Pragmatics, modularity and epistemic vigilance

Diana Mazzarella

Institut des Sciences Cognitive 'Marc Jeannerod'

University College London

Abstract

This paper addresses the question of the place of pragmatic abilities in the overall architecture of the mind. Until recently, pragmatics was assumed to be part of a non-modular, unencapsulated, central system. Sperber and Wilson (2002) have proposed that pragmatics is to be conceived of as a sub-module of the mind-reading module, with its own principles and mechanisms. This is in line with an increasingly modular view of the mind (Cosmides & Tooby 1992, 1994; Sperber 1994b, 2001, 2005; i.a.), according to which cognition consists of many dedicated domain-specific mechanisms or 'conceptual modules', highly interconnected with each other. This paper focuses on the connection between the pragmatics module and epistemic vigilance mechanisms, that is, mechanisms that assess the quality of the incoming information and the reliability of the individual who dispenses it (Sperber *et al.* 2010). The latter take as their proprietary input the output of the pragmatics module and assess its believability. This paper makes two original proposals: first, that epistemic vigilance mechanisms may directly affect the comprehension process, and, second, that the emergence of epistemic vigilance mechanisms targeted at assessing the communicator's competence and benevolence may correlate with different developmental stages in pragmatics.

Keywords: modularity, pragmatics, epistemic vigilance

1. Introduction

The cognitive revolution, which from the early '60s shaped the domains of linguistics, anthropology, psychology and related disciplines, manifested its effect in the field of pragmatics with the seminal work of Sperber and Wilson (1986/1995). Among many other issues, Sperber and Wilson brought to the attention of the pragmatics community the question of the place of pragmatic abilities in the overall architecture of the mind. At that time, Fodor had already suggested that human cognitive architecture is partly modular (Fodor 1983) by introducing the functional and architectural distinction between modular perceptual and linguistic processors, on the one hand, and non-modular higher-

level central processes, on the other. This gave rise to the interesting question of whether pragmatics is to be thought of as a domain-specific modular system or as part of a domain-general central system.

Fodor's (1983) distinction between modular input systems (perception and language) and non-modular central thought processes is based on a precise characterisation of the nature of 'mental modules'. Modular mental systems are task-specific and relatively autonomous systems which contingently share the following properties: they operate in a mandatory and fast way, they are domain-specific, informationally encapsulated and generally associated with a fixed neural architecture, and they exhibit specific breakdown and developmental patterns. Fodor places particular emphasis on informational encapsulation: "The informational encapsulation of the input systems is […] the essence of their modularity" (Fodor 1983: 70). A mental system is informationally encapsulated (or cognitively impenetrable) if it is rigidly restricted in its access to the full range of the organism's knowledge, beliefs and desires. Such a system cannot take account of (potentially relevant) information that does not belong to its proprietary database, that is, which lies outside its own task specific body of information for processing its particular domain of stimuli.

The Fodorian notion of 'module' seems hardly applicable to pragmatic processing given its undisputed property of sensitivity to a wide range of contextual or background information. A constant tenet of post-Gricean approaches to pragmatics (e.g. Relevance Theory and other contextualist accounts) is the assumption that linguistic meaning underdetermines not only what is meant, but also what is said or explicitly communicated, that is, the truth-conditional content of the utterance (Carston 2002). Consider the following example:

(1) Neil has broken his leg.[1]

The hearer of (1) must decide who the referential expressions 'Neil' and 'his' refer to. Does 'Neil' refer to $Neil_1$ (the hearer's son) or $Neil_2$ (her colleague in the linguistics department)? Does 'his' refer to the referent of 'Neil' or to a different male individual who got a broken leg from a scuffle with $Neil_1/Neil_2$? Furthermore, the hearer must decide when the event took place (at some generic time in the past? a few days ago? this morning?). An utterance of (1) may also convey some implicatures (intended implications): for instance, that $Neil_2$ cannot participate in the staff meeting because he is still at the hospital, or that $Neil_1$ cannot run in the school marathon because he has not entirely recovered from the bad accident in which he broke a leg. The linguistically encoded meaning of the utterance, thus, plays a minor, albeit crucial, role in the recovery of the communicated content of the utterance, both at the explicit and at the implicit level. Crucially, there is no principled restriction on the kinds of information that pragmatic processing can call on: perceptual information, background information stored in long-term memory, information that is part of the linguistic context of the utterance. Pragmatic processes appear to be 'informationally unencapsulated' and an account of the structure and function of pragmatics in the mind need to respect this feature of utterance interpretation.

In line with this, pragmatics has traditionally been conceived of as a non-modular central inferential process, i.e. a non-deterministic process of rational

---

[1] Carston 2007: 25.

belief-fixing. Specifically, it has been described as the process of arriving at an interpretation of the utterance, that is, the process of fixing a belief about the speaker's communicative intention. This has indeed been the standard position of Relevance Theory up to the late '90s (Wilson and Sperber 1986, Sperber and Wilson 1986/1995).[2]

The main aim of this paper is to explore recent developments on the issue of the place of pragmatics in the cognitive architecture of the mind from a relevance theoretic perspective. The structure of this paper is as follows: in Section 2, I present the relevance-theoretic account of pragmatics as a module (Sperber and Wilson 2002, Wilson 2005) and discuss the non-Fodorian notion of module that it relies on. In Section 3, I investigate the relationship between pragmatics and mechanisms of 'epistemic vigilance', that is, those mechanisms that assess the quality of incoming information and the trustworthiness of the individual who dispenses it (Sperber *et al*. 2010) and suggest that the interaction between these two systems goes beyond what is currently acknowledged. Finally, I present some developmental implications of this proposal and sketch directions for future research.

2. The pragmatics module

Sperber and Wilson (2002) suggest that the interpretative process is carried out by a dedicated pragmatics or comprehension 'module', with its own principles and mechanisms. How does this fit with the traditional view of pragmatics as an informationally unencapsulated system? In order to answer this question, we first need to consider the conceptual transformation that has characterised the notion of 'mental module' itself. Although pragmatics is now conceived of as a module, it is not a Fodorian module. For this reason, it is worth taking a step back to look at the wider picture of the mind within which Sperber and Wilson's (2002) proposal is located. This is the view of the mind as 'massively modular', a position pioneered by the evolutionary psychologists Cosmides and Tooby (1992, 1994) and advocated by Sperber himself (e.g. Sperber 1994b, 2000, 2005).

Starting from evolutionary considerations, Sperber (1994b) suggests that the mind is modular through and through, that is, that cognition is based on dedicated domain-specific mechanisms, as opposed to domain-general central processes. This view, which takes the name of 'massive modularity', subverts Fodor's architectural taxonomy of psychological processes, and introduces 'conceptual modules' in addition to perceptual ones. While a discussion of the massive modularity thesis goes beyond the purpose of the present paper, it is worth focusing on some of its implications: first, the introduction of a revised and looser notion of mental module, and, second, the hypothesis of a complex network of perceptual and conceptual modules "interconnected in ways that would make an engineer cringe" (Sperber 1994b: 46).

The assumption that cognitive mechanisms, like every biological mechanism, are adaptively specialised for the solution of particular kinds of task is what grounds the conception of 'mental modules' as domain-specific and autonomous computational mechanisms. They are attuned to the regularities of their specific domain and employ dedicated procedures which are justified by

---

[2] But see Kasher 1991 for a modular view of certain components of pragmatic knowledge (i.e. basic speech acts and talk-in-interaction). For a critical discussion, see Carston 1997.

those regularities. Significantly, they do not (necessarily) manifest all the properties that Fodor attributes to mental modules: their operations may not be mandatory (in the sense that an appropriate input may not be sufficient to trigger its own processing (Sperber 2005)) and their informational encapsulation may be conceived of as a matter of degree:

> it may be that we have to rethink the concept of module and allow for a kind of continuum, from peripheral perceptual systems, which are rigidly encapsulated (not diverted from registering what is out there), through a hierarchy of conceptual modules, with the property of encapsulation diminishing progressively at each level as the interconnections among domain-specific processors increase (Carston 1997: 20).

This passage interestingly highlights the connection between this new notion of mental module and the hypothesis that modules are highly interconnected. The output of a perceptual or conceptual module can be fed to other conceptual modules, whose outputs can in turn be fed to further conceptual modules, and so on and so forth. The result consists of a chain of inferences that integrates the contribution of each individual module.

With this picture in mind, let us turn to Sperber and Wilson's (2002) proposal that utterance interpretation is carried out by a dedicated inferential mechanism or 'comprehension module'. Sperber and Wilson suggest that pragmatics is a sub-module of the general mind-reading module, which is responsible for providing explanations of individuals' behaviours in terms of attributed mental states (e.g. beliefs, intentions, etc.).[3] Recognising the intention behind the speaker's communicative behaviour is a particular case of mind-reading (as Grice pointed out long ago). But, while utterances are a type of action and a speaker's meaning is a type of intention (i.e. a communicative intention, which is a second-order informative intention), according to Relevance Theory, the domain of overt communication exhibits such specific regularities and is so important in human life that, instead of employing general mind-reading procedures, it deploys its own dedicated comprehension procedure.

The investigation of such regularities requires the introduction of some technical notions, among which the central one is 'relevance'. Relevance is a property of inputs to cognitive processes (e.g. percepts, utterances) and it is a cost-benefit notion: the smaller the processing effort (cost), the greater the relevance; the greater the cognitive effects (benefit), the greater the relevance. Sperber and Wilson (1986/1995) identify three kinds of cognitive effects: contextual implications (i.e., implications that can be derived from the input and the context, but from neither input nor context alone), strengthening of existing assumptions, and contradiction and elimination of existing assumptions. According to Sperber and Wilson (2002), human cognition has been subject to a continuous evolutionary transformation towards greater cognitive efficiency, so that it tends to be geared to the maximisation of the relevance of the information processed (this is the First, or Cognitive, Principle of Relevance). Given this

---

[3] For a detailed defence of this claim see Sperber and Wilson 2002 and Wilson 2005. According to this view, mind-reading is not a single, homogeneous system but a collection of autonomous dedicated mechanisms, or sub-modules (e.g. the Eye Direction Detector module).

universal tendency to maximise relevance, an audience will pay attention to a stimulus only if it seems relevant enough. By producing an *ostensive* stimulus (e.g. an utterance), a communicator raises a particular expectation in the audience that is not raised by other stimuli. Specifically, her ostensive stimulus conveys a presumption of its own optimal relevance (this is the Second, or Communicative, Principle of Relevance), which is defined as follows:

(2) **Presumption of optimal relevance**
    The ostensive stimulus (e.g. the utterance) is presumed to be (i) at least relevant enough to be worth the hearer's attention and (ii) the most relevant one compatible with the speaker's abilities and preferences.

The Communicative Principle of Relevance expresses the kind of regularity that characterises the domain of overt communication and it is this which, according to Sperber and Wilson (2002), motivates the adoption of the following dedicated comprehension procedure (that works according to an in-built presumption of optimal relevance):

(3) **Relevance-guided comprehension procedure**
    (a) Follow a path of least effort in computing cognitive effects: Test interpretative hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
    (b) Stop when your expectations of relevance are satisfied.

This procedure guides the construction of appropriate hypotheses about explicit content, intended contextual assumptions and intended contextual implications (implicatures) within the overall comprehension process. Note that the regularity expressed by the Communicative Principle of Relevance is specific to the domain of overt communication: in general, an observer is not entitled to expect that the intentional behaviours of others would have any particular level of relevance to him. But since this expectation is warranted in the case of overt communicative behaviours, the relevance-guided comprehension procedure tends to yield reliable conclusions. Here is an illustration of how the relevance-guided comprehension procedure applies to example (1). Consider the following scenario (described by Carston 2007: 25). Robyn is in one of her students' company. At some point during the conversation, the student, Sarah, addresses to Robyn the following utterance:

(1) Neil has broken his leg.

Suppose that Robyn knows two people called "Neil": $Neil_1$, her young son, and $Neil_2$, a colleague in the linguistics department. While Sarah does not know $Neil_1$, she is acquainted with $Neil_2$, who teaches her syntax. Assume that Robyn is a very apprehensive mother and that she is always worried about her son $Neil_1$. When hearing Sarah's utterance, then, the most accessible referent for "Neil" is $Neil_1$. Furthermore, because $Neil_1$ is both a very clumsy and fearless child, the first interpretative hypothesis to come to her mind is that $Neil_1$ has broken his own leg (after she left him at the kindergarten that morning). In following a path of least effort, this is the first hypothesis to be tested. The comprehension procedure stops when it reaches an interpretation that satisfies the hear-

er's expectations of relevance. The interpretative hypothesis that $Neil_1$ has broken his leg this morning, however, does not satisfy Robyn's expectation of relevance. The reason is that she knows it is not compatible with the speaker's abilities (Sarah does not know that Robyn has a son called "Neil"). The interpretative hypothesis is thus discarded in favour of a less accessible one, e.g. $Neil_2$ has recently broken his own leg.

Interestingly, while the construction of interpretative hypotheses calls on various sources of contextual information, the relevance-guided comprehension procedure clearly suggests that not all available contextual information has to be actively taken into consideration. Rather, pragmatic interpretation exploits "whatever information is most highly activated by the automatic workings of the cognitive system at the time" (Wilson 2005: 1141).

3. The pragmatics module and epistemic vigilance

3.1 The output of the pragmatics module

The massive modularity framework of the mind recognises that modules are highly interconnected with each other, that is, they form a network of systems and subsystems connected in such a way that they may take as input the outputs of several other modules. In what follows, I focus on the pragmatics module and its connection with epistemic vigilance module(s) (Sperber *et al.* 2010).

The pragmatics module takes as input an ostensive stimulus and delivers as output an interpretative hypothesis about the communicator's meaning. Importantly, pragmatic interpretation corresponds to the process of fixing a belief about an interpretative hypothesis, i.e. about which propositions the speaker communicated ('comprehension'), rather than to the process of fixing a belief in the propositions themselves ('(doxastic) acceptance'). Sperber *et al.* (2010) have recently investigated the distinction between comprehension and acceptance and suggested that the latter does not automatically follow from the former. Rather, they claim, humans have developed a suite of cognitive mechanisms (the 'epistemic vigilance module(s)'), which assess the believability of a piece of communicated information and act as a filter at the entrance of the 'belief box' of the interpreter. Since communication is open to the risk of misinformation (be it accidental or intentional), the only way for it to remain advantageous (on average) for both communicator and audience is for its outcomes to be assessed by mechanisms that monitor the quality of the incoming information and the reliability of the individuals who dispense it before accepting it.[4]

The metarepresentational output of the pragmatics module, i.e. 'The speaker$_x$ meant that $p_1, .., p_n$' (where $p_1, .., p_n$ is the set of propositions communicated), provides the input to two different kinds of epistemic vigilance mechanisms: mechanisms that focus on the source of information (*who* is to be believed) and mechanisms that focus on the informational content itself (*what* is to be believed). While the former assess the reliability of the speaker, that is, whether the speaker is competent (epistemically reliable) and benevolent (morally reliable),

---

[4] It may seem that epistemic vigilance mechanisms have developed entirely for the sake of the interpreter but they can be advantageous for the communicator too: "from the communicator's point of view, a vigilant addressee is better than one who rejects her testimony outright" (Sperber *et al.* 2010: 376).

the latter assess the degree of believability of the incoming information (the propositions $p_1 ..., p_n$), independently from its source.

Epistemic vigilance mechanisms towards the source of information can deliver either general impressions of trustworthiness (e.g. on the basis of facial clues, see Willis and Todorov 2006) or more costly assessments that result from context-sensitive evaluations of the reliability of the speaker. For instance, given an input of the form 'The speaker$_x$ meant that $p_1, .., p_n$', they may assess whether the speaker may want the audience to believe the set of propositions $p_1, .., p_n$ for reasons that do not concern their (alleged) truth (e.g. because of some deceptive intention). Or they may detect that the set of propositions $p_1, .., p_n$ is warranted by some beliefs of the speaker that are, in fact, false. In both cases, they would prevent $p_1, .., p_n$ from entering the belief box of the interpreter.

Epistemic vigilance mechanisms towards the content, on the other hand, assess the believability of the incoming information relative to the context of the addressee's existing beliefs (which are themselves, of course, open to revision). Specifically, according to Sperber *et al.* (2010), the beliefs against which the communicated information is tested are those that are automatically activated by the comprehension process and used in the pursuit of relevance. These are a subset of the mental encyclopaedia of the addressee, and provide the ground for an "imperfect but cost-effective epistemic assessment" (Sperber *et al.* 2010: 374). When the result of this assessment is a contradiction, there are three possible outcomes: (i) if the source is taken as trustworthy and the background beliefs of the interpreter that conflict with the incoming information are not held with much conviction, these beliefs are corrected; (ii) if the source is not regarded as trustworthy, the new information is rejected; (iii) if the source is regarded as authoritative and the conflicting background beliefs are held confidently, some process of (typically conscious) coherence checking is triggered. Interestingly, the choice among (i), (ii) and (iii) partly depends upon the output of epistemic vigilance mechanisms focused on the source (the speaker).

To sum up, the pragmatics module is conceived of as interconnected with the epistemic vigilance module(s), whose mechanisms assess the reliability of the source of information and the believability of its content. These mechanisms are geared towards preventing the interpreter from being misinformed by filtering the communicated information that he ends up believing. In the next section, I explore the possibility of extending the scope of interaction between the pragmatics module and epistemic vigilance mechanisms (see also Mazzarella 2013) and highlight its implications for the cognitive architecture of the mind.

3.2 Does epistemic vigilance affect the comprehension process?

My proposal is that, not only do epistemic vigilance mechanisms assess the believability of a piece of communicated information (as suggested by Sperber and colleagues), but they also contribute to the assessment of whether an interpretative hypothesis about the speaker's meaning is to be retained and attributed to the speaker as the intended interpretation. In other terms, they play a role in both 'acceptance' of content and the logically prior 'comprehension' of the speaker's meaning (her intended content).[5] That is, as well as assessing whether an interpretation attributed to the speaker (i.e. the output of the comprehension

---

[5] See Padilla-Cruz 2012 for a different proposal along the same line.

procedure) is allowed to enter the 'belief box' of the interpreter, epistemic vigilance mechanisms determine whether an interpretative hypothesis about the speaker's meaning ends up being the output of the comprehension module or not.

This suggestion is grounded on the well-established relationship between interpretation and trust, but it offers a new cognitively oriented perspective in which to frame such a relationship. Before exploring it, let us focus on the relevance-guided comprehension procedure, as presented in **Errore. L'origine riferimento non è stata trovata.**, and its stopping point (b).

(3) **Relevance-guided comprehension procedure**
   (a) Follow a path of least effort in computing cognitive effects: Test interpretative hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
   (b) Stop when your expectations of relevance are satisfied.

Sperber (1994a) suggests that the relevance guided comprehension procedure comes in three different versions: 'naïve optimism', 'cautious optimism' and 'sophisticated understanding'. Interestingly, the difference between the three strategies relies on different assumptions about the communicator's competence and benevolence, which in turn raise different expectations of relevance (hence determine different stopping points in interpretation). A naïvely optimistic hearer takes for granted that the communicator is behaving both benevolently and competently: he takes the communicator to be competent enough to avoid misunderstanding, and benevolent enough not to lead him astray. Thus he expects 'actual optimal relevance'. In contrast, a cautiously optimistic interpreter assumes the communicator to be benevolent, but not necessarily competent. As a consequence, he looks for 'attempted optimal relevance'. Finally, a sophisticated interpreter drops not only the assumption that the communicator is behaving competently, but also that she is behaving benevolently. Then the expectations of relevance that guide the comprehension procedure and determine its stopping point are expectations of 'purported optimal relevance'. The following table illustrates the three different versions of the relevance-guided comprehension procedure (which differ with regard to clause (b)):

Three versions of the relevance-guided comprehension procedure:

	(a)	Follow a path of least effort in computing cognitive effects: Test interpretative hypotheses in order of accessibility
Naïve optimism	(b₁)	Stop when your expectations of *actual optimal relevance* are satisfied (i.e. stop at the first relevant enough interpretation)
Cautious optimism	(b₂)	Stop when your expectations of *attempted optimal relevance* are satisfied (i.e. stop at the first interpretation that the communicator *might have thought* would be relevant enough to you)
Sophisticated understanding	(b₃)	Stop when your expectations of *purported optimal relevance* are satisfied (i.e. stop at the first interpretation that the communicator *might have thought would seem* relevant enough to you)

To appreciate the difference between these interpretative strategies, consider again example (1). If Robyn were a naively optimistic interpreter, she would attribute to Sarah the first interpretative hypothesis that is relevant enough to

her. The first interpretation that comes to Robyn's mind is that Neil$_1$ has broken his own leg this morning. This interpretation is relevant enough to Robyn, in fact, it is highly relevant to her (it allows her to derive many contextual implications, e.g. that she should immediately go to the hospital, cancel her afternoon meeting, etc.). Thus, a naïve interpreter would retain it and mistakenly attribute it to the speaker. But what if Robyn adopted the cautiously optimistic version of the relevance-guided comprehension procedure? Robyn would not take for granted Sarah's competence and she would be vigilant to the possibility that Sarah may not know what Robyn knows (and may consequently fail in her attempt to make the relevant information that she intends to convey more accessible than any other possible interpretation). Robyn would realise that Sarah could not have intended the interpretative hypothesis 'Neil$_1$ has broken his own leg this morning' to occur to her (precisely because she does not know that Robyn has a son). Sarah could not have thought that this interpretation would be relevant enough to her as she, Sarah, has no thoughts of any sort involving Neil$_1$. Thus, the comprehension procedure would go further and test the next most accessible interpretative hypothesis. For instance, it would access and assess the interpretation that Neil$_2$ has recently broken his own leg. Since Robyn takes it that Sarah might have thought this interpretation to be relevant enough to her (as in fact it is), the interpretation is selected and attributed to Sarah.

Utterance interpretation, thus, may depend on considerations about the speaker's competence (as in the example discussed above) and/or benevolence. The issue of what brings such considerations to bear on the interpretative process, however, has not been addressed within the literature. I suggest that epistemic vigilance mechanisms can modulate the hearer's expectations of relevance (i.e. from 'actual' to 'attempted' or 'purported' optimal relevance) and assess whether the interpretative hypothesis under construction satisfies these expectations. If the interpreter is vigilant towards the speaker's competence, for instance, he will expect 'attempted' optimal relevance'. As a consequence, he will stop at the first relevant interpretation that the speaker might have thought would be relevant to him (as described in the cautiously optimistic version of the relevance-guided comprehension procedure).

The issue to be addressed now is whether such extended interaction between the comprehension system and epistemic vigilance mechanisms is compatible with the thesis that pragmatics is a modular system. In what follows, I put forth a tentative suggestion to implement my proposal within a massively modular framework such as the one adopted by Relevance Theory.

From the perspective of the information flow through the architecture of the cognitive systems, the role of epistemic vigilance in the comprehension processes suggests that the epistemic vigilance module(s) does not receive its input only when the comprehension process is over. Rather, during the comprehension process, subparts of the interpretation are fed to the epistemic vigilance module(s) for its assessment. As a consequence, it may filter out interpretative hypotheses that are incompatible with the speaker's mental states (i.e. her beliefs and desires).

Consider an utterance of (1) in the context described above. Following a path of least effort, the interpreter starts "fleshing out" the propositional schema (delivered by the language decoding processor) 'Neil$_x$ has broken his$_y$ leg at some time t' by assigning to the proper name "Neil" the referent Neil$_1$ (this is because Robyn's concept of Neil$_1$ is more highly activated than her concept of

Neil$_2$). This part of the overall interpretative hypothesis is fed to the epistemic vigilance modules(s). It provides a hypothesised topic of conversation, that is, Neil$_1$, which in turn triggers epistemic vigilance mechanisms targeted at assessing the competence of the speaker (Sarah) on particular topics. These mechanisms access the piece of information that Sarah does not know that Robyn has a son called "Neil". Epistemic vigilance thus detects an incompatibility between the speaker's system of beliefs and the interpretative (referential) hypothesis under construction. As a consequence, it inhibits the comprehension procedure and prompts it to access (and assess) the next most accessible referential hypothesis.[6]

In general, epistemic vigilance mechanisms that monitor the speaker's competence and benevolence may restrict and direct the operations of the comprehension module. The role played by assumptions about the speaker's competence and benevolence in pragmatic interpretation has long been recognised (see Sperber 1994a) but its implications for the location of pragmatic abilities in the overall architecture of the mind have not been explored yet. I suggest that this role might be explained in terms of the interaction, as just discussed, between the comprehension module and epistemic vigilance mechanisms. I think this also opens up an interesting direction of research for developmental pragmatics, to which I now turn.

3.3 Developmental implications

Sperber (1994a) suggests that the three versions of the relevance guided comprehension procedure ('naïve optimism, 'cautious optimism' and 'sophisticated understanding') may correspond to different stages in pragmatic development. That is, children may start out as naïvely optimistic interpreters and progressively acquire the ability to monitor the speaker's competence and benevolence and to adapt their interpretative behaviours to these.

This gives rise to the interesting question of what allows the progression from naïve optimism to the further developmental stages. Sperber (1994a) claims that the three versions of the relevance-guided comprehension procedure require the interpreter to manipulate increasingly higher order representations of mental states. For this reason, Carston (1997) and Wilson (2005) have suggested that the development of more sophisticated interpretative strategies may correlate with the emergence of more complex mind-reading abilities: the move from naïve optimism to cautious optimism may correlate with the emergence of first-order mind-reading ability, the one from cautious optimism to sophisticated understanding with the emergence of second-order mind-reading abilities.

While this suggestion is certainly worth exploring further, the recent work on epistemic vigilance by Sperber *et al.* (2010) seems to open further interesting scenarios for developmental pragmatics. In light of the hypothesized interaction between the comprehension system and epistemic vigilance mechanisms, it seems plausible to assume that these three stages in the development of pragmat-

---

[6] A similar account is proposed by Sperber and Wilson (1986/1995: 186-87) with regard to the relationship between a decoding module and a central inferential system. Thanks to Deirdre Wilson for pointing this out to me and suggesting that the relation between decoding and inferential comprehension and between comprehension and epistemic vigilance mechanisms could be framed in the same way.

ic abilities may follow a similar developmental trajectory to that of epistemic vigilance capacities. Naïve optimism, cautious optimism and sophisticated understanding involve different assumptions about the communicator's competence and benevolence. As noted above, epistemic vigilance mechanisms focused on the source of information (*who* to believe) monitor the speaker's epistemic and moral reliability, that is, her competence and benevolence. Thus, it seems plausible to hypothesise that the emergence of epistemic vigilance mechanisms targeted at the assessment of the speaker's competence may correlate with (and prompt) the development of a cautiously optimistic interpretative strategy. Similarly, the move to sophisticated understanding may be triggered by the emergence of epistemic vigilance mechanisms monitoring the speaker's benevolence.

A very interesting and plausible picture emerges: the three interpretative strategies described above may be nothing but an *epiphenomenon* of the interaction between a *single* comprehension procedure and the epistemic vigilance mechanisms. On this construal, a 'cautiously optimistic interpreter' would be nothing but an interpreter who is vigilant towards the speaker's incompetence (i.e. her lack of knowledge or dependence on false belief), whereas vigilance towards the speaker's malevolence (i.e. her possibly deceptive intentions) would underpin the interpretative behaviour of a 'sophisticated interpreter'.

While the development of epistemic vigilance has been the subject of recent experimental investigation (e.g. Clèment, Koening and Harris 2004, Mascaro and Sperber 2009), an explicit comparison between the development of epistemic vigilance and of pragmatic competence remains to be carried out. Furthermore, there is a growing body of research on children's ability to track the communicator's epistemic state and use this to infer what she intends to refer to (Carpenter, Call and Tomasello 2002, Southgate, Chevallier and Csibra 2010, i.a.). The implications of this literature for Sperber's (1994a) theoretical distinction between 'naïve optimism', 'cautious optimism' and 'sophisticated understanding' have not been assessed yet. Crucially, Southgate *et al.* (2010) show that 17-month-old infants can take account of the speaker's epistemic state (i.e. her false-belief) in reference resolution. In Southgate *et al.*'s study, the infants see the experimenter place two novel objects in different boxes and leave the room. An accomplice changes the position of the objects in her absence. When the experimenter returns, she points towards one of the boxes and says to the infant: "Do you know what's in here? Shall we play with it? Shall we play with it? Let's play with it!" Finally she says, "Can you get it for me?" The issue here is which of the two objects the infant would take the word 'it' to refer to. The results showed that infants as young as 17-month-old were significantly more likely to choose the object in the box that the experimenter had not pointed to. What I would like to point out here is that the infant assigning the appropriate referent to the pronoun 'it' seems to require a cautiously optimistic interpretative strategy. That is, the infant should not stop at the first relevant referential interpretation (which corresponds to one where the referent is taken to be inside the pointed-to box), but rather at the first relevant referential interpretation that the experimenter could have thought would be relevant to him. This requires the infant to take account of the experimenter's epistemic state (i.e. her false belief that the intended object is in the pointed-to box) and reason that she could have not intended to refer to the object in the pointed-to box because she does not know that it has been swapped with the object in the non-pointed-to box. The

results raise the question of whether there is in fact any developmental stage corresponding to naïve optimism or whether naïve optimism is a theoretical construction without any empirical counterpart.

4. Conclusions

Relevance Theory has recently presented an account of pragmatics that is in line with an increasingly modular view of the mind. Pragmatics is conceived of as a sub-module of the mind-reading module, which exploits a dedicated inferential procedure that is attuned to the regularities of the domain of overt communication. As for any conceptual module, the comprehension module is part of a network of mental modules, which are highly interconnected with each other. This paper has focused on one of these connections, that is, the connection between the comprehension module and epistemic vigilance module(s). According to Sperber *et al.* (2010), epistemic vigilance mechanisms take as their proprietary input the output of the pragmatic module and assess its believability. In this paper, I make two further proposals: first, that epistemic vigilance mechanisms may affect not only the believability of the output of the relevance-guided comprehension procedure, but also the comprehension procedure itself; second, that the emergence of epistemic vigilance targeted at assessing the speaker's competence and benevolence may correlate with different developmental stages in pragmatics. These proposals suggest a programme of future research in cognitive and developmental pragmatics.

References

Carpenter, M., Call, J., and Tomasello, M. 2002, "A new false belief test for 36-month-olds", *British Journal of Developmental Psychology*, 20, 393-420.

Carston, R. 1997, "Relevance-theoretic pragmatics and modularity", *UCL Working Papers in Linguistics*, 9, 1-27.

Carston, R. 2002, *Thoughts and utterances: The pragmatics of explicit communication*, Oxford: Blackwell.

Carston, R. 2007, "How many pragmatic systems are there?", in Frápolli M.J. (ed.), S*aying, meaning, referring. Essays on the philosophy of François Recanati*, New York: Palgrave Macmillan, 18-48.

Clément, F., Koenig, M., and Harris, P. 2004, "The ontogeny of trust", *Mind & Language*, 19, 360-79.

Cosmides, L., and Tooby, J. 1992, "Cognitive adaptations for social exchange", in Barkow, J.H., Cosmides, L. and Tooby, J. (eds.), *The adapted mind: Evolutionary psychology and the generation of culture*, New York: Oxford University Press, 163-228.

Cosmides, L. and Tooby, J. 1994, "Origins of domain specificity: The evolution of functional organization", in Hirschfeld, L.A. and Gelman, S.A. (eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture*, Cambridge: Cambridge University Press, 85-116.

Fodor, J.A. 1983, *The Modularity of Mind*, Cambridge, MA: MIT Press.

Kasher, A. 1991, "On the pragmatic modules: A lecture", *Journal of Pragmatics*, 16, 381-97.

Mascaro, O. and Sperber, D. 2009, "The moral, epistemic and mindreading components of children's vigilance towards deception", *Cognition*, 112, 367-80.

Mazzarella, D. 2013, "'Optimal relevance' as a pragmatic criterion: the role of epistemic vigilance", *UCL Working Papers in Linguistics*, 25, 20-45.

Padilla Cruz, M. 2012, "Epistemic vigilance, cautious optimism and sophisticated understanding", *Research in Language*, 10 (4), 365-86.

Southgate, V. Chevallier, C. and Csibra, G. 2010, "Seventeen-month-olds appeal to false beliefs to interpret others' referential communication", *Developmental Science*, 13 (6), 907-12.

Sperber, D. 1994a, "Understanding verbal understanding", in Khalfa, J. (ed.), *What is Intelligence?*, Cambridge: Cambridge University Press,179-98.

Sperber, D. 1994b, "The modularity of thought and the epidemiology of representations", in Hirschfeld, L.A. and Gelman, S.A. (eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture*, Cambridge: Cambridge University Press, 39-67.

Sperber, D. 2001, "In defense of massive modularity", in Dupoux, E. (ed.), *Language, brain and cognitive development: Essays in honor of Jacques Mehler*, Cambridge, MA: MIT Press, 47-57.

Sperber, D. 2005, "Modularity and relevance: How can a massively modular mind be flexible and context-sensitive?", in Carruthers, P., Laurence, S. and Stich, S. (eds.), *The Innate Mind: Structure and Content*, New York: Oxford University Press, 53-68.

Sperber, D. & Wilson, D. 1986/1995, *Relevance: Communication and Cognition*, Oxford: Blackwell.

Sperber, D. and Wilson, D. 2002, "Pragmatics, modularity and mind-reading", *Mind & Language*, 17 (1&2), 3-23.

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G. and Wilson, D. 2010, "Epistemic vigilance", *Mind & Language*, 24 (4), 359-93.

Willis, J. and Todorov, A. 2006, "First impression: making up your mind after a 100-ms exposure to a face", *Psychological Science*, 17, 592-98.

Wilson, D. 2005, "New directions for research on pragmatics and modularity", *Lingua*, 115, 1129-1146.

Wilson, D. and Sperber, D. 1986, "Pragmatics and modularity", in Farley, A.M., Farley, P.T. and McCullough, K.E. (eds.), *Chicago Linguistic Society Parasession on Pragmatics and Grammatical Theory*, 67-84, reprinted in Davies, S. (ed.), *Pragmatics: A Reader*, New York: Oxford University Press, 1991, 583-95.