Minds and Machines manuscript No.
(will be inserted by the editor)

Relevant Information and Relevant Questions

Comment on Floridi's "Understanding Epistemic Relevance"

Received: date / Accepted: date

Prologue: Why Relevance?

Floridi's chapter on relevant information bridges the analysis of "being informed" (which itself depends on a theory of strongly semantic information, and presupposes an analysis of semantic information that encapsulates truth) with the analysis of knowledge as "relevant information that is accounted for." Yet, unlike the work that precedes the development of a theory of subjective relevance, and unlike the work that depends on such a theory, the proposed analysis of relevant information in terms of what an agent might ask, were he or she informed of the availability of a certain piece of information, looks rather uncontroversial. It doesn't spark a controversy—as the veridicality thesis did—or even contain an implicit critique on the present state of a discipline—as the network theory of account does for mainstream (post-Gettier) epistemology. All we find is a certain amount of clarification (epistemic relevance is relevance *for an agent*, relevance depends on context, level of abstraction, ...), and a number of incremental improvements (the relevant issues aren't exhausted by the questions that are actually asked). Why then choose this specific chapter as the focus of a critical appraisal?

One reason for devoting my comment to the topic of relevance is related to my own interest in the question of how knowledge and information should be related. In particular, what does it mean for a theory of knowledge to put *information first*, and what does it mean for information to be a stepping stone to knowledge? If we want a viable information-based epistemology, every component of our theory should function as intended, both in isolation and in interaction with the other components. By scrutinising the proposed analysis of epistemic relevance, I want to find out whether one specific component of a broader theory delivers its goods. The slogan for this motivation might therefore be: "care about the details."

Another reason for taking a closer look at the notion of epistemic relevance is that it allows me to be at the same time constructive and critical; critical because I identify some crucial flaws in Floridi's analysis of epistemic relevance; constructive because I give an outline of a solution, and thus contribute to one of the core projects within the philosophy of information. Here too, we can summarise this with a slogan: "progress from new or better models, not merely from counterexamples."

A final reason is that when I first read a draft of the paper on which this chapter is based I already suspected that the proposed account of subjective relevance might be incomplete, but I never made this suspicion precise. As it turns out, the worries that form the basis of the present contribution are quite remote from what I initially thought to be the problem. In the epilogue, I shall briefly comment on these earlier doubts.

Overview

The paper is structured as follows. Section 1 contains a summary of how Floridi arrived at his proposal of relevant information as information an agent might ask for, were she or he informed of its availability. In Section 2 I diagnose the main flaw in Floridi's proposal, and subsequently explain (§3) why an easy fix isn't available. In Section 4 I formulate two potential defences on behalf of (conservative revisions of) Floridi's proposal. These defences allow me to improve the initial diagnosis of why Floridi's proposal doesn't lead to a good measure of subjective relevance. An attempt to do better is given in Sections 5 and 6. In the epilogue, I conclude by offering some more general remarks on the relation between (bounded) rationality, the need to ask the right questions, and the ability to ask the right questions.

1 Epistemic relevance

According to Floridi's analysis of epistemic relevance, the relevance of a certain piece of information for a certain agent can be reduced to two independent factors: (i) how well a piece of information answers a given question; (ii) the probability that this question is asked. The main part of the chapter devoted to this topic contains an analysis and a series of successive revisions of each of these factors.

Starting from the initial proposal that a piece of information is relevant just when an agent asks a question that can be answered by that piece of information (the basic case), Floridi introduces a number of refinements and modifications. Since the flaw I want to expose in Floridi's proposal is related to its precise formulation, I here only give an informal description of the outcome of the successive revisions of the basic case. These revisions primarily broaden the scope of the information that is relevant for a given agent.

1. We shouldn't only look at questions that are actually asked, but also at questions that might be asked. This is a first part of the *probabilistic* revision.
2. In fact, we shouldn't only look at the questions an agent might ask in his or her present epistemic state, but also include those questions an agent might ask if she or he were informed of the availability (but not of the content) of a given piece of information. This is the *counterfactual* and *metatheoretical* revision.
3. Given that a certain question is asked, a piece of information isn't only relevant if it perfectly answers that question, but also if it only partially answers that question by either being incomplete or inaccurate. This proviso for partial answers motivates a second part of the *probabilistic* revision: We should consider how probable it is that a given piece of information answers a question.

2 The problem

As explained in the introduction, my main concern with the proposed analysis of epistemic relevance only bears on how it is formalised. As a matter of fact, the issue I will point out is entirely independent of the proposed revisions of the basic case: We find it in the basic (non-probabilistic) case where relevance is characterised by an equivalence as well as in the successive (probabilistic) revisions where relevance is characterised by an equality. Consider, first, the basic case. Here, Floridi writes:

> 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 (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. $S(i, q, d, c, l)$
>
> (p. 249)[1]

This analysis, he claims, is summarised by the following equivalence:

$$R(i) \leftrightarrow Q(a, q, d, c, l) \wedge S(i, q, d, c, l) \tag{BC}$$

A first thing one should notice is that the initial description includes an existential quantification "asks *a* question" (emphasis added) that is absent from the equivalence that intends to summarise the whole proposed analysis. A second thing to note is that the variables for agent, domain, context and level of abstraction occur in the *explanans* (right-hand side of the equivalence), but not in the *explanandum* (left-hand side). Both these features can easily be exploited to derive a contradiction from (BC).

[1] Unless explicitly mentioned, page-numbers refer to Floridi's "The Philosophy of Information."

Indeed, we can assume that i is relevant in virtue of

$$Q(a,q,d,c,l) \land S(i,q,d,c,l)$$

while any of the following could be the case

$$\neg\big(Q(a,q',d,c,l) \land S(i,q',d,c,l)\big)$$
$$\neg Q(a',q,d,c,l) \land S(i,q,d,c,l)$$
$$Q(a,q,d',c',l') \land \neg S(i,q,d',c',l')$$

thereby allowing us to derive by means of (BC) the contradictory conclusion

$$R(i) \land \neg R(i) \tag{$\perp$}$$

to the effect that i is both relevant and irrelevant.

To be sure, the idea that a single piece of information can both be relevant and irrelevant seems an essential part of any theory of epistemic or subjective relevance, but ($\perp$) is hardly a good way to reflect this feature. What we want to say is that this or that piece of information is relevant for *some* agent, in *some* context, but isn't necessarily relevant for another agent, in another context.

Restoring consistency across the board is, fortunately, a fairly straightforward matter. We only need to include the relevant agent and the remaining contextual factors in the explanandum, and to reintroduce the existential quantification in the explanans.

$$\forall(i,a,d,c,l)(R(i,a,d,c,l) \leftrightarrow \exists q(Q(a,q,d,c,l) \land S(i,q,d,c,l))) \tag{BC$^\exists$}$$

What we obtain is precisely what the informal description of the basic case was meant to be in the first place.[2] Presumably, this is just how the charitable reader should have understood Floridi's presentation of the basic case. In short, something one should hardly complain about . . . except for the fact that the problem I just pointed out gets transferred to the probabilistic revisions of the basic case, where an analogous charitable reading is not readily available.

Consider, next, the first probabilistic revision (p. 252) of the basic case (the remaining arguments are left out to improve readability)

$$R(i) = \Pr(Q(q)) \times \Pr(A(i,q)), \tag{PR}$$

where the predicate Q is true for all questions q that are asked, and the relation A is true for all question/information-pairs i, q such that i adequately answers q.

As before, we can imagine that $R(i)$ is high because both factors, the probability that a question q is being asked and the probability that i is an

[2] One of the ambiguities that are removed by replacing (BC) with (BC$^\exists$) is the status of the letters $i, \ldots, l$. In the original version, they could both be understood as constants and as variables. In the revised version they are clearly variables of a *multi-sorted* language.

adequate answer to q, are high. Still, this does not preclude the possibility of there being another question q' such that

$$\Pr(Q(q)) \times \Pr(A(i,q)) \neq \Pr(Q(q')) \times \Pr(A(i,q'))$$

and hence

$$R(i) \neq R(i), \tag{$\neq$}$$

which is as much of a contradiction as ($\perp$).

3 Analysis and diagnosis

There are at least two ways to avoid this conclusion. We can drop the assumption that the probabilistic revision of the basic case should result in an equality, or we can deny the reasoning that led to ($\neq$) by pointing out that the equality only holds for a restricted range of questions. Both options are worth exploring. I start with the second.

The standard reading of an equation like (PR) is as an equality that holds in general (that is, for all i and q). This type of reading doesn't leave much room for a restriction on the admissible values of q, and even less room for a restriction that should be based on i (i.e. the questions that are somehow related to i). More exactly, unless we assume that $R(i)$ already includes an implicit restriction on the range of admissible values, we have no reason to assume that $\Pr(Q(q))$ is an admissible factor of $R(i)$, while $\Pr(Q(q'))$ isn't. Furthermore, since the restriction is meant to block the reasoning that led to ($\neq$), the only non-circular restriction is one that picks out exactly one question.[3] Consequently, if $R(i)$ includes such a restriction, we'd better make it explicit by letting R take two arguments.

$$R(i,q) = \Pr(Q(q)) \times \Pr(A(i,q)) \tag{PR1^q}$$

While formally sound, this is hardly an acceptable formalisation of the relevance of i. (PR1^q) captures at best one aspect of the relevance of i.

When we replace (PR) with an inequality (our first option),

$$R(i) \geq \Pr(Q(q)) \times \Pr(A(i,q)) \tag{PR1$^\geq$}$$

we do exploit (PR1^q): If there is a question q such that $R(i,q) = r$, then $R(i)$ is at least as high as r.

By taking the maximum (assuming it exists), we can obtain a new equality from (PR1$^\geq$):

$$R(i) = \max\{R(i,q) \mid q \text{ is a question}\} \tag{PR1$^{\max}$}$$

If we take this path, we unify our two ways of avoiding the contradictory conclusion ($\neq$). The move from (PR1$^\geq$) to (PR1$^{\max}$) depends, however, on

[3] Non-circular in the sense of not being defined in terms of the avoidance of the unwanted inequality.

the implicit but crucial assumption that the relevance of i can be reduced to the relevance of i relative to some q. When this assumption is made explicit

$$R(i) = x \text{ iff } \underline{\text{there is a } q} \text{ such that } R(i, q) = x \tag{$\exists$}$$

we immediately notice the resemblance with ($\mathrm{BC}^{\exists}$). Yet, what works for a Boolean analysis of relevance, as illustrated in this further variant of ($\mathrm{BC}^{\exists}$):

$$R(i) \leftrightarrow \exists q(R(i, q) \geq k) \tag{$\mathrm{PR}^{\exists}$}$$

doesn't necessarily work for a probabilistic analysis.

There is surprisingly much to be said in favour of an analysis of relevance along the lines of ($\mathrm{PR}^{\exists}$). Since each $R(i, q)$ includes all the features of Floridi's proposal, it has all the advantages of the counterfactual and metatheoretical revisions, and even retains several virtues of the probabilistic revision. It obviously fails to distinguish between degrees of relevance, but doesn't share any of the other limitations of (BC) and ($\mathrm{BC}^{\exists}$).

The lack of "degrees of relevance" makes this proposal similar to to that proposed by Gabbay and Woods under the heading of "agenda relevance," where:

> [R]elevance is defined over ordered triples $\langle \mathbf{I}, \mathbf{X}, \mathbf{A} \rangle$ of items of information **I**, cognitive agents **X**, and agendas **A**. (...) We shall propose that **I** is relevant for **X** with regard to his or her agenda **A** if and only if in processing **I**, **X** is affected in ways that advance or close **A**. (Gabbay & Woods 2003, 158)

More importantly, ($\mathrm{PR}^{\exists}$) is also an adequate explanation of what "relevance" means in the definition of knowledge as "relevant semantic information that is accounted for."

Let me, now, introduce some additional terminology to make clear why the assumption of ($\exists$) is problematic in the full-fledged probabilistic case. Following Floridi's usage of the term (see Chapter 8), we can say that an answer *saturates* a question just when it "erases the data deficit" (p. 189) of that question.[4] As such, the saturation of q by i is a necessary (and presumably sufficient) condition for $A(i, q)$. In a query-oriented context, we primarily care about saturation. Here, we should also care about its converse. We shall therefore say that a question q or set of questions Q *exhausts* a piece of information i just when the (combined) data deficit of the question(s) contains the deficit that can be erased by i.

Seen from one side, the i's that saturate some q are the expected outcomes of any information retrieval system. Seen from the other side, when some Q exhausts i, some $\mathsf{Q}' \subseteq \mathsf{Q}$ will presumably adequately capture a's interest in i. But if that's the case, and if Q' contains at least two questions, ($\mathrm{PR1}^{\max}$) will fail to take into account at least one question that reflects a's interest in

[4] Keep in mind that these notions only make sense in a context, at a particular LoA.

i. The latter fact can easily lead us to misevaluate the relevance of two pieces of information that, for instance, satisfy the following two conditions

$$\max\{R(i_1, q) \mid q \text{ is a question}\} = \max\{R(i_2, q) \mid q \text{ is a question}\}$$
$$\mathsf{Q}_1 \subset \mathsf{Q}_2$$

where Q_1 and Q_2 respectively exhaust i_1 and i_2, and no $\mathsf{Q}_1{}' \subset \mathsf{Q}_1$ or $\mathsf{Q}_2{}' \subset \mathsf{Q}_2$ does. For indeed, if we need more questions to exhaust i_2 than to exhaust i_1, there *might* be more in i_2 that interests a than there is in i_1.[5] Yet, by only considering the question that maximises $R(i, q)$ we apparently cannot account for our intuition that i_2 could in that case be more relevant than i_1 for a.

With this in mind, we can understand Floridi's proposal as an analysis of $R(i, q)$ (the relevance of i relative to a question q) instead of an analysis of $R(i)$ (the relevance of i), and consider his successive probabilistic, counterfactual and metatheoretical revisions as attempts to come up with a more representative (multi-)set of $R(i, q)$'s. Moreover, since the relevance of i should depend on the set of such $R(i, q)$'s, Floridi's final proposal does contribute to our understanding of epistemic relevance *simpliciter*, but, in view of the problems we reported with regard to (PR1$^{\max}$), it also fails to deliver a complete analysis.[6]

4 Two Defences

Before moving on to a new proposal, I would like to consider two potential defences against the worries I raised in the previous section. The first defence advances that the proposed measure of epistemic relevance should only apply to atomic pieces of information. The second defence advances that the problems that were identified are at least partly solved by the counterfactual and meta-theoretical revisions of (PR).

The guiding intuition behind the first defence is that atomic pieces of information can always be exhausted by a single (and presumably fairly simple) question. Once this restriction is in place, (PR1$^{\max}$) is no longer objectionable. Assuming that $R(i) = R(i, q)$, the q that is singled out will presumably be (a) a sub-question of some question q' that exhausts i, and (b) adequately reflect a's interest in i. All of this seems largely correct, but the restriction imposed on (PR1$^{\max}$) also deprives it of its interest. The problem is that we cannot simply sum $R(i)$ and $R(i')$ to compute the relevance of the complex piece of information that contains i and i', because the relevance of i and the relevance of i' need not be independent (i' may, for instance, contain information that is redundant in view of i).[7] As a result, (PR1$^{\max}$) cannot even be the starting

[5] The "might" qualification is essential since the interest of a in i_1 and i_2 is captured by subsets of Q_1 and Q_2.

[6] I assume here that (PR1$^{\max}$) is the *obvious* way of fixing Floridi's probabilistic versions, just like I assumed that (BC$^{\exists}$) was the intended reading of (BC).

[7] Atomic pieces of information shouldn't be understood in the same way as atomic propositions, as this would exclude basic disjunctive information.

point of a generally applicable measure of epistemic relevance; it is only a limiting case of a still to be given more general measure.[8]

The second defence suggests that if we use a probabilistic (and meta-theoretical) revision of ($PR1^{max}$), the thus obtained measure no longer ignores questions that shouldn't be ignored. Consider the following adaptation of Floridi's final proposal (where the expression $I_a \Pr(ni, l_n)$ stands for "a is informed of the probability that there is new information, ni, available (about a given domain, at some LoA, *etc.*)", see p. 255):

$$R(i,q,a,l_m) = \begin{cases} \Pr(A(i,q,l_m)) & \text{if } \Pr(Q(a,q,l_m)) = 1 \\ \Pr(I_a \Pr(ni,l_n) \Box\!\!\rightarrow Q(a,q,l_m)) \times \Pr(A(i,q,l_m)) & \text{if } 0 \leq \Pr(Q(a,q,l_m)) < 1 \end{cases},$$

and let $R(i)$ be defined as before by taking the maximum. The idea would then be that, as in the first defence, the question that is singled out by taking the maximum is a question that adequately reflects a's interest in i (and thus a sub-question of some question that exhausts i). The contribution of the counterfactual revision is precisely that it focuses on the questions an agent would ask when informed of the availability of i, and that this focus is sufficient to let

$$\max\{R(i,q,a,l_m) \mid q \text{ is a question}\}$$

single out the best or most efficient query, given a's interests. This idea is reinforced by the stipulations that (a) the focus is on rational agents, i.e. agents that would pick out the most appropriate question, and (b) the questions we're talking about are best seen as abstract queries rather than as specific questions.

A defence of this type is flawed for at least two reasons. The first reason is connected to the fact that q doesn't need to exhaust i to make $A(i,q,l_m)$ true. Hence, $A(i,q,l_m)$ is insensitive to redundancy in the sense that $A(i,q,l_m)$ and $A(i,q',l_m)$ may both hold even though q exhausts i but q' doesn't. A more problematic consequence of this fact is that the probabilities will favour the *easier* questions. To wit, if

$$A(i,q,l_m) \rightarrow A(i,q',l_m)$$

is valid,[9] we ought to accept

$$\Pr(A(i,q',l_m) \geq \Pr(A(i,q,l_m))$$

as well.

The second reason is that we have no reason to assume that more encompassing questions are, even given the counterfactual condition, more likely to be asked than the less encompassing ones. More exactly, if we straightforwardly identify complex questions with sets of more basic questions (and do not further worry about their internal structure), we know that the probability

[8] The question of how we should compute the sum of multiple $R(i,q)$'s (for dependent i's) will come back in a different guise in the next section.

[9] The underlying intuitive principle is that if i answers the more encompassing question, it surely also answers the less encompassing one.

that a set of questions is asked can never be higher than the probability that one of the basic questions it contains is asked. That is, if $\mathsf{Q} = \{q_1, \ldots, q_n\}$, we have

$$\Pr(I_a P(ni, l_n) \mathrel{\Box\!\!\rightarrow} Q(a, q_i, l_m)) \geq \Pr(I_a P(ni, l_n) \mathrel{\Box\!\!\rightarrow} Q(a, \mathsf{Q}, l_m))$$

for each $1 \leq i \leq n$.

Taken together, these two reasons show that $R(i, q') > R(i, q)$ doesn't imply that q' takes more advantage of the content of i than q does: The best question is neither the question that is strictly more likely to be asked, nor the question that is strictly better answered by i. Consequently, asking the right questions is something that cannot be explained in terms of the question that maximises $R(i, q)$.

5 Constraints on relevance

The main lesson of the preceding section is that the intuitively plausible principle ($\exists$) is false. If we assume that the relevance of i depends on the value of $R(i, q)$ for some q, there are only a limited number of ways of selecting such a q. If it depends on $\Pr(I_a P(ni, l_n) \mathrel{\Box\!\!\rightarrow} Q(a, q, l_m))$, on $A(i, q, l_m)$ or on their product, we end up with the problems that were exposed in the previous sections. If, by contrast, we require that q be such that (a) it exhausts i, and (b) none of its sub-questions exhausts i, then we in fact reduce the relevance of i to the probability that the (subjectively) best question will be asked. Yet, even that approach has unwelcome consequences. Let, by way of illustration, i_1 and i_2 be two pieces of information such that every question that can be answered by i_1 can also be answered by i_2, but not *vice versa*. Assume, moreover that q_1 and q_2 are, respectively, the best questions for these pieces of information, and that

$$\Pr(I_a P(ni_1, l_n) \mathrel{\Box\!\!\rightarrow} Q(a, q_1, l_m)) > \Pr(I_a P(ni_2, l_n) \mathrel{\Box\!\!\rightarrow} Q(a, q_2, l_m)),$$

which indicates that a is more interested in an answer to q_1 than in an answer to q_2. Still, this is consistent with (note the presence of q_1 in both consequents!)

$$\Pr(I_a P(ni_1, l_n) \mathrel{\Box\!\!\rightarrow} Q(a, q_1, l_m)) = \Pr(I_a P(ni_2, l_n) \mathrel{\Box\!\!\rightarrow} Q(a, q_1, l_m)),$$

which indicates that a doesn't really consider i_2 less relevant than i_1. This reveals that an implementation of ($\exists$) that is based on the best question agrees with the following two constraints:

1. If i_1 answers questions that are likely to be asked while i_2 doesn't, then (all else being equal) i_1 is more relevant than i_2.
2. If i_1 answers questions that are *not* likely to be asked while i_2 doesn't, then (all else being equal) i_1 is *less* relevant than i_2.

As a consequence, even though i_1 and i_2 may be equally useful to a, i_1 would be considered less relevant just because it also answers questions that a wouldn't ask. Whereas the second type of constraint is a standard component of definitions of epistemic justification (as in the common view that we not only want to maximise true beliefs, but also want to minimise false beliefs), it doesn't seem appropriate for relevance.[10]

Instead, the following principle seems a more accurate implementation of our attitude towards the presence of irrelevant information.

3. If i_1 answers questions that are *not* likely to be asked while i_2 doesn't, then (all else being equal) i_1 cannot be more relevant than i_2.

This constraint is motivated by the consideration that irrelevant content should not make a piece of information less relevant, but it shouldn't make it more relevant either. Consequently, a good measure of relevance should agree with the first and the third constraint, but not with the second. As I repeatedly argued, this cannot be achieved on the basis of (∃).

6 Outline of an alternative

If we move to an analysis of relevance that takes into account multiple questions, these constraints should be modified accordingly. Thus, the first constraint becomes: If i answers more questions that a might ask (in the sense of $\Pr(I_a P(ni, l_n) \square\!\!\rightarrow Q(a, q, l_m)))$ than i' does, then, for a, i is more relevant than i'. If the posing of different questions were independent, this constraint could be formalised as:

$$\sum_{j \in \mathbb{N}} (R(i, q_j, a, l_m)) > \sum_{j \in \mathbb{N}} (R(i', q_j, a, l_m)) \tag{C1}$$

with $\{q_j \mid j \in \mathbb{N}\}$ the set of all questions. Unfortunately, this isn't the case. There could, and often will be multiple questions q, q' such that $R(i, q)$ and $R(i, q')$ are both high, but depend on each other.

Such dependencies can be understood along two different (and incompatible) lines. We can think of the conditional probability $\Pr(q \mid q')$ as the probability that one would ask q given that one already asked q'. This reading of conditional probabilities easily leads to undesirable results. Indeed, if one assumes that agents wouldn't ask the same question twice, this reading entails $\Pr(q \mid q) = 0$, which is clearly false. Alternatively, we can think of the conditional probability $\Pr(q \mid q')$ as the probability that one would be interested in an answer to q, given that one is already interested in an answer to q'. On that account, we trivially have $\Pr(q \mid q) = 1$. More importantly, we can now give a dynamic interpretation of the dependence between questions q and q' as the

[10] This counterexample relies on the fact that relevance has different features when it applies to declarative information than when it applies to questions: A question can be *less relevant* (in the sense of being a worse question) than some of its sub-questions, but a piece of information is always *at least as relevant* as any of its parts.

probability that one would refrain from asking q, if one already obtained an answer to q'.

Clearly, the latter option is the kind of dependence between questions that we need to track to understand the dependence between the relevance of a piece of information relative to multiple questions. Writing down the general sum for many such dependent $R(i, q)$'s is then a tedious, but otherwise straightforward task.[11] The resulting approach is related, but also more sensible than the one hinted at the beginning of Section 4 because (a) it yields a unified approach for the relevance of complex as well as atomic pieces of information, and (b) it acknowledges the complex interaction between pieces of information and queries for information. It is precisely this complexity that an assumption like ($\exists$) ignores.

Epilogue: Relevance and Limited Rationality

The value we accord to relevant information cannot be separated from our cognitive limitations. It is precisely because our resources are limited that we should only devote attention to information that is relevant for us: Our success as cognitive agents critically depends on our ability to ask the right questions. This is why Floridi identifies relevant questions with questions that a rational agent (without further qualification) would ask (p. 262). Yet, as we shall see, certain differences between real and idealised agents makes this identification less straightforward.

When we contrast limited and ideally rational agents, the posing of questions is at least in one crucial respect different from having beliefs. This can be seen as follows. When it comes to beliefs, the development of more realistic models can be motivated by the principle that what is feasible for an ideal agent, isn't necessarily feasible for a real agent. Conversely, the normative import of idealised models can be explained with an appeal to the intuitively plausible principle that what is rational for an ideal agent will surely be rational for a real agent as well (Hawthorne & Bovens 1999, 243). At first blush, we can understand Floridi's appeal to rational agents in the explication of epistemic relevance along similar lines.[12] However, if we rely on the contrapositive version that "what is irrational for a real agent is also irrational for an ideal agent," we can readily construct a counterexample for the application of this principle to the posing of questions. Consider, first, the following conditional:

1. If it is irrational for a real agent *not* to ask a certain question, it is equally so for an ideal agent,

but (and this is the counterexample):

[11] Because it is unrelated to my final point, I'm deliberately ignoring the further dependence between $A(i, q)$ and $A(i, q')$, which obviously should also be taken into account.

[12] Granted, the assumption that the model concerns ideally rational agents need not follow from the description of rationality Floridi gives (p. 264), but the reliance on a probabilistic model surely pulls in that direction.

2. it is irrational for an agent with limited resources to ask superfluous or redundant questions, but this can be entirely unproblematic for an agent with unlimited resources.

As a consequence, we cannot entirely explain the notion of a relevant question in terms of what an ideally rational agent would do, because an ideal agent values asking the right questions, but doesn't need to value the avoidance of the wrong questions.

Crucially, this type of argument does not depend on the fact that it is harder for real agents to maximise their expected benefits than it is for idealised agents (condition R4 in Floridi's description of rational agents, p. 264). Instead, what I want to emphasise is that with respect to what counts as relevant information non-ideal agents have other preferences than ideal agents (condition R3 in the same description). Arguably, the preferences of an ideal agent may be so that as long as all relevant questions are asked, the agents expected utility isn't negatively influenced by asking further superfluous questions.

In sum: The value of relevant information (and the threat of irrelevant information) can only be understood in a context where resources are scarce. But how does this diagnosis affect the value of a probabilistic account of epistemic relevance? Here, I do not have a complete satisfactory answer.

As we have seen, the probabilities that figure in the analysis of epistemic relevance need not reflect logical relations between the actual posing of different questions (or, more neutrally, between different actual queries), but only logical relations between the questions or queries themselves. As a result, there is no worry about resources relative to the number of questions that are being asked. Such resources lie outside the scope of the model we use.

One might, however, worry that the intended interpretation of the probabilities, together with a logic of questions that is based on classical logic (as in Wiśniewski 1995), leads to a formal theory that cannot prevent the assignment of high probabilities to certain intuitively irrelevant questions. The core of this concern is that if the probability of asking a question is constrained by such an erotetic logic, the resulting probabilities need not be a good indicator of the relevance of its answers.

Such problems arise, amongst others, because a classical account of question-evocation[13] yields many intuitively irrelevant questions (De Clercq & Verhoeven 2004). For instance, the set $\Gamma = \{p \vee q, r\}$ not only evokes sensible questions like $?\{p, q\}$ or $?\{\neg p, \neg q, p \wedge q\}$, but also totally unrelated questions like $?\{t, \neg t\}$. Whether we use the erotetic notion of question-evocation to constrain $\Pr(I_a P(ni_1, l_n) \Box\!\!\rightarrow Q(a, q, l_m))$, or the implication relation between questions (which is plagued by similar irrelevancies) to constrain conditional probabilities between questions, the resulting probabilities succeed in assigning high probabilities to relevant questions, but fail to assign low probabilities to ir-

[13] Γ is a set of declarative premises, we say that a question $Q = ?\{A_1, \ldots, A_n\}$ is evoked by Γ iff (i) $\Gamma \vdash_{CL} A_1 \vee \ldots \vee A_n$, while (ii) for each A_i we have $\Gamma \nvdash A_i$.

relevant questions. In the example we gave, t would (assuming it is true) be considered relevant just because p and q are deemed relevant.

With this in mind, the often heard concern that classical logic is good for the ideal agent, but less so for the real agent appears to generalise to the realm of questions, and thus to the problem of epistemic relevance. The adoption of a non-classical account of the logical relations between questions and declarative sentences (De Clercq & Verhoeven 2004) is just one way out.

The Rejection of classical logic (or at least its erotetic extension) is not our only option. Logical notions like those of erotetic implication and question evocation should not be confused with epistemic notions like the probabilities that figure in the analysis of epistemic relevance.[14] This creates room for an alternative response. The deficiencies of Wiśniewski's logic of questions can, according to this view, be used to reject the connection between a logical analysis of questions and the probabilities we use to refer to the questions a rational agent might ask (given some counterfactual condition).[15]

This diagnosis reveals at least the following. Because relevant questions are best understood in terms of what limited rational agents would do, and because the formal modelling of limited rational agents is notoriously hard, the probabilities that figure in the different analyses of epistemic relevance presuppose a lot more than the description of rational agents suggests. This holds even though the description of rational agents (p. 264) does not presuppose agents with unlimited resources. As I read the definition, rational agents should try to maximise benefits and minimise costs, but these are requirements that have totally different implications for agents with limited resources than for agents with unlimited resources. This diagnosis remains, however, consistent with Floridi's defence that the identification of relevant questions with the questions that a rational agent would ask is non-circular. It only reveals that there is still a lot to be said about the questions a non-ideal agent should ask and (especially) the questions such an agent should *not* ask.

Acknowledgements

References

De Clercq, K. & Verhoeven, L. (2004), 'Sieving out relevant and efficient questions', *Logique et analyse* **47**(185-188), 189–216.

Fitelson, B. (2008), 'Goodman's "new riddle"', *Journal of Philosophical Logic* **37**(6), 613–643.

[14] See Fitelson (2008) for this type of diagnosis in the context of evidential support and confirmation.

[15] Such a looser connection between logical principles and probabilities need not block my arguments from Section 4, for there I only relied on logical connections between questions and sets (or conjunctions) of questions. These are principles that belong to standard classical logic; not to its erotetic extension.

Gabbay, D. & Woods, J. (2003), Agenda relevance, *in* 'Agenda Relevance. A Study in Formal Pragmatics', Elsevier, Amsterdam, pp. 155–193.

Hawthorne, J. & Bovens, L. (1999), 'The preface, the lottery, and the logic of belief', *Mind* **108**(430), 241–264.

Wiśniewski, A. (1995), *The posing of questions : logical foundations of erotetic inferences*, Kluwer Academic Publishers, Dordrecht ; Boston.