4

EXPERT KNOWLEDGE, EVIDENCE, AND INFORMATION

CONHECIMENTO POR ESPECIALISTA, EVIDÊNCIA E INFORMAÇÃO

Tommaso Piazza*

ABSTRACT – In this paper I argue that the existence of expert knowledge potentially poses a problem for Evidentialism, the view that a person's justification supervenes on the evidence this person has. An expert is the kind of person from which knowledge (or justified belief) is expected in situations in which a non-expert would normally not attain knowledge (or justified belief); so, potentially, the epistemic status of their beliefs differ even if the evidence they possess seems to be the same. A viable solution to this problem has to show that, appearances to the contrary, the expert and the non-expert, in the problematic cases, do not possess the same evidence. I propose a solution to this problem by defending a principle that specifies the conditions under which a piece of information should be counted as evidence.

KEYWORDS – Expert knowledge. Evidentialism. Evidence. Epistemic justification. Information. Perception.

RESUMO – Neste estudo, argumento que a existência do conhecimento por especialista potencialmente coloca um problema para o evidencialismo, a visão de que a justificação que uma pessoa tem sobrevem à evidência que essa pessoa possui. Um especialista é o tipo de pessoa a partir de quem o conhecimento (ou a crença justificada) é esperado em situações em que um não-especialista normalmente não atingiria o conhecimento (ou a crença justificada). Assim, potencialmente, o estatuto epistêmico das crenças deles diferem, mesmo que a evidência que possuem pareça ser a mesma. Uma solução viável a esse problema tem de mostrar que, apesar das aparências, o especialista e o não-especialista, nos casos problemáticos, não possuem a mesma evidência. Proponho uma solução a esse problema, ao defender um princípio que especifica as condições sob as quais um item de informação deveria contar como evidência.

PALAVRAS-CHAVE – Conhecimento por especialista. Evidencialismo. Evidência. Justificação epistêmica. Informação. Percepção.

* PhD. in Philosophy at the University of Florence; Professor of Philosophy at the University of Porto; Researcher C2007, co-financed FSE and POPH, of the Institute of Philosophy of the University of Porto.

Evidentialism about epistemic justification is the view that the (quantity of) justification a subject has for believing a given proposition depends on the evidence that person possesses. More in particular, the idea is conveyed by the following principle:

EJ Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* if and only if having *D* toward *p* fits the evidence *S* has at *t*. (C&F 2004: 83)

As such, this view is committal to the claim that whenever two subjects have different degrees of justification for believing the same proposition, this must depend on a difference between the evidence the first and the second possesses; that which implies, in particular, that whenever two persons believe the same proposition, and just the first knows it because she has a stronger justification, this will have to be explained in terms of a difference in the evidence both possess.

In this paper, I want to explore an implication of EJ that, on the face of the intuitive character of this principle, would seem to be rather problematic. Under a rather uncontroversial characterization, a subject *S* can be regarded as an *expert* in some domain DOM if, for some *p* ranging on some feature of the environment that is central to DOM, *S* may attain knowledge that *p* in situations in which a different subject *S1*, who is not an expert in DOM, could not attain knowledge of *p*. This different propensity to attain knowledge (form justified belief) is what in principle reveals *S*'s expertise in DOM, and *S1*'s lack of expertise in this domain. Central, to this characterization, is the idea of a situation's being the *same*. It is because *S* and *S1* are in the *same situation* that their different epistemic performances reveal their different expertise in the domain. The expert is precisely the one who knows (or form a justified belief) when the non-expert, at the very best, may just attain true belief. If this characterization of expertise is accepted, it is not difficult to figure out what the offending implication of EJ anticipated above may be. For it is intuitive to think that the *same situation* in which the expert possibly knows (forms a justified belief) and the non-expert doesn't is to be specified in informational terms. In other words, it is intuitive to think that the expert is the subject who potentially knows (forms a justified belief) when the information which is available to her is the same as the information available to the non-expert, who doesn't know (doesn't forma justified belief); and if the informational situation in which the expert knows and the non-expert does not know is *the same*, how could there be a difference in the evidence they possess which explains the different

epistemic status of what they eventually come to believe? Indeed, it is not difficult to imagine situations in which an expert, by solely considering some piece of perceptual evidence, may come to know something falling within her area of expertise which a non-expert would not come to know even if, by mere chance, she were to form a true belief on the basis of the very same evidence. As an example, just imagine an expert bird-watcher briefly looking at a bird, and concluding on the *sole* basis of what she sees that it is an X, and compare this with a novice who transitions to the same conclusion on the basis of the same evidence; if the bird is really an X, it seems safe to maintain that the expert comes to know this fact; yet the expert has the same evidence as the non-expert, to which it does not seem intuitive to attribute knowledge of the fact that the bird is an X: whence the tentative conclusion that the difference among the expert and the non-expert, at least in the case under discussion, could not be accounted for in terms of the different amount of evidence they possess.

In the face of similar scenarios, contemporary evidentialists have reacted by contending that the difference among the expert and the non-expert still admits an explanation proceeding in evidential terms: despite impressions to the contrary, the expert has always more evidence than the non-expert, and in particular some knowledge about her area of expertise which, although temporarily set in the background, nonetheless makes the difference with respect to the non-expert (who does not possess that knowledge). The main difficulty with the reply at issue, from an Evidentialist point of view, is however that evidence temporarily set in the back-ground, as contrasted with the evidence which a subject is currently aware of, may not necessarily qualify as evidence which is "possessed"; and so, that it could not be invoked to explain the different justificatory status of the expert and the non-expert's beliefs, which is claimed to supervene on the evidence they *have*[1].

[1] R. Feldman contends that, in so far as the relevant background knowledge actively sustains the expert's belief, it should be counted as evidence possessed, despite its unconscious nature (2004: 239). So he tentatively draws the conclusion that a subject might be currently thinking of some evidence, and hence possess it, even if thinking of it unconsciously. I doubt that this solution may offer to the Evidentialist any effective way out: in the very first place, in fact, it is not very clear what it does mean for an unconscious belief *actively to sustain* a second belief; to say that it does so, in the absence of any further specification, may just seem to be an alternative way of describing the intuitive fact the expert's background knowledge somehow makes a difference for the justificatory status of her classificatory belief. How this should be explained, however, is precisely what is at stake among the Evidentialist, who contends that it contributes *qua* evidence possessed, and those who oppose Evidentialism on the ground that, given its unconscious nature, its contribution could not be understood

This difficulty, in many cases, can be overcome by adverting to a dispositional sense in which a subject can be credited with the possession of evidence, and so with justification: if a subject has already gathered some evidence E, and in spite of failing to currently consider this evidence retains the disposition to retrieve it when a question to which it is highly pertinent arises, then no principled difficulty seems to stand in the way of crediting a subject with the dispositional possession of E, and then with (dispositional) justification which E may happen to bestow on her beliefs. For instance, the dispositional sense in which some evidence may be possessed can be used to explain, consistently with Evidentialism, the knowledge-status of many stored beliefs for which a subject is currently considering no evidence: these beliefs will qualify as knowledge to the extent to which it will be reasonable to credit the subject who (dispositionally) entertains them with the dispositional possession of the relevant evidence. In the same way, it could be suggested, an expert transitioning to the true belief that *p* on the sole basis of a given piece of perceptual evidence will be justified to a degree sufficient for knowledge in believing that *p* provided that it will be reasonable to credit her with the dispositional possession of the relevant additional evidence; and a non-expert, transitioning from the same perceptual evidence to the belief that *p* will not come in that way to know this proposition because she will not possess the relevant evidence, either currently, or dispositionally.

One problem with this line of resistance is however the following: the idea that dispositional evidence may make the difference with respect to the justificatory status of a belief which is unsupported in the most central (occurrent) sense owes its credibility to counterfactual considerations; the idea is that, although that evidence is not used to justify a given belief, it would be made available and used if a question with respect to which it is pertinent did arise. More precisely, the principle which governs the notion of dispositional justification seems to be something like the following:

---

in that way. So, Feldman might be charged of begging her opponent's question when he substantiates the view that unconscious evidence may be evidence possessed in light of the observation that this evidence may at times actively sustain other beliefs. Moreover, it might be contended that the most central feature of evidence possessed which makes it an ideal candidate for making the difference for the justificatory status of a belief is that it is evidence of which a subject is currently aware; the problem, once this is accepted, is that although there is a reasonable sense in which it could be said that something is being thought unconsciously, no similar sense seems to sustain the suggestion that something might be the object of an unconscious awareness.

(D-j***) If the O-evidence E would justify the belief that *p*, the D-evidence E D-justifies the belief that *p*. (Piazza 2009)

For instance, the reason why I can be credited dispositional knowledge of the name of my mother-in-law even if, absorbed by different things, I am currently entertaining no evidence bearing on that issue is that, if asked, I could readily retrieve the pertinent information and return, on its basis, the fully (and then occurrently) justified answer that her name is Carla. In the case of an expert which forms a belief on the sole basis of perceptual evidence, however, it could not do to essay an explanation of the belief's knowledge-status which relied on the observation that she has background knowledge, and so that she possesses (additional) dispositional evidence: for the very good reason that the expert does know *occurrently* what she comes to believe on the basis of the perceptual evidence she is aware of, and the best we can expect from dispositional evidence is dispositional justification, and, only under certain conditions, dispositional knowledge. So, although it may be true that the expert will have D-evidence which the non-expert will fail to have, and so – contrary to the initial assumption – it will be true that the expert possesses a larger amount of evidence, it will be nonetheless the case that the evidence which could explain the different justificatory status of their respective beliefs will be the *same*; and this could well seem enough to counter, on a principled basis, the general principle advertised above.

A better strategy, it could be suggested, would then seem to need quite different materials. Just consider the following suggestion: so far, the expert and the non-expert have been shown to have the same evidence (at least that which is relevant to the occurrent justificatory status of the beliefs they entertain) and yet to differ with respect to what they know. This diagnosis crucially rests on the (suppressed) premise that the perception of the bird is evidence for the novice as it is evidence in the hands of the expert. As we have seen, the Evidentialist reply considered above did share this presupposition, as it tried to vindicate the claim that the expert and the non-expert have at their disposal different amounts of evidence by showing that the expert has at her disposal *additional* evidence which the non-expert fails to possess. However, this way of describing the situation is not mandatory. Much depends on what one takes to be *constitutive* of evidence.

In response to this question E. Conee has once returned the answer that any *indication* of a proposition's truth should be regarded as evidence for that proposition (Conee and Feldman 2004: 15, 252). As to the related question about what "sorts of things count as genuine indications of the truth of propositions" (BonJour 2007: 159) the idea – conveyed under the

heading of "seeming evidentialism" (SE) – is that anything which might seem to a person to bear on the truth of a proposition should be regarded as an *indication* in a sense relevant for its constituting genuine evidence (Conee and Feldman 2004: 15). SE implies that a piece of information held by a person could not be taken to be evidence for a given proposition *p* unless it seems to that person that that piece of information bears on the truth of *p*; so SE would seem to make available to the Evidentialist a different diagnosis of the scenario considered above.

Since she is a fresh trainee, the non-expert bird-watcher does not have the slightest idea as to how X-birds typically look like (one can take this to be somewhat constitutive of her lack of expertise). Therefore, even if she perceptually discriminates the X-exemplar in all its relevant details, she is not in a position to appraise the visual information she gathers as an indication of the fact that the bird in view belongs to species X. If it is so, however, SE will predict that this information is not evidence for the proposition that the bird is an X, and this might well seem enough to explain, in evidential terms, the different epistemological merits of the expert and the non-expert's beliefs: even if the relevant evidence is set in the background, and so is unsuited to explain the epistemic merits of what the expert is currently believing *qua* evidence, it seems safe to assume that its role will be reflected by the expert's *unmediated* disposition to believe, on the basis of what she is seeing, that the bird is an X; so, although both beliefs will have been formed on the basis of the same information, just the first will have been formed on the basis of evidence; so the second, predictably enough, will not constitute knowledge, while the first almost certainly will.

Despite it points in the right direction, also this reply seems to have its own problems though. Sure, the distinction it introduces between information possessed, and evidence seems to be correct: although there seems to be no denying that the non-expert bird-watcher does acquire visual information from the X-bird upon attentively looking at it, the claim that she thereby comes in possession of evidence for the proposition that it is an X seems to be much more dubious. However, SE seems unsuited to deliver a satisfactory explanation of why it is so.

To begin with, SE seems to have its own problems, and in particular the problem – outlined by L. BonJour – that it seems to allow for an undue proliferation of evidence: to name just one dubious consequence of the view under discussion, for instance, SE seems to imply that "for someone to whom it seems that a certain sort of physical appearance shows that a person is a witch, that sort of appearance really is evidence of witchhood" (2007: 160). To the extent to which we are unwilling to qualify as "evidence" subjective and probably misguided impressions, like the

delusional thought that a woman with a gothic look and dressed all black will most probably be a witch, SE then seems to provide no attractive answer to the question about what is constitutive of evidence.

The latter point should not be underestimated: under one *modest* reading, it could just seem to convey a terminological remark, i.e. a plead for reserving the label "evidence" just to those things, whatever they are, which are able to confer justification on someone's beliefs. Were it so, this point would be most welcomed by those who like philosophical theories which are conservative with respect to common usage. In fact we commonly say things like "Mr President was condemned on the evidence delivered by an eye witness", or "these footprints in the snow are evidence that someone passed by during the night", and so just a theory which identified evidence with *justification-providing* evidence could seem able to comply with the pre-philosophical understanding of the label. However, the problem is not just a matter of *style*, as it would be if the relevant disagreement occurred between those who care, and those who do not care for philosophy's continuity with respect to common usage. The most urgent problem is that SE, by issuing the prediction that the witch-belief would be based on evidence, threatens to imply that the belief is epistemically justified, a consequence which could fly in the face of the general agreement that witches do not exist. Sure, if by "evidence" one does not mean right from the start justification-providing evidence, the latter consequence is not immediate: a person might have evidence for a proposition, and yet fail to be justified because believing the proposition might be an *unfitting* attitude to take (coming back from work I see my wife's car parked in front of our house; if I were to take it as an indication of the fact that, despite she arrived at home first, she is not anymore there, because in the meanwhile she has been abducted by the aliens, SE would imply that I have evidence for that proposition, yet, since the belief does not fit this evidence, that I would not be justified); however, if one also believes – together with the Evidentialist – that justification is a matter of *fit* with the evidence, it seems to become unavoidable: for the very simple reason that if a person believes that every woman with a gothic look is a witch – and this belief is rightly counted as being part of the person's evidence[2] – she will form a belief which fits the evidence

[2] I think that the classification of the (arguably delusional) belief that every gothic-looking girl is a witch as evidence could not be contested, not at least from the point of view of a considered Evidentialist. One may be led to deny that this belief is actually evidence for the proposition that someone is a witch on the ground that, differently than a perception as of someone with a gothic look, it is a generalization which, on its own terms, is *neutral* as to whether someone is or is not a witch. However, remember that the most promising strategy which the Evidentialist has at her disposal to preserve EJ

whenever she will come to believe, upon seeing a woman with a similar look, that that woman is actually a witch. Whatever it is meant by saying that belief in a proposition *fits* the evidence, in fact, it seems that belief in a logical consequence of one's evidence should be regarded as an epistemic attitude which *fits* this evidence. The same problem can be illustrated in a variety of different ways: suppose, just to give a further example, that a person believes that tea leaves arranged in a certain way indicate that the weather will turn to rain; if it is so, whenever she will recognise that pattern of arrangement, she will apprehend the tea leaves as evidence for the proposition that the weather is going to change, and her belief that the weather will turn to rain will *fit* the evidence, and so it will be justified.

In addition, even if we were to accept it as an intrinsically acceptable, it seems that SE could not deliver the materials needed to counter on a principled basis the objection levelled at the beginning of this paper against EJ. The problem is that, even if SE may sustain the thought that the non-expert bird watcher does not have evidence for the belief that the bird is an X, and so explain, consistently with EJ, why just the expert knows, it seems unsuited to counter, in the same way, a suitably modified version of the counterexample at issue. Just suppose, in addition, that the novice bird-watcher, upon receiving the visual information from the bird, also happens (for no reason whatsoever) to have the inclination to take it as an indication of the fact that the bird is an X. In a similar possible (unlikely, yet conceivable) situation, SE still commits to regarding the visual presentation of the bird, as experienced by the non-expert, as evidence she has for (and upon which she bases the belief in) the proposition that the bird is an X; yet, also in this case we seem compelled to deny to her belief the status of knowledge. So, again, we are back to the initial problem: since the expert, upon basing her belief on the visual presentation of the bird, comes to know that it is an X, and the non-expert, upon basing the same belief on the same evidence, does not, we seem to have a counterexample to the claim that any such difference will always have to be explained in terms of the different amount of evidence they possess. So, I believe that the Evidentialist, if she is to have any chance of solving the problem, should really look in a different direction. One possibility is the following.

---

on the face of the bird-watcher example is to credit the expert with additional evidence, and in particular with some generalization as to how X-birds typically look like. So, an Evidentialist which is in the business to pursue this strategy could not *coherently* deny the evidence-status to the belief about how witches look like on the grounds that it is a generalization, because she is committed to acknowledging that status to the generalization as to how X-birds typically look like.

Also from the standpoint of an Evidentialist, it should sound as a truism that a person's justification does not solely depend on the evidence one has, but also on the way in which different pieces of evidence actually *interact* with one another. With this, I just want to exploit the rather harmless observation (which EJ, in the form I have considered it, somewhat conceals) that whether a person, on the basis of her global evidence (E_1 ... E_n), is justified in entertaining a given proposition does not solely depend on the nature of each and every E_k which belongs therein; it obviously also depends on the way in which these evidential items combine together. Whether a given E_k does supply a justification for believing a proposition, in other words, often depends on what further evidence E_{k1} ... E_{kn} one has, and this in either of the following two senses: believing the target proposition may fit *less* E_k alone, than it fits E_k plus other evidence one has; or believing the target proposition may fit *more* E_k alone, than it fits E_k *plus* other evidence one has: in the latter case the additional evidence will probably defeat the justification which supervenes on E_k alone. Let us however concentrate on the first case, the case in which the degree of fit (and, along with EJ, the degree to which a belief is justified) is greater when E_k is combined with other evidence one possesses than it would have been otherwise. Let us dub the case at issue one in which the relevant pieces of evidence *interact constructively*.

It is arguable that constructive interaction is a vast category, and that it comes in at least two varieties which deserve being singled out separately. One mode of interaction is *addition*: this is arguably the way in which two distinct and *convergent* testimonies as to some *p* combine, thereby supplying a justification for the belief that *p* which is stronger than the justification supplied by just one of them; or it is the way in which *more* instances of a given generalization combine, thereby supplying an (inductive) justification for it which is stronger than the justification supplied by *less* instances of the same generalization. Arguably, one fundamental feature of interaction by addition – above and beyond the fact that it results in a degree of justification which is greater than it would have been otherwise – is that each piece of information could not so interact unless, in itself, it were not capable of supplying some positive degree of justification: for a subject who knows nothing about a given generalization – so for a subject committed to stay neutral as to whether it is true – the observation of just one instance of it should be taken to supply some positive degree of justification, one which, although in a small measure, will raise the generalization's probability. In the same way, for a subject who knows nothing about whether *p* has been the case, so for a subject committed to stay neutral as to whether the proposition

is true, one testimony in favour of *p* supplies some positive degree of justification, one which will raise *p*'s probability. What happens as new evidence is acquired (either new testimonies are delivered, or more observational data which confirm the generalization are gathered) is just that the strength of the previous justification increases. It is important to focus on this feature of interaction by addition, because, arguably, it is what differentiates it from the second variety of constructive interaction I want to present.

Just consider the following (rather artificial) case: an astronomer has at her disposal just two telescopes, one of which is an ultrareliable instrument, while the other is not reliable at all. Say that the first one, call it A, is such that whenever you see (what seems to be) a star while observing the night sky through it, the probability of the star's being there is very, very high; B, the other telescope, has the opposite feature: whenever you see (what seems to be) a star while observing the night sky through it, the probability of the star's not being there (your impression to the contrary being aroused by any of many possible malfunctioning) is very, very high. Now suppose that the astronomer looks through one of the telescopes, and sees something which looks very much like a star (call this evidence E). It is sensible to believe that if the astronomer is using A, E justifies the astronomer in believing that there is a star over there to a rather high degree; on the other hand, if the astronomer is using B, the opposite conclusion would seem to be better justified: in that case it would be very unlikely that there really is a star over there. The problem, however, is that A and B are indistinguishable: the astronomer, by simply looking at it, cannot tell whether it is A or B which she is using. However suppose that a clever student of the astronomer figured out a method to tell A apart from B, and that after having put this method into practice, she tells the astronomer that he has been using A (call this additional evidence E*). Obviously, E and E* interact constructively: for the astronomer, the belief that there is a star over there is now justified, on the basis of E *plus* E*, better than it was justified on the basis of E alone. This time, however, this is true much for the reason that the belief that there is a star over there is justified for the *first time*, as, E alone (namely, unaided with the information as to which telescope has been used) supplies to the astronomer no degree of justification for either believing, or disbelieving that there is a star over there. For this very reason I propose to label this way in which two pieces of information my combine interaction by *enabling function*: the contribution of E* is not just to add to the strength of a justification that has already been supplied by E; as it were, its function is to disclose the epistemic potential of E which, if E* had not been produced, would have remained un-actualized.

Now, the intuition underlying the Evidentialist's response to the bird-watcher's example, reviewed at the beginning of this paper, would seem to be that the difference between the novice and the expert is made by a piece of information which *enables* (in the expert's case), and would have enabled, yet since it is not possessed it does not enable (in the novice's case), the visual presentation of the bird to supply a justification for the belief that it is an X. For it seems intuitive that the novice, who lacks the additional information (the ornithological back-ground knowledge), is justified to *no* degree in believing that the bird, as it happens to be the case, is an X; that which sustains the diagnosis that the information which is missing (in the case of the novice) and is arguably possessed (in the case of the expert) could not have interacted with the information which both have (the visual presentation they possess) by *addition*. The problem which has emerged, as it will be remembered, was however that the additional *enabling* evidence, since it is not entertained consciously, could not be taken to be possessed in some central sense; and that the obvious alternative, that it was possessed, yet in the (derivative) dispositional sense, in turn has proven useless; for the best we can expect from dispositional evidence is dispositional justification, whereas it has appeared clear that the expert, who transitions to the belief that the bird is an X on the basis of her perception of it, is occurrently, and not just dispositionally justified.

However, once the model of interaction by enabling function has been brought into plain view, there seems to be a clear way in which the basic suggestion proposed by the Evidentialist could be preserved consistently with the observation that the relevant background knowledge is not (occurrently) possessed; in fact, until now, we have placed no significant constrain as to what could *enable* one piece of information to perform an evidential function; so, although the model has been illustrated in light of an example in which this role is discharged by an additional piece of evidence (the information that the telescope used by the astronomer was A) there is no apparent reason why we should believe that something else could not discharge that role.

Consider the following example, which is supposed to pave the way for a broadened model of the interaction by enabling function. Mafalda, a four-years girl, is brought to the zoo by his father; since he has already been there many times, Mafalda has now learned to tell a leopard apart from a cheetah. So, when they both look into the cage where five cheetahs are sleeping, Mafalda comes to know that it is cheetahs, and not – say – leopards, that she is seeing. However Mafalda has not yet learned to count, so although what he sees allows her to know what animals there are in the cage, the very same information doesn't allow her to know

that the cage contains *five* animals. What we have here, it would seem, is a case in which the very same piece of information *I* which Mafalda acquires by looking into a cage is evidence for a proposition, yet is not evidence for a different proposition; and this makes almost irresistible the thought that the difference must be explained by supposing that an enabling condition is operative, in so far as the evidential status of *I* relative to the first proposition is concerned, and that a different condition, whose satisfaction would have enabled *I* to count as evidence for the second proposition, is not operative[3]. By the same token, it appears clear that both the first and the second enabling conditions do not require Mafalda to possess additional evidence; what they require is that Mafalda possesses a *cognitive ability*: it is because Mafalda is able to tell a leopard apart from a cheetah that she is allowed to learn, upon looking into the cage, that it contains cheetahs; and it is because she does not have the ability to count that the very same visual information has no cognitive significance relative to how many cheetah it does contain. So, we have a case in which it is (the exercise of) a cognitive ability which interacts with a piece of information by enabling its evidential function, and not, as in the astronomer's case, some additional piece of evidence. So, we seem to have all the materials for explaining what is going on in the expert's case consistently with the observation that he does not possess (in some central sense) additional enabling evidence; for even if the function of the expert's background knowledge is not to contribute additional *evidence*, it is sensible to believe that what constitutes her very expertise is the possession of an ability – which at times may be manifested in the form of a piece of propositional knowledge, yet does not necessarily coincide

[3] I think that a person like Mafalda who is not able to count need not be disqualified as one who also fails to possess the relevant numerical concepts. Were it so, Mafalda's inability to discern, on the basis of what she is seeing, that the cage in view contains five animals could be explained, more straightforwardly, by observing that she is not able to entertain thoughts having the concept of five as an ingredient. Also in this case, I believe, the advertised dependence of the evidential status of the visual information acquired by Mafalda on the possession of the relevant ability would be vindicated; the general suggestion, in this case, would become that the evidential status of information should be taken to depend on the possession of the concepts which are needed to grasp the propositions this information is possibly evidence for (as we will see, this suggestion is almost implicit whenever the relevant concept, differently than a numerical concept, is a recognitional one). However, it is also fair to admit that the possession of the concept of five does not seem to be constrained by the possession of the ability to identify, perceptually, groups of five members. So, Mafalda might well be able to grasp a proposition having as ingredient the concept of five, and yet, as in the present case, fail to possess the ability to count up to five.

with it[4] – which *enables* the visual information she has acquired to perform an evidential function; and what explains the fact that the novice, who holds the same information, does not come to know on its basis what the expert knows, is simply that the novice, differently than the expert, does not possess the relevant ability, and so for this reason just possesses information, and not evidence.

Of course, as it is well-known, it is not very clear, from a philosophical point of view, what an ability (in particular, as contrasted with an opportunity) should be supposed to be. However, one thing seems clear: the relevant ability is an ability to process information in a certain way; so, it is sensible to expect that considerations of an information-theoretic nature may help clarify, at least to a certain extent, what, in the particular case at issue, an ability should be taken to be. Here is something which F. Dretske (1999) may offer to us which is highly relevant to the case at issue.

Quite interestingly, Dretske distinguishes among *sensory experience* and *cognition*, the conceptual use of the information made available by the former. According to the proposal at issue, the main difference between the both is that they *code* information in different ways. It is customary to notice that the information relative to some property of a given source may be coded either in a *analogical* or in a *digital* way, depending on whether the value of this property is represented continuously or discretely[5]. Dretske's distinction, although it is conveyed by resorting to the very same vocabulary, concerns the way in which a system may encode the distinct piece of information *that* a given source instantiates a given property with a given magnitude: in other words, not information about *properties*, but about the *facts* that may have those properties as constituents. So what Dretske is after is a distinction, essentially paralleling the first one, among the way a given system may

[4] I take the relevant ability to be *practical*; in particular, I believe that it should be identified with a recognitional ability of some sort. This is the reason why I believe that the ability which makes the difference between the expert and the novice may at times be manifested in the form of explicit knowledge as to how X-birds look like, yet does not need coincide with it. What is constitutive of a practical ability is that whoever possesses it *knows how* to do a certain thing – in the case at issue *how* to recognize an X-bird –, not that she is also able, in each and every occasion in which she exercises it successfully, to state in explicit terms what criteria govern its exercise.

[5] The speedometer of a car encodes in an analogical way information about its speed because, at least to some extent, to every state of the speedometer there corresponds a value of the car's speed; the light of a dashboard which registers oil-pressure, on the other hand, encodes information about the pressure of the oil in a digital way, because the encoding system has just two states at its disposal, one that encodes the information that the pressure has a smaller value than some suitably chosen *v*, and another that encodes the information that the pressure has a value which is equal to, or greater than *v*.

represent the information, for instance, that a car has a speed of 40 km per hour (as contrasted with a system representing a speed of 40 km per hour).

According to Dretske, in order for the information that *S is P* to be encoded in a digital way by a given structure the most specific information encoded by this structure must be the information that *S is P*; by contrast, the very same information will be encoded in an analogical way whenever there will be some more specific information about *S* carried by the structure other than the information that *S is P*. One case in point is the different way in which a picture, and a statement, typically code information, say the information that a given car is red: a picture normally (even if not necessarily) encodes this information analogically because it typically portrays the car within a context, as having certain dimensions, a given shape, etc. Moreover, and most inescapably, it represents the car as being of a *particular* red, *illuminated* in a certain way, etc. Contrast now a picture of the red car with the sentence "the car is red": the sentence, too, conveys the information that the car is red; yet it *just* conveys the information that the car is red, so, *a fortiori*, this is information is the most specific information about the car it encodes; this is why the sentence, differently than the picture, codes this information in a digital way.

Dretske actually puts the distinction already introduced at the service of explaining, in information-theoretic terms, what happens when, on the basis of a perception, someone acquires full cognitive awareness of something which is presented through it. To begin with, Dretske maintains – plausibly enough – that a perception normally encodes far more information than a belief, based on it, may finally encode. So a perception will typically code this information in an analogical way, while the belief will encode this information in a digital way. The transition can then be portrayed as one whereby the information which is embedded in analogical form within our experience is *extracted* from it and *converted* in digital form. This analogical-digital conversion qualifies as generalization and a fully epistemic classification because it involves a loss of information that is consequent on a process whereby individual differences (the features of the particular red of the car, for instance) are disregarded, and treated as being of essentially the same kind. As Dretske puts it:

> "Digital conversion is a process in which irrelevant pieces of information are pruned away and discarded. Until information has been lost, or discarded, an information processing system has failed to treat *different* things as essentially the *same*. It has failed to classify or categorize,

> failed to generalize, failed to 'recognize' the input as being an instance (token) of a more general type. […] Cognitive activity is the *conceptual* mobilization of incoming information, and this conceptual treatment is fundamentally a matter of ignoring differences (as irrelevant to the underlying sameness), of going from the concrete to the abstract, of passing from the particular to the general. It is, in short, a matter of making the analog-digital transformation." (*ibidem*, 141, 142)

What is important, in the present context, is that Dretske's account of conceptualization and classification allows us to give a more precise content to the intuitive characterization of a cognitive ability introduced above. The simple suggestion is that the relevant cognitive ability must be one which encapsulates the digital-conversion unit which is selectively responsive to the information, coded in an analogical way by the relevant experiential presentation, which (belief in) the proposition eventually comes to represent in a digital way. A concept will typically embody such conversion unit: according to C. Peacocke (1992), for instance, the conditions which someone must satisfy if she is to possess the concept of red is that she must be primitively disposed to transition to belief that something is red if presented with something which appears to her as being red. So the concept of red, according to Peacocke, is a mental structure whose proper function is to extract pertinent information to the effect that something is red from its analogical sensory presentation, and to encode it in a digital way. This, as hinted at before, would be a case in which the evidential import of the information encoded by perception is constrained by the possession of the relevant ability in the more straightforward sense that a person unequipped by such an ability is unable to grasp the propositions this information is evidence for. At the same time, the ability to count can now be illuminated as the digital-conversion unit which constrains the epistemic significance of Mafalda's visual presentation of the animals in the cage: unless Mafalda's psychological architecture sustains the relevant conversion unit, the information that there are five cheetahs in the cage, which is coded in an analogical way by his perception, although one which she will possess (she sees five cheetahs), will be cognitively unavailable to her (she does not see *that* the cheetahs are five).

The general moral that the preceding considerations make available is that the evidential import of a given body of information I, relative to a given proposition p, is conditional on the possession of a mental structure S needed to select from I the relevant part which, eventually coded in digital form, becomes cognitive possession in the form of a belief that p.

It is important, in the present context, to stress that the presence of the relevant conversion unit does not constrain in any significant way the nature of the information *I* which is part of the "sensory information store" potentially serving as evidence for the purpose of justifying one's beliefs. A child which we would ordinarily characterize as lacking the capacity to discriminate a daffodil need not be regarded as "*perceptually* deficient" when he looks at a yellow flower, and yet is unable to return the answer that it is a daffodil.

> "What the pupil needs is not more information of the sort that could be supplied by the use of a magnifying glass. [...] The requisite (requisite to identifying the flower *as* a daffodil) information is getting in. What is lacking is an ability to extract this information, an ability to decode or interpret the sensory message" (*ibidem*, 144).

In other words, for that pupil it is not *information* to the effect that the flower is a daffodil that is lacking; rather, it is *evidence* for that proposition that, courtesy of the missing ability, is not there to sustain an intelligent transition to the corresponding belief. The expert's (the novice's) situation can be described in essentially the same terms. The novice does enjoy all the information which would be sufficient, for someone endowed with the relevant ability, also to enjoy the evidence in light of which the belief that the bird is an X. However, he does not possess the relevant pieces of ornithological knowledge that sustain the expert's cognitive routine, and so fails to have evidence in the very first place. Therefore, predictably enough, he does not know either.

***

In this paper, my principal concern has been the notion of expert knowledge, and the account of it which the current view on epistemic justification known as Evidentialism is bound to propose. In particular, I have concentrated on the standard implication of Evidentialism, according to which the same belief could not be justified differently for two different persons (or for the same person at different times) unless these persons (or the same person at different times) had not different evidence. As far as the way in which expert knowledge is understood in this paper, the latter implication is *prima facie* problematic: for an expert is exactly the kind of person from which knowledge (and so *better* justification) is expected in situations in which a non-expert would not know (would have a *worse* justification). If a situation is understood as defined by the information which is available in it (so as an informational

situation), the latter understanding of expert knowledge makes almost irresistible the conclusion that an expert who knows a given proposition in an informational situation in which a non-expert does not could not differ from the second for the evidence she has. The final part of the paper has been devoted at resisting the latter line of thought: in particular, drawing on a suggestion presented by Conee, I have suggested that the information held by a subject should be distinguished from the evidence this person has. Conee, however, contends that the information available to a subject becomes evidence for a proposition when the person has an inclination to regard it as an indication of the truth of a proposition. My proposal differs from Conee's. First I differentiate between two ways in which distinct pieces of evidence may interact *constructively*, i.e. by supplying – in conjunction – a justification for a proposition that is stronger than the justification which each of them – in isolation – would have made available. The first more intuitive way is by addition: in this case, the strength of the (conjunctive) justification is a function of the strength of the justification independently supplied by each conjunct. A second way of constructive interaction is less apparent, and it is the way in which two pieces of information interact when one becomes *evidence* for the first time only when (and only because) a second information is added. I have found it congenial to dub this kind of interaction by *enabling function*. As I have tried to show, this model of interaction generalizes interestingly when we see that an enabling condition, which sometimes is met because a second piece of (enabling) information is added, may sometimes be fulfilled because a subject possesses a cognitive ability, which is the ability to process a first piece of information in a full cognitive way. In some central cases, the relevant ability will be the one which underlies the possession of the concepts involved by some target proposition; and in these cases it will be straightforward that a first piece of information will be evidence for the proposition only if processed by resorting to the that cognitive ability: in a similar situation, a person unequipped with this ability will be most directly unable to grasp the proposition at issue, and so, predictably enough, the information at issue will not count as evidence for that proposition which she will have at her disposal. In other cases this will not be so, and a person lacking the relevant ability will be able to grasp a proposition, and simply fail to possibly entertain evidence for it. Both cases are explained, information-theoretically, by supposing that the ability at issue sustains the conversion-unit which is needed to process the first information in a fully cognitive way.

Once the distinction at issue is brought in plain view, Evidentialism is not anymore incapable to account for expert knowledge: once we see that information and evidence are not one and the same thing, the observation

that an expert knows precisely in those informational situations in which a non-expert would not know no longer establishes the claim that the different epistemic merit of their beliefs cannot be explained in evidential terms. As far as an expert has cognitive abilities which a non-expert lacks, she will have more evidence than a non-expert, even if both will possess the same information.

References

BONJOUR. L. "Review of *Evidentialism* by Earl Conee and Richard Feldman". *Mind*, 116 (2007), p. 157-161.

DRETSKE, F. *Knowledge and the Flow of information*. CSLI Publications, 1999.

CONEE, E.B.; FELDMAN, R. *Evidentialism. Essays in Epistemology*. Oxford University Press, 2004.

PEACOCKE, C. *A Study of Concepts*. MIT Press, 1992.

PIAZZA, T. "Evidentialism, and the Problem of Stored Beliefs". *Philosophical Studies*, 145, 2 (2009), p. 311-324.