i, q, d, c, l_m)) & \text{if } 0 \leq P(Q(a, q, d, c, l_m)) < 1 \end{cases} \quad (4)$$

or, by simplifying our notation:

$$R(i) = \begin{cases} P(A(i, l_m)) & \text{if } P(Q(a, l_m)) = 1 \\ P(IaP(ni, l_n) \square \rightarrow Q(a, l_m)) \times P(A(i, l_m)) & \text{if } 0 \leq P(Q(a, l_m)) < 1 \end{cases} \quad (5)$$

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, $P(IaP(ni, l_n) \square \rightarrow Q(a, l_m))$. 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:

$$R(i) = \begin{cases} P(A(i, l_m)) & \text{if } P(Q(a, l_m)) = 1 \\ P(IaP(ni, l_n) \square \rightarrow Q(a, l_m)) \times P(A(i, l_m)) & \text{if } 0 \leq P(Q(a, l_m)) < 1 \end{cases} \quad (6)$$

[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. This is most convenient. Concentrating on the interesting case in which $0 \leq P(Q(a, q, d, c, l_m)) < 1$, and for variables N (corresponding to $IaP(ni, l_n)$), A (corresponding to $P(A(i, q, d, c, l_m))$) and Q (corresponding to $P(IaP(ni, l_n) \square \rightarrow Q(a, l_m))$), given some interpretation of the conditional probabilities:

$P(n)$, the probability of variable N ;

$P(q | n)$, the probability of variable Q given n ;

$P(q | \bar{n})$, the probability of variable Q given \bar{n} ;

$P(a)$, the probability of variable A ;

$P(q | n)$, the probability of variable Q given n ;

$P(q | \bar{n})$, the probability of variable Q given \bar{n} ;

$P(r | a, q)$, the probability of variable R given a and q ;

$P(r | a, \bar{q})$, the probability of variable R given a and \bar{q} ;

$P(r | \bar{a}, q)$, the probability of variable R given \bar{a} and q ;

$P(r | \bar{a}, \bar{q})$, the probability of variable R given \bar{a} and \bar{q} ;

we obtain the corresponding value of the joint probability function:

$$P(r) = \sum_{N,A,Q} P(R|A, Q) \cdot P(Q|N) \cdot P(N) \cdot P(A)$$

For example, if we assume¹⁹ that the probability of new information occurring is very high (0.9), that if there is new information, the probability that the agent will ask a question about it is also very high (0.9), and that, if the question is asked, the probability that the answer is adequate is also high (0.8), then the node probability (see Table 1) yields a degree of epistemically relevant semantic information of 0.6868. This can be graphically shown as in Fig. 2.

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

¹⁹ The mere assumption of these values is justified here because this is only an illustrative example. The identification of the right set of Bayesian priors is a hard problem faced by any analysis of real-life phenomena. Of course, the formulation of a prior distribution over the unknown parameters of the model should be based on the available data (including subjective beliefs) about the modelled phenomena, yet this is easier said than done, see for example Dongen (2006).

Table 1 Example of a node probability table for a Bayesian interpretation of epistemic relevance

N	Q	A	R	$P(N)$	$P(Q N)$	$P(A)$	$P(R Q, A)$
Yes	Yes	Yes	Yes	0.9	0.9	0.8	1.0
Yes	No	Yes	Yes	0.9	0.1	0.8	0.1
No	Yes	Yes	Yes	0.1	0.1	0.8	1.0
No	No	Yes	Yes	0.1	0.9	0.8	0.1
Yes	Yes	No	Yes	0.9	0.9	0.2	0.1
Yes	No	No	Yes	0.9	0.1	0.2	0.0
No	Yes	No	Yes	0.1	0.1	0.2	0.1
No	No	No	Yes	0.1	0.9	0.2	0.0
Yes	Yes	Yes	No	0.9	0.9	0.8	0.0
Yes	No	Yes	No	0.9	0.1	0.8	0.9
No	Yes	Yes	No	0.1	0.1	0.8	0.0
No	No	Yes	No	0.1	0.9	0.8	0.9
Yes	Yes	No	No	0.9	0.9	0.2	0.9
Yes	No	No	No	0.9	0.1	0.2	1.0
No	Yes	No	No	0.1	0.1	0.2	0.9
No	No	No	No	0.1	0.9	0.2	0.1

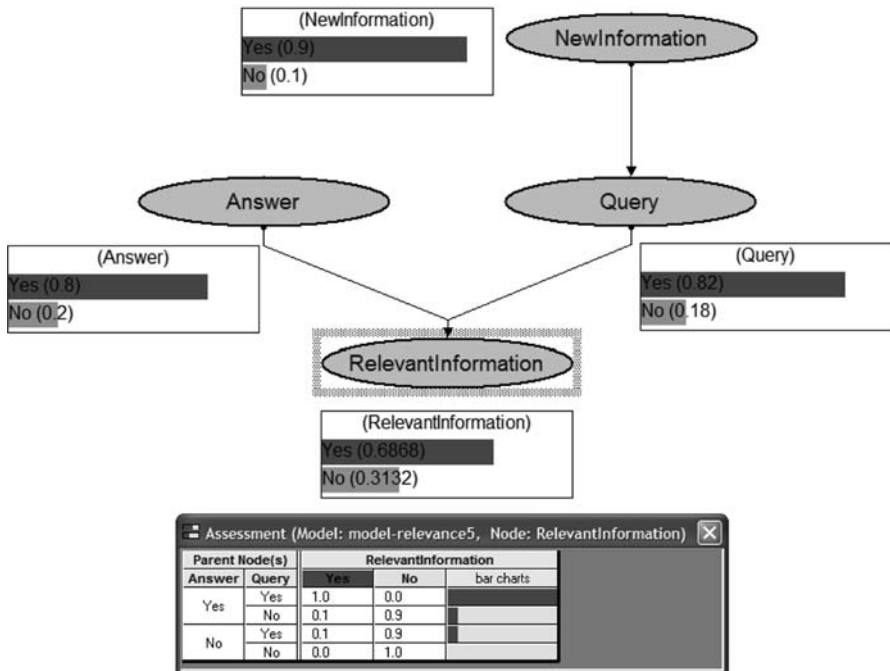


Fig. 2 Example of an implementation of [6] by means of a Bayesian network. The variables N , A , Q and R have been given more intuitive names. The assessment in the smaller window shows the conditional probabilities of variable $R = \text{"RelevantInformation"}$. The graphic was produced with MSBNx Version 1.4.2, Microsoft Research's Bayesian network authoring and evaluation tool

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,²⁰ it is easy to see how several elements in the previous quotation can also be found included in

²⁰ “Relevance here is a technical term (though clearly related to the natural language homonym), whereby an interpretation 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).

[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.²¹

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”.²² If we analyse epistemic relevance in terms of cognitive efforts, clearly misinformation 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.²³ 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 skeptical issues. That misinformation may turn out to be useful in some serendipitous way is also a red

²¹ 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.

²² “Disinformation” is misinformation purposefully conveyed to mislead the receiver into believing that it is information.

²³ This is consistent with the truth requirement established in Cohen (1994).

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 a 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 Two Objections and Replies

The subjectivist interpretation of epistemically relevant information is not entirely uncontroversial and has been subject to some criticisms,²⁴ which may be summarised into two objections. Each of them cast further light on the proposal. Fortunately, they are both answerable.

10.1 The Interpretation is Complete: No Relevant Semantic Information for Semantically Unable Agents

The first objection argues that the subjectivist interpretation in [6] relies too heavily on the semantic capacities of the agent *a*. How can some specific information *i* (say the location of some organic debris) be relevant to an amoeba, if the amoeba cannot ask any question, not even in principle (as when one might wish to say: “imagine an amoeba could speak, then ...”). There is plenty of information that is epistemically

²⁴ I am summarising here a variety of questions and objections raised, with some consistency, at several meetings where I presented the ideas laid out in this article (for a complete list see the acknowledgments). I discussed with Fred Dretske the first objection during the 30th Wittgenstein Symposium. I am very grateful to Jeremy Seligman for having suggested the second objection and for our conversations about it, which greatly helped me to clarify the issue.

relevant to semantically unable agents, but [6] fails to take this into account, so it is incomplete at best.

The objection deserves three replies. First, as it was specified at the outset, the proposed interpretation concerns *semantic* information, not any kind of information. The location of some food is a vitally important yet physical fact, which might be conveyed by a message, and hence be transformed into semantic information, but it is not in itself semantic information. For agents entirely incapable of any semantic interactions, such translation is impossible or rather meaningless, and so should be any analysis of epistemically relevant semantic information as a consequence. In the example, the amoeba and its environment do not interact semantically and this is just the end of the story. No such difficulties arise with animals with higher cognitive faculties, capable of interacting with other animals and the environment semantically.

The second reply is that one might think of some cases of *relevant facts* or uninterpreted *signals/data* as being interpretable in terms of hard-wired questions, posed by the agents involved, which receive equally hard-wired answers, offered by the environment, where the latter are interpretable as affordances. Take the interactions between the environment and organisms even simpler than amoebae, for example heliotropic plants, such as snow buttercups or sunflowers, which react to light. In this case, the “question” may be seen to have been hard-wired by evolutionary processes into the motor cells located in flexible segments of the plant specialized in pumping potassium ions into nearby tissues (thus changing the turgor pressure) reversibly. The direction of the sun works as the factual answer. The result is a diurnal motion of flowers or leaves. The hard-wiring of questions and answers is just another way of conceptualising utility functions associated with behavioural strategies. The important difference is that the “questioning” is entirely externalized. This then allows a quantitative approach to Shannon-type information and its relevance that is coherent with the subjectivist interpretation of semantic information proposed in this paper, as is clearly shown in Polani et al. (2006) (see also Polani et al. 2001), who have recently tried to use the bottleneck method (Tishby et al. 1999) to “study a scenario where a multicellular colony has to trade-off between *utility* of strategies for investment in persistence or progeny and the (Shannon-type) relevant information necessary to realize these strategies.” (p. 337).

The third reply concerns AI. Heliotropic plants act as analogue computational agents. Do digital computational agents fare any better when it comes to epistemically relevant information as defined in [6]? Not really. Identifying relevant information is just a case of the “frame problem” (McCarthy and Hayes 1969), a notorious obstacle for any form of artificial intelligence developed so far. Hard- or soft-wiring “questions”—which might enable artificial agents to act as if they could process relevant semantic information—is a good but limited strategy. The subjectivist interpretation of relevant information cannot really work for artificial agents simply because the latter are not semantic engines. The same intelligence that leads an agent to *ask* questions is what allows that agent to spot subjectively relevant information. Amoeba, sunflowers and Turing Machines have no semantic abilities, no intelligence, no curiosity, ask no questions and cannot therefore be used as counterarguments against [6] because nothing can be semantically relevant to

them. On the contrary, [6] may provide a criterion to discriminate between intelligent, semantically able agents and non-intelligent, semantically unable ones. For one may run a sort of reverse-engineered Turing test, in which the agents being tested do not answer, but rather ask questions. One may then check whether the human and the artificial agent's capacities to grasp the relevance of the answers provided are sufficiently indistinguishable. If one day they will be, that is, if one day artificial agents will perform on average as well as humans in asking questions and dealing with the relevance of the semantic information they receive as answers, some people (this author included) will consider that nothing short of a miracle, which will usher in a dramatically new era in human history.

10.2 The Interpretation is Sound: Rationality Does Not Presuppose Relevance

The second objection still focuses on the role played by the agent *a*, but with a different strategy. We have seen that throughout the paper the agent *a* was assumed to be rational. Without this normative condition, *a* would not be pursuing her interests consistently and she would not be asking those questions that she considers most helpful in order to gather the sort of information that would be useful to satisfy her needs and interests. At the train station, Mary would ignore the announcements, would not ask when her train leaves, would carelessly buy a ticket to some place, wait at a random platform and jump on the first train she fancies. The whole world would be a matter of indifference to her and hence irrelevant to her. So far so good. The objection, however, is not that the informee *a* in [6] should not be assumed to be a sufficiently rational agent, but that this necessary condition begs the question. For surely a rational agent must also be one who is capable of discriminating between more or less relevant alternatives, that is, pieces of semantic information, in order to formulate and guide her choices. But then, we are back to square one: we are explaining epistemic relevance by presupposing it, and, even worse, by doing so necessarily, i.e. inescapably.

There are two replies to this objection. The first turns out to be unsatisfactory but helps one to understand the second, which actually answers the objection.

As a first attempt, one may accept the circularity but reject its viciousness. In order to clarify what it means for some information *i* to be epistemically relevant to *a*, one has to refer to a rational agent *a* to whom *i* is relevant, hence *a* must be an agent capable of detecting and weighing relevance. Yet, what is being presupposed is not some pre-formed or innate quantity of relevant information in the head of the informee, but just the presence in *a* of some relevance-detecting capacity, implicit in the description of *a* as a rational agent. Some information *i* is relevant to *a* if and only if *a* behaves in such a way as to show that *i* is relevant to her and this includes asking questions to retrieve *i*. We are using a rational agent *a* to identify information epistemically relevant to *a* in the same way as a chemist may use litmus to detect acid substances. As Polani et al. (2001) put it: "One can regard the decision system or agent as an estimator for the optimal action to take" and hence the corresponding, relevant information that determines it.

The previous reply would be entirely convincing were it not for the fact that the chemist does not stop at the successful litmus test, but actually explains its success

through a well-supported theory on the nature and nomic behaviour of acids as substances that are proton donors and accept electrons to form ionic bonds. It is this scientific explanation that allows her to avoid any circularity. Litmus does not tell her what it means for a substance to be an acid; it merely tells her whether it is an acid. If this were all she could say about acidity it would be just a matter of circular and conventional definition, not much better than the *virtus dormitiva* used to explain why camomile makes one sleep. Likewise, if all we could offer were a reduction of epistemically relevant information to *a*'s capacity to detect it, we would be merely shifting the problem. We would be saying that some information *i* is epistemically relevant to an agent *a* because *a* can detect it as such, and *a* detects *i* as such because it is so. Clearly, the viciousness of the circularity kicks in, unless we have something comparable to the chemist's safety exit. The first reply is unsatisfactory. Fortunately, a second reply comes to its rescue.

It is insufficient to rely on the rationality of the agent if this is just another way of speaking of epistemic relevance, but rational agents need not be defined in terms of their capacities to detect relevant information. The "surely" used in formulating the objection above is merely rhetorical and unjustified. On the contrary, the standard way in which a rational agent *a* is defined refers only to the following four conditions:

- R.1) information input: *a* can perceive her environment through sensors;
- R.2) action output: *a* can act upon her environment through effectors in order to modify and/or control it;
- R.3) preferences about outcomes: *a* is not indifferent with respect to the new states of the environment and of herself that she may obtain (outcomes); she has desires, likes, dislikes and so forth, which may make her privilege some outcomes over others;
- R.4) optimization of outcomes with respect to preferences: *a* acts in her own best interest, always preferring the outcomes that have the highest *expected utility*.

No condition contains, or requires a reference to, epistemically relevant information; the Harsanyi doctrine, regarding multiagent systems, is also relevance-independent; and the logic of rational choice in general or Bayesian learning in particular does not rely on epistemic relevance. So *it seems* that, by referring to rational agents in [6], one can validly reduce epistemically relevant semantic information to a set of conditions none of which is based in its turn on relevance or other cognate concepts.

"It seems" because our skeptical opponent may think he has a last card up his sleeve. He may retort that irrational agents are normally defined in terms of *ir/* relevant information. In particular, he may stress that it is a standard move in the literature on rational choice to consider an agent *a* irrational if

- Ir.1) *a* deliberately ignores (what *a* considers to be) relevant information; and/or
- Ir.2) *a* deliberately uses (what *a* considers to be) irrelevant information (here the typical case is known as the sunk cost fallacy).

His objection is now that (Ir.1) and (Ir.2) are not cases of mistaken evaluation (the clauses “deliberately” and the bracketed conditions take sufficient care of this) but rather evidence that the very concept of ir/rational agent presupposes that of ir/relevant semantic information and therefore cannot ground it.

This final version of the objection is based on a conflation between irrationality and incoherence, which unfortunately seems to be widespread in the literature on rational choice. We have seen above that two or more agents are rational if they satisfy the Harsanyi doctrine. A single agent *a* is then rational if she satisfies conditions R.1–R.4. If R.1 or R.2 are unsatisfied, *a* is not an agent; and if R.3 or R.4 are unsatisfied, *a* is an agent but not rational. This is textbook material. Once all these conditions are satisfied, the rational agent may be embedded in [6] to yield a non-circular definition of epistemically relevant information. And once it is established that some information *i* is epistemically relevant to the rational agent *a*, then this can be used to establish whether *a* is a *coherent* agent. Now coherence may be loosely discussed in terms of rationality but, given the standard definition above, it is important to be more precise and careful, by looking at the course of actions of our agent through time.

Suppose *a* is irrational because *a* fails to satisfy R.3 or R.4. According to [6], there is no information that is epistemically relevant to *a*. So there is no problem about *a* deliberately ignoring relevant information, a behaviour which is an effect of *a*'s irrationality, not a cause. The agent *a* is irrational because of other shortcomings, unrelated to the notion of relevance.

Suppose *a* is rational because *a* initially satisfies R.3 and R.4. Then, according to [6], there is some information *i* that is epistemically relevant to *a*, and some other information *l* that is epistemically irrelevant. But then, if, after this stage, *a* goes on deliberately ignoring *i* or deliberately using *l*, this means that her course of action becomes incoherent: in particular, she no longer acts in her own best interest, an interest that was defined by the role she played in identifying epistemically ir/relevant information in the earliest stage of the analysis using [6]. At this point, the agent is not entirely irrational, but she is certainly incoherent: she is initially rational, insofar as her behaviour may lead to the definition of epistemically ir/relevant information, but she is then irrational, insofar as her subsequent behaviour does not take her previous behaviour and its outcome into account. In other words, she is incoherent and her incoherence results from a comparison between her behaviour before and after the identification and gathering of some information *i* as epistemically ir/relevant.

In either case, there is no circularity. We are defining coherence in terms of deliberate consideration/usage of epistemically ir/relevant information and ir/relevant information in terms of rational agency, and this in terms of R.1–R.4 (stand-alone agent) and the Harsanyi doctrine (multiagent systems). What is gained, rather than a fallacious definition, is a useful way to quantify the mismatch between the rational behaviour of asking *q* in [6] first, and the irrational behaviour of disregarding the answer to it or using some information that does not satisfy [6], later. Basically, the higher the probability that *a* might have asked *q*, the more irrational *a* is if she then deliberately ignores the resulting *i* or uses *l*.

11 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. It is defensible in terms of its completeness and soundness. 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.

Acknowledgements The first time I discussed the topic of a theory of epistemic relevance was during a talk I gave at the University of Regensburg (Regensburg, Germany 9 November 2005). I owe to Rainer Hammwoehner and Hans Rott not only the kind invitation but also the conceptual pressure that made me start working on this paper. A first version of the paper was then presented at the Department of Communication Science of the University of Salerno (Fisciano, Italy, 10 May 2006), and I wish to thank Roberto Cordeschi for that opportunity and the feedback I received in that occasion. The paper was further improved and discussed at the "Workshop on Information Theories", organized by Juerg Kohlas and Giovanni Sommaruga at Fribourg University (Münchenwiler, Switzerland, 17–18 May 2006). They, the attendees, and especially Rolf Haenni and Jeremy Seligman provided some very helpful comments. A new version was the subject of an invited talk at the Department of Philosophy of the University of Siena (Siena, 14 June 2006), where I took advantage of a long discussion with Claudio Pizzi on secondorder probabilities. This led to a paper presented at a seminar organised by the Computer Science Department of Mälardalen University (Västerås, Sweden, September 2006), where I was kindly invited by Gordana Dodig Crnkovic. The discussion with the participants and especially with Gordana, Susan Stuart and Lars-Göran Johansson generated several improvements. The issue of hardwired questions was discussed there. The final version of the article then became the ISI Samuel Lazerow Memorial Lecture I delivered at the University of Arizona (Tucson, 8 February 2007). I am grateful to Don Fallis for the invitation and to the Research Group on the History and Philosophy of Information Access, the School of Information Resources and Library Science and The International Visitors Fund for the kind support. The last opportunity I had to discuss this paper was as an invited lecture at the 30th Wittgenstein Symposium and at the 48th Boston Colloquium for Philosophy of Science. Finally, I would like to acknowledge the help, useful comments and criticisms by Pia Borlund, Ken Herold, Karen Mather, Paul Oldfield and Federica Russo and the journal's anonymous referees. As usual, all the aforementioned people are responsible only for the improvements and not for any remaining mistakes.

References

- Aumann, R. J. (1976). Agreeing to disagree. *Annals of Statistics*, 4(6), 1236–1239.
- Bar-Hillel, Y. (1964). *Language and information: Selected essays on their theory and application*. Reading: Addison-Wesley.
- Bar-Hillel, Y., & Carnap, R. (1953). An outline of a theory of semantic information. rep. in Bar-Hillel (1964).
- 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., & Cohnitz, D. (2004). *Information and information flow – An introduction*. Frankfurt: Ontos Verlag.
- Choo, C. W., Detlor, B., & Turnbull, D. (2000). *Web work: Information seeking and knowledge work on the world wide web*. Dordrecht: 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: M.E. Sharpe.
- Cooper, W. S. (1971). A definition of relevance for information retrieval. *Information Storage and Retrieval*, 7, 19–37.
- Cover, T. M., & Thomas, J. A. (1991). *Elements of information theory*. New York: Wiley.
- Crestani, F., Lalmas, M., Van Rijsbergen, C. J., & 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., & 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. In J. Bubenko, O. Eriksson, H. Fernlund, & M. Lind (Eds.), *Proceedings of the Fifth Promote IT Conference*, Borlänge, Sweden. Studentlitteratur: Lund.
- Dodig-Crnkovic, G. (2006). Model validity and semantics of information. In L. Magnani (Ed.), *Model-based reasoning. Science and engineering abduction, visualization, and simulation*. 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 L. Floridi (Ed.), *The Blackwell guide to the philosophy of computing and information* (pp. 40–61). Oxford, New York: Blackwell.
- Floridi, L. (2004b). Outline of a theory of strongly semantic information. *Minds and Machines*, 14(2), 197–222.
- Floridi, L. (2005a). Semantic conceptions of information. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. <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., & Sanders, J. W. (2004). The method of abstraction. In M. Negrotti (Ed.), *Yearbook of the artificial-nature, culture and technology, models in contemporary sciences* (pp. 177–220). Bern: Peter Lang.
- Gaifman, H. (1988). A theory of higher-order probabilities. In B. Skyrms & W. Harper (Eds.), *Causation, chance and credence*. London, Ontario: University of Western Ontario Press.
- Gärdenfors, P. (1978). On the logic of relevance. *Synthese*, 37, 351–367.
- 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 R. Bernardi & M. Moortgat (Eds.), *2nd Colognet-Elsnet Symposium – Questions and answers: Theoretical and applied perspectives*. Amsterdam: OTS.
- Harsanyi, J. (1968). Games with incomplete information played by “Bayesian” Players – Parts 1, 2, 3. *Management Science*, 14, 159–182, 320–334, 486–502.
- Hintikka, J., & Suppes, P. (Eds.) (1970). *Information and inference*. Dordrecht: Reidel.

- Hitchcock, D. (1992). Relevance. *Argumentation*, 6(2), 251–270.
- Israel, D., & Perry, J. (1990). What is information? In P. Hanson (Ed.), *Information, language, and cognition*. Vancouver: British Columbia Press.
- 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., & Smith, B. C. (Eds.) (2006). *The Oxford handbook of philosophy of language*. Oxford: Clarendon Press.
- McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In D. Michie & B. Meltzer (Eds.), *Machine intelligence* (pp. 463–502). Edinburgh: Edinburgh University Press.
- Mizzaro, S. (1996). Relevance: The whole (Hi)story. *Technical Report UDMI/12/96/RR (Dec.)* Dipartimento di Matematica e Informatica, Università di Udine, Italy.
- Moser, P. K. (2002). *The Oxford handbook of epistemology*. Oxford, New York: Oxford University Press.
- Polani, D., Martinetz, T., & Kim, J. T. (2001). An information-theoretic approach for the quantification of relevance. In J. Kelemen & P. Sosik (Eds.), *Proceedings of the 6th European Conference on Advances in Artificial Life* (pp. 704–713). Springer.
- Polani, D., Nehaniv, C., Martinetz, T., & Kim, J. (2006). Relevant information in optimized persistence vs. progeny strategies. In L. Rocha, L. Yaeger, M. Bedau, D. Floreano, R. Goldstone, & A. Vespignani (Eds.), *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems* (pp. 337–343). MIT.
- Saracevic, T. (1970). The concept of “relevance” in information science: A historical review. In T. Saracevic (Ed.), *Introduction to information science*. 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. In R. M. Zaner (Ed.), *Annotated, and with an introduction*. New Haven: Yale University Press.
- Sequoiah-Grayson, S. forthcoming. The Metaphilosophy of Information. *Minds and Machines*.
- Shannon, C. E., & Weaver, W. (1949, rep. 1998). *The mathematical theory of communication*. Urbana: University of Illinois Press.
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Malden: Basil Blackwell.
- Strawson, P. (1964). Identifying reference and truth-value. *Theoria*, 30, 96–118. Reprinted in P. Strawson, (1971). *Logicolinguistic papers* (pp. 75–95). London: Methuen.
- Subramanian, D., Greiner, R., & Pearl, J. (1997). The relevance of relevance. *Artificial Intelligence*, 97(1–2), 1–5.
- Tishby, N., Pereira, F., & Bialek, W. (1999). The information bottleneck method. In B. Hajek & R. S. Sreenivas (Eds.), *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing* (pp. 368–377). The University of Illinois.
- Walton, D. N. (1991). Critical faults and fallacies of questioning. *Journal of Pragmatics*, 15, 337–366.
- Weingartner, P., & Schurz, G. (1986). Paradoxes solved by simple relevance criteria. *Logique et Analyse*, 29, 3–40.
- Wilson, D., & Sperber, D. (2004). Relevance theory. In L. R. Horn, & G. L. Ward (Eds.), *The handbook of pragmatics* (pp. 607–632). Malden: Blackwell.
- 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.