

# Feelings of (Un)Certainty and Margins for Error\*

Jérôme Dokic

*Abstract:* So-called epistemic or noetic feelings are often recruited in one's reasoning, and we may wonder how this recruitment is realized at the psychological level, and whether it is epistemologically warranted. I tackle these issues by focusing on feelings of subjective certainty and uncertainty in the context of ordinary perceptual categorizations. I first locate epistemic feelings within our cognitive architecture, by reference to the influential two-system framework of reasoning and decision-making as well as recent empirical models of our metacognitive abilities. I then put forward the thesis that in a normal context, feelings of perceptual certainty track the safety of our perceptual beliefs, whereas feelings of perceptual uncertainty track the fact that these beliefs are not safe. In other words, our felt certainty or uncertainty about the category of what we perceive is an indication of the fact that a margin for error has or has not been provided. I conclude by discussing two distinctions relevant to the account presented here, namely the distinction between perceptual and conceptual certainty (or uncertainty), and the distinction between objective and subjective certainty (or uncertainty).

## 1. Introduction

There is little doubt that ordinary theoretical and practical reasoning often relies on epistemic feelings, which include *feelings of knowing*, *feelings of doubt*, *feelings of truth*, *feelings of rationality*, *feelings of confidence*, and *feelings of (un)certainity*. Epistemic feelings are not mere epiphenomena, but can play a substantial role in the rational processes underlying our judgments and decisions. In deliberating, we sometimes exploit them as original (though defeasible) premises, which as a first approximation concern our own epistemic skills, states, and processes.<sup>1</sup>

\* I would like to thank the editors, Santiago Arango and Kirk Michaelian, for their helpful comments on an earlier draft of this essay.

<sup>1</sup> See de Sousa (2008) and Hookway (2008). Epistemic feelings are also called “noetic feelings” and “metacognitive feelings”; see Arango (2011; 2013) and Dokic (2012). Note that the use of the parenthesized phrase “(un)certainity” throughout this essay is not an homage to French Theory but should simply be read as an abbreviation of either “certainty and uncertainty” or “certainty or uncertainty” (depending on the context).

As an illustration, suppose that I am a participant in a TV game show. I have to press a certain key if the object that is briefly shown to me belongs to a given category (“picture of an animal”, say, or “picture of an artefact”). I receive some amount of money for a correct response, but I lose twice this amount for an incorrect response. So I decide to be careful. I form the intention to press the key only when I *feel certain* that the designated object belongs to the right category. In a scenario of this kind, my feeling of certainty intuitively contributes to the rationality of my action.

The general issues which form the background of this essay are why epistemic feelings are recruited in our reasoning and whether their recruitment is epistemologically warranted. In what follows I would like to make some preliminary remarks on these issues, and to explore the psychology and epistemology of epistemic feelings. My main focus will be on feelings of subjective certainty and uncertainty in the context of perceptual categorization.<sup>2</sup>

The essay is structured as follows. In the first two sections, I locate the relevant feelings within our cognitive architecture, by reference to the two-system framework of reasoning and decision-making as well as recent empirical models of our metacognitive abilities. In the next section, I introduce a few distinctions relevant to the characterization of the reliability of feelings of (un)certainty. In Sections 5 to 7, I suggest that, in normal contexts, feelings of perceptual certainty track the safety of our perceptual beliefs, whereas feelings of perceptual uncertainty track the fact that these beliefs are not safe. The epistemological notion of safety is related to so-called margin for error principles (at least when perception is at stake). Thus, our felt certainty (respectively, uncertainty) about the category of what we perceive is an indication of the fact that a margin for error has been provided (respectively, has not been provided). I surmise that the tracking properties of the relevant feelings contribute to explaining why they often provide premises to our reasoning, for instance in order to enhance the reliability of our judgmental skills when the context demands it. Section 8 contrasts perceptual and conceptual (un)certainty, and offers some speculation about whether feelings of conceptual (un)certainty can be seen as tracking the safety, or lack of safety, of beliefs grounded on non-perceptual processes. In Section 9, I comment on the distinction between objective and subjective (un)certainty, and suggest that the relevant feelings can sometimes reflect mere subjective or internal (un)certainty.

<sup>2</sup> In fact, I will only deal with visual categorizations, but at least some of my conclusions should extend to other sensory modalities. Indeed, some of David Smith’s uncertainty monitoring experiments (described in Section 5) involve auditory categorization tasks (for instance in dolphins).

## 2. *Epistemic feelings and dual cognition*

According to an influential framework of reasoning and decision-making, cognitive processes underlying judgments and decisions split into two kinds.<sup>3</sup> Although the details of this framework are a matter of great controversy, it is generally admitted that there is a distinction between automatic, largely non-conscious processes, which belong to “System 1” (henceforth S1), and more controlled, conscious processes, which belong to “System 2” (henceforth S2).<sup>4</sup>

S1 exhibits at least some symptoms of cognitive modularity, such as encapsulation and inaccessibility (Fodor 1983). First, S1 is encapsulated in the sense that its cognitive processes are indifferent to the subject’s background beliefs. Second, in a case of S1-type reasoning, the subject does not have access to the intermediary cognitive processes leading to the spontaneous outcome, whether it is a judgment or an action. The latter merely appears to the subject as the right thing to believe or to do in the relevant context. Whatever justification she may give for her judgment or action has a *post hoc* character, and may not correspond to what actually caused the outcome within S1.<sup>5</sup>

In the current understanding of the two-system framework, S1 is in fact a rag-bag collection of cognitive processes and activities that are defined in contrast to S2-type processes. It involves various kinds of heuristics and rules of thumb, innate as well as acquired modules, contextualized strategies, internalized short-cuts, and complex perceptual recognitional abilities. As many authors have observed (see, e.g., Evans 2003), S2 appears to be a more coherent category, linked to general intelligence and working memory abilities. It represents a higher form of rationality, and is the locus of abstract, analytic and hypothetical reasoning.

The relationship between S1 and S2 is still under discussion, but let me mention here what seems to be the most promising account, namely what Jonathan Evans calls the “default-interventionist” model. According to this model, S1 cues default behaviors that S2 may approve, inhibit, or otherwise intervene in with more effortful reasoning. S1 competes directly for control of our judgments and actions. Our default tendency is to stick to S1-type reasoning and behave in apparently spontaneous ways. Even when S1-type

<sup>3</sup> See for instance Stanovich and West (2000), Kahneman and Frederick (2002), Evans (2003; 2007; 2008), Gigerenzer (2007), Evans and Frankish (2008), Kahneman (2011).

<sup>4</sup> I use the standard terminology of “System 1” and “System 2” coined by Stanovich (1999), but as Evans and Frankish (2008) make clear in their introduction, the relevant distinction is better conceived as a distinction between two kinds of cognitive *process*.

<sup>5</sup> For a defense of the modularity of S1-type systems, as well as the (more controversial) claim that S2 is itself a special, metarepresentational module, see Sperber and Mercier (2008).

responses are inhibited by S2, they can persist as mere inclinations to act. To take a typical example from social psychology, I consciously know that I have no reason to doubt the sincerity of the person I am conversing with, yet (for reasons I do not fully understand) I may still be inclined to act as if she is insincere.

What is the place of epistemic feelings in this two-level cognitive architecture? Do they belong to S1 or S2? The answer is “both”. On the one hand, they are conscious states apt to participate in S2-type reasoning just like other consciously accessible beliefs or pieces of information. On the other hand, they are based on implicit processes characteristic of S1. For instance, I believe that I know a person’s name because I have a *feeling of knowing*. However, I do not know *why* I have this feeling, for it is causally grounded on unconscious monitoring mechanisms. In other words, epistemic feelings are *cross-level* states; they have one leg in S1 and the other in S2 (Koriat 2000; 2006).

One might object that the existence of cross-level states such as epistemic feelings blurs the very distinction between two systems of reasoning. For instance, Price and Norman claim that the existence of epistemic feelings makes “the inter-relation between the types of processes subsumed under S1 and S2 [...] so intimate as to melt the usefulness of the dichotomy” (2008: 37). A related objection is that epistemic feelings threaten the modularity of S1-type processes relative to S2-type processes. More precisely, they seem to show that intermediary cognitive processes within S1 are at least partially accessible from S2.

In response to the first objection, it should be noted that epistemic feelings do not jeopardize the autonomy of S1, which controls behavior most of the time. Indeed, no epistemic feelings need arise when the subject acts spontaneously – or at least, no feeling need *mediate* the implicit processes leading to her behavior.

The second objection is relevant only if we assume that S1-type systems are modular, at least in the sense that their intermediary cognitive processes are inaccessible to other cognitive systems. Even on this assumption, though, it is not clear that epistemic feelings are incompatible with *informational* inaccessibility. For at least some of them are based on sub-personal mechanisms that are sensitive to the quality of S1-type processes independently of their contents. Strictly speaking, then, epistemic feelings need not have access to the information that is processed within S1, but only to the *way* it is processed, i.e., to *non-semantic* properties. S1 may still remain informationally inaccessible from S2’s point of view.

It would be a mistake to conclude from the foregoing considerations that epistemic feelings are generated only in situations in which S1-type processes

are either inhibited by S2 or interrupted for another reason.<sup>6</sup> Of course this is what happens in many situations. For instance, the subject tries to remember a person's name, but fails to retrieve the relevant information. A characteristically frustrating feeling of knowing is then produced, namely the tip-of-the-tongue experience (Brown 2003). This feeling in turn motivates more sophisticated behavior, such as explicitly asking oneself questions of the form "Does this person's name begin with an X?" in order to provoke the correct response. In this case, the subject's feeling is subserved by monitoring an S1-type process, namely the spontaneous attempt to remember something. However, epistemic feelings might also arise from monitoring S2-type processes. For instance, explicit reasoning to a given conclusion can give rise to a feeling of truth or validity, which might be based on implicit monitoring of carefully controlled processes.

### 3. *Epistemic feelings and metacognition*

Epistemic feelings are also called "metacognitive feelings", but what does "metacognitive" mean in this context? Psychologists usually define metacognition as "cognition about one's own cognition".<sup>7</sup> Philosophers, on the other hand, tend to equate the notion of metacognition with that of metarepresentation, which refers to the ability to form representations about other mental representations, and is associated with possession of a "theory of mind".<sup>8</sup> Correspondingly, contents are metarepresentational when they are explicitly about representations as such. For instance, the content of the belief that Pierre believes that it is raining is metarepresentational, because of the presence in it of the mental state of believing that it is raining.

In what sense, then, can epistemic feelings be said to be metacognitive? In fact, two different senses are at stake, depending on whether we are talking about the implicit *causal bases* of epistemic feelings or their consciously experienced *contents*. Moreover, as we shall see, these two senses are relatively independent from each other.

The causal bases of epistemic feelings can be said to be metacognitive in the sense that they involve implicit monitoring mechanisms that are sensitive to properties of other cognitive processes. Many epistemic feelings seem to

<sup>6</sup> Such a view is suggested by Thompson (2007), at least in the case of what she calls "feelings of rightness".

<sup>7</sup> See Metcalfe and Shimamura (1994).

<sup>8</sup> Exceptions include Proust (2006; 2007), who sharply severs metacognition from metarepresentation. See also Arango (2011) and various essays in Beran *et al.* (2012).

be sensitive to the *fluency* of perceptual and/or conceptual processes.<sup>9</sup> For instance, feelings of knowing can be based on an implicit evaluation of the fluency of our spontaneous attempt to remember something.<sup>10</sup>

Note that the operations of the relevant monitoring mechanisms do not require metarepresentational abilities. To begin with, they are sensitive to properties of cognitive processing independently of the *contents* that are processed. Even if they involve representations of what are in fact other representations, they do not involve metarepresentations, i.e., representations of representations *as such*. Analogously, although a dog who sees the word “dog” on a blackboard may visually represent another representation, it does not metarepresent, for it does not represent the word “dog” as something that itself represents something else.

We should carefully distinguish between judgments of knowledge based on epistemic feelings (which are experience-based judgments) and those based on explicit reasoning (theory-based judgments). As Koriat writes, “whereas the latter judgments draw upon the *content* of domain-specific beliefs and knowledge that are retrieved from memory, the former rely on contentless mnemonic cues that pertain to the quality of processing, in particular, the fluency with which information is encoded and retrieved.” (2006: 19-20).

There is another, more speculative reason for doubting that implicit metacognitive mechanisms require metarepresentational abilities. It can be argued that by definition metarepresentations are either actually or potentially conscious. Metarepresentations rely on a relatively sophisticated conception of mental representation, which may be available only to conscious creatures and not to sub-personal mechanisms. In contrast, the metacognitive mechanisms underlying epistemic feelings involve only implicit representations, which are not and cannot become conscious. These are first-order representations which happen to be about internal rather than external states. They are first-order yet self-directed, as opposed to world-directed.

The other sense in which epistemic feelings can be conceived as metacognitive concerns their apparent representational contents, which determine how they can rationally guide decision-making and the fixation of belief. The question is whether these contents are explicitly about one’s own epistemic states and processes. Here, we face two alternatives. The first alternative gives a “yes” answer to this question, and maintains that epistemic feelings are or involve metarepresentations. For instance, the feeling of knowing a person’s name is what it seems to be, namely the feeling *that one knows this person’s name*. Its

<sup>9</sup> See, e.g., Jacoby and Dallas (1981), Wurtz *et al.* (2007), and Oppenheimer (2008).

<sup>10</sup> See Koriat (2005: 312).

content is metarepresentational insofar as it involves an epistemic state, namely one's knowledge of the name (see Dienes and Perner 1999). The second alternative promotes a "no" answer to the question, and insists that we can exploit epistemic feelings in reasoning without deploying metarepresentational abilities. At this point, the challenge is to show that at least some epistemic feelings really have first-order contents.<sup>11</sup>

It is important to realize that the two senses in which epistemic feelings can be described as metacognitive are largely independent. Even if one acknowledges that the causal basis of some epistemic feeling is sensitive to non-semantic properties of lower level processes, the question of whether this feeling has a metarepresentational content at the personal level remains entirely open.

#### 4. *The reliability of epistemic feelings*

The fact that epistemic feelings are used in ordinary reasoning as (defeasible) premises that somehow concern our own epistemic skills might raise a puzzle. If epistemic feelings are sensitive to non-semantic properties of cognitive processing, how could they have any content concerning our own epistemic predicament, let alone any epistemic value with respect to such a content?

In order to answer this question, we should distinguish three sets of facts about a given implicit metacognitive mechanism MM:

1. The facts to which MM is directly sensitive.
2. The facts that are tracked by MM (relative to normal contexts).
3. The facts that are explicitly associated with an epistemic feeling arising from the operations of MM.

As an illustration, consider feelings of knowing. First, we can assume that their underlying metacognitive mechanisms are directly sensitive to something like the quality of the attempt to remember some piece of information, like the name of one's interlocutor. As Koriat has argued, these mechanisms need not monitor the memory trace itself. They can be causally disconnected from the subject's relevant state of knowledge. Insofar as the notion of sensitivity is a

<sup>11</sup> A possible strategy would be to argue that what appears to be metarepresentational information is carried at the level of the intentional *mode* of the feeling rather than explicitly represented at the level of its content. For instance, the content of the feeling of uncertainty relative to the state of affairs *p* is not *that I feel uncertain that p*, but simply *p* itself. The relevant attitude is feeling-uncertain(*p*) rather than feeling(uncertain that *p*). My main worry with this strategy is that it does not explain what premises feelings of uncertainty add to explicit reasoning. Of course it cannot be the premise that *p* itself.

causal-informational one, feelings of knowing are not directly sensitive (do not have direct access) to the presence of the name in long-term memory.<sup>12</sup>

Second, as psychologists commonly acknowledge (see, e.g., Koriat 2006), feelings of knowing are by and large reliable. In a normal context, I would not have a feeling of knowing a person's name if I did not possess the relevant knowledge. In other words, feelings of knowing can “track” states of knowledge, in the sense that (roughly) the former normally occur only in the context of the latter. Unlike the notion of sensitivity, the notion of tracking does not require that there be a causal-informational link between what is tracked and the tracking mechanism. (Sensitivity is a species of tracking, but there can be tracking without sensitivity.) The implicit cues on which my feeling is based are contingently but stably associated with my state of knowledge.<sup>13</sup> This association holds in a normal (ecological) context, but it can be artificially severed by psychologists, who can easily produce “illusory” feelings of knowing (Bjork 1999).

The third set of facts determines what contents epistemic feelings have for the subject who uses them in her reasoning. On this count, psychologists have shown that the contents implicitly or explicitly attributed to epistemic feelings are highly flexible and context-dependent, even in normal contexts. They can vary with education, past associations, and local saliency effects. For instance, the same type of fluency-based feeling can be attributed to features as diverse as the aesthetic value of the object under consideration, the truth of a proposition, or the validity of an inference.<sup>14</sup>

Still, the fact that the recruitment of epistemic feelings in theoretical or practical reasoning is often adaptive (see next section) shows that it is generally *constrained* by what epistemic feelings objectively track. Although the full epistemological implications of this constraint remain to be articulated, the core idea is that part of the explanation of why epistemic feelings can be recruited as symptoms of states of affairs beyond the body is that they co-vary reliably with the latter, and the subject is at least implicitly aware of this correlation. It does not follow that the subject must be able to represent her own state of

<sup>12</sup> Of course, other types of epistemic feelings may be such that their underlying metacognitive mechanisms are directly sensitive to the relevant target in memory. Metcalfe (2000) argues that this is the case with “feelings of imminence”, such as those involved in tip-of-the-tongue experiences.

<sup>13</sup> So my feeling can be considered as an “internal local sign” in something like Millikan's (2004) sense.

<sup>14</sup> See Unkelbach and Greifeneder (2013). One hypothesis (see Dokic 2012) is that epistemic feelings do not have *intrinsic* contents beyond the body. In themselves, they are nothing but subtle bodily experiences, and do not disclose epistemic states and processes in the way perceptual experiences disclose the sensible world to us.

knowledge as such. For instance, a given feeling of knowing can be recruited as the feeling that I *know* the answer to this question but also, perhaps more cautiously, as the feeling that I *can* answer this question. The latter content, unlike the former, is not obviously metarepresentational.

### 5. (Un)certainty, safety and margins for error

In the remainder of this essay, I would like to contribute to the analysis of the epistemic value of feelings of certainty and uncertainty by using tools from formal epistemology. More precisely, I will offer a general suggestion about what facts are (normally) *tracked* by such feelings.

As a test case, consider one of David Smith's numerous experiments on uncertainty monitoring in both humans and non-human animals (Smith *et al.* 2003; Smith 2005). In one (forced-choice) condition, the subject has to touch a visual pattern on the screen when it is judged to be dense, and the symbol "S" when the pattern is judged to be sparse instead. In another condition, the subject is also allowed to press a third, so-called "uncertainty" key, which simply advances her to the next trial. Humans and some species of non-human animals (such as rhesus monkeys) are able to use the uncertainty key adaptively by reducing the number of errors that they would make in the forced-choice condition.

As Smith and his collaborators have observed, human subjects spontaneously explained that they hit the uncertainty key when they were "unsure", or "felt uncertain" about the category of the stimulus. We may speculate that some non-human animals, who showed psycho-physical patterns very similar to those of human subjects in doing the task, were also able to exploit their feelings of uncertainty in more sophisticated practical reasoning (though not in a metarepresentational way). We may say that the successful subjects felt either competent or incompetent to press the right key, whether the correct response was in fact "dense", "sparse", or neither.<sup>15</sup>

Success in doing the categorization task is indeed guaranteed if the subject *knows* that the stimulus falls under the appropriate category. Knowledge involves a safe belief rather than a lucky guess. Mark Sainsbury helpfully characterizes a lucky guess as being, "not a proposition which might easily not have

<sup>15</sup> For simplicity's sake, I shall assume that the relevant concepts *dense* and *sparse* are *vague*, in the sense that there are borderline cases of dense and sparse patterns. In contrast, Smith and his collaborators consider that the "dense" response is correct if and only if the visual pattern contains *exactly* 2950 illuminated pixels. The subjects are not informed of this precise definition of "dense", so that it is somewhat unclear what concepts they operate with when they do the categorization task. This should not affect either Smith's or my main conclusions, though.

been true, but a *way of reaching a belief* which might easily not have delivered a true one” (1995: 595; my italics). Conversely, a belief is safe when it is obtained by a *mechanism* (or *method*) that delivers true beliefs not only in the actual world but also in nearby possible worlds. This makes safety relative to a mechanism or method of knowledge:

*Mechanism-safety:*

The true belief B is safe iff the mechanism that produced B also produces true beliefs in nearby possible worlds.<sup>16</sup>

In some contexts, safety can be shown to be associated with so-called “margin for error” principles.<sup>17</sup> Consider the visual recognitional mechanism M that produces beliefs about the density of the stimulus (seen in optimal lighting conditions). M is simply whatever mechanism in the brain underlies conscious visual perception of an object as being dense or sparse. Suppose that, as a result of M, I truly believe that the stimulus *i* in front of me is dense. My belief is safe iff M might easily not have produced false beliefs. Now this requires that stimuli that are visually indiscriminable in density from *i* be dense too. In other words, M is subject to the following margin for error principle:

*Margin for error principle:*

If one knows by visual means that *i* is dense, *i*-1 and *i*+1 are also dense.

Here, *i*-1 and *i*+1 are possible stimuli adjacent to *i* in a series of progressively denser patterns. Let us assume that the subject can visually discriminate between them, but cannot discriminate between either of them and *i*. The argument for the margin for error principle goes as follows. Suppose that although *i* is dense, *i*-1 is not.<sup>18</sup> It follows that even though M produced the true belief that *i* is dense, it might easily have produced the false belief that *i*-1 is dense (since M would have also classified *i*-1 as dense). In this case, my belief that *i* is dense is true, but it is too risky to constitute knowledge. It does not leave a margin for error.

<sup>16</sup> Similar principles are discussed by Peacocke (1986) and Weatherson (2004). Other classical discussions of safety include Pritchard (2007) and Sosa (2009).

<sup>17</sup> See Williamson (1994; 2000), Engel (2008), Dokic and Égré (2008). The latter work argues against Williamson’s claim that the margin for error principle is incompatible with the KK principle, according to which if one knows that A, one knows that one knows that A.

<sup>18</sup> There is no commitment here to the view that there is a sharp boundary between dense and sparse patterns (although it is sometimes convenient to pretend that this is so). For a defence of such a view, as known as the epistemic conception of vagueness, see Williamson (1994; 2000). As Williamson himself makes clear, margin for error principles are independent from the epistemic conception of vagueness.

To sum up, competence at a visual categorization task hangs on the subject's knowing, and thus safely believing, that the stimulus falls under the right category. This in turn implies that a margin for error has been provided.

## 6. *Tracking safety*

What I want now to suggest is that feelings of certainty, or more precisely their underlying metacognitive mechanisms, track the fact that the subject's perceptual belief is safe. Normally, when she feels certain about the category of the stimulus, her belief that the stimulus belongs to the category is safe, and thus leaves a margin for error.

One may ask whether *M* is a perfectly reliable recognitional mechanism, in the sense that it always produces safe beliefs about the density of the stimulus. The answer will be "no" if *M* sometimes overshoots and classifies some stimuli as dense while the subject does not know them to be dense. For instance, a marginally dense stimulus can still be perceived as dense, even though it is too close to the boundary between dense and sparse patterns for there to be a margin for error. As a consequence, the perceptual belief that the stimulus is dense is true but not safe; the subject does not *know* that it is dense.

Even though this is partly an empirical matter, I suspect that our visual recognitional mechanisms are not perfectly reliable in this sense. We can then speculate that some feelings of certainty have evolved or developed to *enhance* the reliability of our perceptual beliefs. Suppose that a stimulus *i* is clearly dense. Not only does it visually *seem* dense, but I may also feel *certain* that it is dense. The suggestion is that my feeling of certainty tracks the fact that a belief about the density of *i* based on my visual experience is safe, i.e., that it leaves a margin for error. In other words, it tracks the fact that no "small" change (including non-noticeable ones) could turn *i* into something sparse. In general, whenever I feel certain that *n* is dense, my belief that *n* is dense is safe. The mechanism underlying my feeling of certainty – let us call it *MM* – produces a response *only if*  $K(p)$ .<sup>19</sup>

Whether or not *M* is perfectly reliable, *MM* may remain *silent* in many cases in which my visual abilities classify (safely or not) the presented object as dense. In particular, the present account does not entail that there is a situation *S<sub>n</sub>* in which a subject feels certain, and thus knows, that the stimulus *n* is dense, while being unable to know that the possible stimulus *n-1* is dense because *n-2* is sparse. In other words, there is no commitment to the claim that *MM* produces a response *if* and only if  $K(p)$ .

<sup>19</sup> I deliberately ignore the existence of "illusory" feelings. A more careful claim is that the *function* of feelings of certainty is to track safety, which allows for particular feelings to be misleading.

If the function of all feelings of certainty is to track the presence of a margin for error, perhaps we can consider that *degrees* of felt certainty indicate the relative modal distance between the stimulus and the boundary between cases in which the margin for error principle is satisfied and cases in which it is not. The more certain we feel, the less there is a risk that the margin for error principle is not satisfied. In a Williamsonian framework, a strong feeling of certainty that  $p$  should correlate, not only with  $K(p)$ , but with  $K...K(p)$ , even if the precise number (if there is one) of the iterations of the knowledge operator  $K$  cannot be deduced from the feeling itself.

On this account, feelings of certainty can be exploited by the subject to inhibit spontaneous, S1-type actions based on visual categorizations, such as pressing one key rather than another. Suppose, for instance, that the subject decides to act in accordance with the reflective rule that she should follow her inclinations only if she is *quite certain* about the matter. Perhaps the stakes of the situation are very high, and the consequences of a mistake would be disastrous. Then she can use her conscious feelings of certainty at the level of S2 and control the outputs of S1 by filtering out some spontaneous actions based on visual categorizations (even if the latter are in fact safe).

Psycho-physical experiments show that our spontaneous visual categorizations, as they can be revealed for instance in forced-choice conditions, are more frequently correct than we tend to think. This is the well-known phenomenon of *perceptual underconfidence* (Björkman *et al.* 1993). One way of interpreting this phenomenon is that one tends to self-report visual categorizations (presumably a S2-type achievement) only when one has some feeling of certainty, or at least when one does not feel any uncertainty at all, about the category of the stimulus.

It is essential to the present account that M and MM are *different* cognitive mechanisms. For there will be a pair of situations  $S_n$  and  $S_{n-1}$  between which M cannot discriminate (because both stimuli  $n$  and  $n-1$  seem dense) but between which MM can discriminate, eliciting a feeling of certainty in  $S_n$  but not in  $S_{n-1}$ . This should be possible if MM can operate independently of the *contents* of the visual experience itself (the fact that the object is dense rather than sparse, for instance).

An empirically plausible suggestion, which might draw on an analogy with metamemory, is that MM is a metacognitive mechanism, or set of mechanisms, capable of monitoring the *quality* of the subject's visual categorizations. How these mechanisms actually operate is still an open question. We may speculate that they somehow take into account the stability of the visual system's responses and levels of activation in the appropriate neural networks. Stimulus cues such as contrast, crowdedness or blurriness are known to be

relevant,<sup>20</sup> but other independent factors have been put forward. For instance, within a Bayesian approach to perception, visual subjective uncertainty has been represented as a function of both sensory uncertainty and expected uncertainty (so-called “priors”) (Barthelmé and Mamassian 2010; Mamassian and Barthelmé 2011).

### 7. *Tracking lack of safety*

Given the foregoing account of feelings of certainty, it is natural to construe feelings of uncertainty as tracking the fact that the subject’s perceptual belief is *not* safe, and thus that the margin for error principle is *not* satisfied. Whenever the subject feels uncertain about the category of the stimulus, any belief that she might form on the basis of her visual experience would not be safe, because a margin for error is not provided.

A subject can see a stimulus under perfect viewing conditions and yet be unable to tell for sure whether it is dense or sparse. In such a situation, she may feel uncertain that the stimulus is one way or the other. Suppose, for instance, that the subject’s recognitional mechanism *M* rapidly oscillates between “dense” and “sparse” responses, in a situation in which the stimulus itself is constant. On the basis of her visual experience, she is inclined to believe that the object is dense but, at the same time, she is also inclined to believe that it is sparse. This supposition makes sense only if *M* is not perfectly reliable, in the sense introduced in the previous section, for neither the belief that the stimulus is dense nor the belief that it is sparse is safe.

Let us say that when such a situation occurs, the subject is in a narrow state of uncertainty. Now we might consider the possibility that her brain hosts a mechanism that is *sensitive* to whether she is in such a state.<sup>21</sup> Such a mechanism would be able to track a situation in which the margin for error principle is not satisfied relative to both the (possible) belief that the stimulus is dense and the (possible) belief that it is sparse. Insofar as the relevant mechanism produces a characteristic feeling that can be exploited in explicit reasoning, it grounds the ability to detect borderline cases, such as cases at the boundary between dense and sparse patterns. When a reflective subject seeing a border-

<sup>20</sup> Note that these visual cues can still be conceived as non-semantic, to the extent that they concern the proximal stimulus and not the conscious content of visual experience.

<sup>21</sup> Carruthers (2008) describes one such mechanism, which he calls “the gate-keeping mechanism”: “when confronted with conflicting plans that are too close to one another in strength [it] will refrain from acting on the one that happens to be strongest at that moment, and will initiate alternative information-gathering behaviour instead” (2008: 66). For critical discussion of Carruthers’ proposal, see Dokic (2012: 313-315).

line stimulus *i* is aware of the ambiguity of her visual responses, she should be able to work out the following conditionals: if *i* is dense, *i*-1 is sparse, and if *i* is sparse, *i*+1 is dense. This does not violate any margin for error principle, for the subject does not know, at least by visual means, whether the stimulus *i* is in fact dense or sparse.

### 8. *Perceptual vs. conceptual (un)certainty*

Our discussion so far has focused on feelings of (un)certainty with respect to perceptual categorizations – what we may describe as feelings of *perceptual* (un)certainty. Now we should acknowledge not only feelings of perceptual (un)certainty but also feelings of *conceptual* (un)certainty. We can have a feeling of conceptual (un)certainty relative to a conclusion that we have just drawn, or more precisely to the non-perceptual method we have used to reach that conclusion.

It is plausible, although further elaboration is certainly needed, that feelings of conceptual certainty track the safety of the relevant belief. In other words, they track the fact that the method used to reach that belief is reliable: it produces true beliefs also in nearby possible worlds. Analogous remarks apply, *mutatis mutandis*, to feelings of conceptual uncertainty. Whether margin for error principles apply in these cases and what form they should take are important questions beyond the scope of this essay. Of course they cannot take the form that they have in the perceptual case. Suppose that I reach the conclusion that the stimulus *i* is dense on non-perceptual grounds – for instance, I deduce it from testimony. I may know that *i* is dense by these means even if *i*-1 is not dense. This is not a counterexample to the margin for error principle stated in Section 5 above since the latter principle is about visually acquired knowledge. In other words, what counts as a nearby possible world can vary depending of the method used to produce the relevant belief.

It is equally plausible that feelings of conceptual (un)certainty track safety, or lack of safety, by being directly sensitive to the fluency of the cognitive processes leading to the belief. Psychologists working on metacognition draw a distinction between kinds of fluency depending on the underlying processes at stake, so that metacognitive mechanisms can monitor either perceptual processes (perceptual fluency) or processes concerned with conceptual content and semantic knowledge (conceptual fluency).<sup>22</sup> In the context of the two-system framework, we should draw a further distinction between two kinds of

<sup>22</sup> See, for instance, Bullock and Reber (2013).

feelings of conceptual (un)certainty, or at least two different cases in which such feelings can arise. Suppose that I feel certain that a given judgment of mine, say about a mathematical equation, is true. The judgment itself can be either spontaneous or explicitly inferred from other mathematical premises. In the first case, my feeling of conceptual certainty is based on monitoring the fluency of the S1-type heuristic process underlying my judgment. In the second case, my feeling of conceptual certainty is based on monitoring the fluency of S2-type analytic processes, which correspond to the carefully controlled steps leading to my judgment.

One might wonder what the relationship is between perceptual and conceptual (un)certainty. Can a feeling of perceptual certainty co-exist with a feeling of conceptual uncertainty, or vice versa? For instance, can I feel perceptually certain that what I am seeing is a horse, while feeling conceptually uncertain about that (I may have some reason to think that it is a donkey cleverly disguised as a horse)? Alternatively, can I feel perceptually uncertain that what I am seeing is a horse (it is too far away, or there is a mist), while feeling conceptually certain that it is a horse (perhaps because I remember having seen it closer)?

Part of what is at issue here is the modularity of the mechanism monitoring feelings of perceptual (un)certainty. If this mechanism is indeed modular, it will produce either a feeling of certainty or a feeling of uncertainty whatever further, higher-level feelings are experienced by the subject (if any). Alternatively, if this mechanism is not modular, it will be sensitive to the quality of non-perceptual processes as well, and will produce an overall feeling of certainty or uncertainty.

There is some reason to think that feelings of perceptual (un)certainty are based on relatively modular mechanisms, which are sensitive only to perceptual processes. These mechanisms produce characteristic spontaneous behavior, whose function is to optimize the subject's perceptual experience, such as moving one's head from side to side to resolve visual ambiguity. However, it does not follow that the subject can have a cognitive impression of perceptual certainty (respectively uncertainty) co-existing with a cognitive impression of conceptual uncertainty (respectively certainty). It is probable that the subject will associate a single content to her overall bodily experience instead of attributing two contrary contents to different aspects of this experience. In this respect, the situation is very unlike that of perceptual illusions, such as the Müller-Lyer illusion, where a perceptual impression (that two lines are unequal) co-exists with a contrary cognitive impression (that these lines are in fact equal).

An analogy might be helpful at this point. Bullot and Reber (2013; see also Reber *et al.* 2004) argue that feelings of beauty or at least aesthetic preference are sensitive to the fluency of the cognitive processes at stake. Now given that

visually disfluent artworks (for instance, Bridget Riley's paintings, which in a sense are visually irritating) can provoke positive aesthetic experience (at least in some of us), these authors suggest that high conceptual fluency, grounded for instance on intellectual knowledge of the artwork and art history, can override disfluency at lower levels. On their account, what counts for the feeling of beauty is *overall* processing fluency. Thus, when we face the painting called "Movement in Squares" (1961), we do not enjoy two aesthetic feelings: a feeling of visual ugliness accompanied with a feeling of intellectual beauty. Rather, a single aesthetic content is attributed to our whole bodily experience, which is recruited in the formation of aesthetic judgments.

### 9. *Objective vs. subjective uncertainty*

It might be objected that the present account of felt uncertainty does not deal with the important distinction between objective and subjective uncertainty. There is objective uncertainty when the state of the world is not fully determined, for instance when it is about the future (if the latter is considered to be metaphysically open). By contrast, subjective uncertainty is about one's own cognitive state, and can arise even in situations in which the relevant state of the world is fully determined at the metaphysical level. Proust (2007) gives the following utterances as expressing respectively objective and subjective uncertainty:

- (1) The circulation is so difficult that I don't know any more whether Peter will come.
- (2) I am so confused that I don't know any more whether Peter will come.

Now as described above, the mechanism underlying feelings of uncertainty seems to be indifferent to the distinction between objective and subjective uncertainty. It detects narrow states of uncertainty, i.e., competing or hesitating responses that can be due either to internal perceptual noise or to slightly changing stimuli. Genuine feelings of uncertainty, one might argue, are feelings of *subjective* uncertainty.

However, feelings of uncertainty might be triggered in situations in which there are no narrow states of uncertainty. Suppose that the recognitional mechanism *M* does not "hesitate" between several responses. The stimulus is not ambiguous; it visually seems more dense than sparse. However, the margin for error principle is not satisfied, because the stimulus is too close to the boundary between dense and sparse patterns. In this situation, the subject may feel

a bit uncertain about the category of the stimulus, even though the visual response is stable. In general, a feeling of uncertainty as to whether  $p$  might coexist, in a particular context, with a weak belief that  $p$ , even in the absence of a comparably weak belief to the contrary. This comes close to tracking the fact that the subject's spontaneous categorization in this context resembles guessing more than knowing, and that she is in a state of *subjective* uncertainty.

In the context of Smith's experiments, we can imagine that the subject is naturally inclined to press the "dense" key yet refrains from doing so because she feels uncertain about it. A related strategy consists in not pressing the "dense" key in the absence of a feeling of certainty. These strategies might yield slightly different results if there is a middle zone between feelings of certainty and feelings of uncertainty, where the subject feels neither certain nor uncertain about the category of the stimulus. In fact, we already allowed for the possibility that there is such a zone by acknowledging that the margin for error can be provided without felt certainty. Analogously, there might be cases in which the margin is not provided but the subject does not feel uncertain.

It does not follow that the subject *knows* the boundary between feelings of certainty and the middle zone, or between the latter and feelings of uncertainty. In general, the relevant feelings might not be "luminous" in Williamson's (2000) sense. A mental condition is luminous when as soon as one enjoys it, one *knows* that one is enjoying it. For instance, if one has a feeling of certainty relative to  $p$ , one would immediately know that one feels certain that  $p$ . Arguably, one can feel certain or uncertain without possessing the concept of feeling, which is required for knowing that one has feelings at all. However, feelings of certainty and uncertainty are not luminous even if the subject actually possesses the concept of feeling. The state of knowing that one feels certain might bring in its own zone of uncertainty. Perhaps one can be uncertain about whether one feels certain about the category of the stimulus – a higher-order form of uncertainty. In this case, some feelings of certainty will be such that we cannot know them to be feelings of certainty, perhaps because they are too close to the phenomenological boundary between feelings of certainty and other types of feelings (or no feeling at all). Similar remarks apply to feelings of uncertainty.<sup>23</sup>

There is no reason to exclude *a priori* the possibility that different feelings of uncertainty are based on different implicit mechanisms. A plausible suggestion is that a mechanism capable of monitoring the quality of the subject's perceptual response will be common to feelings of certainty and at least some feelings of uncertainty. An additional mechanism, such as Carruthers'

<sup>23</sup> For further discussions of the effects of iterating certainty and uncertainty, see Égré and Bonnay (2012).

(2008) gate-keeping mechanism, might have the function of detecting narrow states of uncertainty, in which two contrary beliefs compete for behavior. This mechanism may also generate a feeling of uncertainty. Since it has a distinctive phenomenological profile, we might describe it as a feeling of hesitation or puzzlement, whether it is caused by either merely subjective or both subjective and objective uncertainty (i.e., internal noise or external small changes).

## 10. *Conclusion*

This essay was about a family of epistemic feelings, namely those which have to do with subjective certainty or uncertainty, especially in the context of perceptual categorization tasks. Drawing on both empirical research and formal epistemology, I have suggested that the relevant feelings of certainty (respectively uncertainty) are phenomenological reflections of the safety (respectively lack of safety) of the subject's spontaneous perceptual beliefs. This claim relies on a realist account of at least a core component of perceptual knowledge, according to which there is a fact of the matter as to whether a given perceptual belief is safe and leaves an appropriate "margin for error".

On the present account, the main epistemological advantage of recruiting feelings of (un)certainly in our reasoning is to enhance the reliability, not of perception itself, but of perceptual beliefs. Depending on the stakes, we can rationally form different perceptual beliefs on the basis of one and the same sensory experience, by drawing on our feelings about the trustworthiness of our perceptual categorizations.

One might claim that epistemic feelings not only track states of knowledge, but are intentionally about them. However, the fact that subjects discriminate between knowledge and ignorance shows at best that they know *when* they know (at least sometimes), and can behave adaptively because of this knowledge, without necessarily knowing *that* they know, i.e., without deploying the concept of knowledge (or directly related concepts, such as perception). Since the safety of a subject's perceptual belief is an essential aspect of what guarantees success in doing the categorization task at hand, her feeling of certainty can be interpreted as an explicit indication of her competence, rather than a conscious reflection on her epistemic predicament. To feel certain may just be to feel competent to *do* something. Depending on the subject's conceptual sophistication, this may or may not be re-interpreted metarepresentationally, as the feeling that one knows something.

Obviously, what I have offered here is at best a preliminary analysis of the epistemic value of feelings of (un)certainly. Much more has to be said about the

epistemological contribution of epistemic feelings to theoretical and practical reasoning. However, I hope to have made clear that an adequate analysis should be constrained by an account of what epistemic feelings normally track, which partly explains why we are entitled to let them modulate our judgments and decisions. This is an instance of a general strategy, according to which epistemic properties such as justification, entitlement or knowledge are shown to be grounded, at least partly (if not wholly), on ontological features of our mental states.<sup>24</sup>

### References

- Arango-Muñoz, Santiago, 2011, “Two Levels of Metacognition”, in *Philosophia*, 39, 1: 71-82.
- Arango-Muñoz, Santiago, 2013, “The Nature of Epistemic Feelings”, in *Philosophical Psychology*: 1-19.
- Barthelmé, Simon, Pascal, Mamassian, 2010, “Flexible mechanisms underlie the evaluation of visual confidence”, in *Proc. Natl. Acad. Sci.*, 107, 48: 20834-20839.
- Beran, Michael J., Johannes, Brandl, Josef, Perner, Joëlle, Proust (eds.), 2012, *Foundations of Metacognition*, Oxford University Press, Oxford.
- Bjork, Robert A., 1999, “Assessing our own competence: Heuristics and illusions”, in Daniel Gopher, Asher Koriat (eds.), *Attention and performance XVII: Cognitive regulation of performance: Interaction of theory and application*, MIT Press, Cambridge, MA: 435-459.
- Björkman, Mats, Peter, Juslin, Anders, Winman, 1993, “Realism of confidence in sensory discrimination: The underconfidence phenomenon”, in *Perception and Psychophysics*, 54: 75-81.
- Brown, Alan S., 2003, “A review of the déjà vu experience”, in *Psychological Bulletin*, 129: 394-413.
- Brun, Georg, Ulvi, Doguoglu, Dominique, Kuenzle (eds.), 2008, *Epistemology and emotions*, Ashgate, Aldershot.
- Bullot, Nicolas J., Rolf, Reber, 2013, “The Artful Mind Meets Art History: Toward a Psycho-Historical Framework for the Science of Art in Appreciation”, in *Behavioral and Brain Sciences*, 36, 2: 123-137.
- Carruthers, Peter, 2005, *Consciousness. Essays from a Higher-Order Perspective*, Clarendon Press, Oxford.
- Carruthers, Peter, 2006, *The architecture of the mind*, Clarendon Press, Oxford.

<sup>24</sup> Michaelian (2012) can be seen as another relevantly close illustration of the same strategy. He offers an ontological, metacognitive solution to an epistemological question (which he calls “the endorsement problem”), namely, “How can beliefs produced by endorsing information received from imperfect sources be formed in a justification-conferring manner?” (p. 284).

- Carruthers, Peter, 2008, "Meta-cognition in Animals: A Skeptical Look", in *Mind and Language*, 23, 1: 58-89.
- de Sousa, Ronald, 2008, "Inference and epistemic feelings", in Brun, Doguoglu, Kuenzle (eds.): 185-204.
- Dienes, Zoltan, Perner, Josef, 1999, "A theory of implicit and explicit knowledge", in *Behavioral and Brain Sciences* 22: 735-808.
- Dokic, Jérôme, 2012, "Seeds of Knowledge: Noetic Feelings and Metacognition", in Beran, Brandl, Perner, Proust: 302-320.
- Dokic, Jérôme, Egré, Paul, 2009, "Margin for error and the transparency of knowledge", in *Synthese*, 166: 1-20.
- Égré, Paul, Bonnay, Denis, 2012, "Metacognitive perspectives on unawareness and uncertainty", in Beran, Brandl, Perner, Proust: 322-341.
- Engel, Pascal, 2008, *Va Savoir!*, Hermann, Paris.
- Evans, Jonathan St. B.T., 2003, "In two minds: dual-process accounts of reasoning", in *Trends in Cognitive Sciences*, 7, 10: 454-59.
- Evans, Jonathan St. B.T., 2007, "On the resolution of conflict in dual process theories of reasoning", in *Thinking & Reasoning* 13, 4: 321-39.
- Evans, Jonathan St. B.T., 2008, "Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition", in *Annu. Rev. Psychol.*, 59: 255-78.
- Evans, Jonathan St. B.T., Frankish, Keith (eds.), 2008, *In two minds: Dual processes and beyond*, Oxford University Press, Oxford.
- Fodor, Jerry. 1983, *The Modularity of Mind*, MIT Press, Cambridge (Mass.).
- Gigerenzer, Gerd, 2007, *Gut Feelings: The Intelligence of the Unconscious*, Viking Adult, New York.
- Hookway, Christopher, 2008, "Epistemic immediacy, doubt and anxiety: on the role of affective states in epistemic evaluation", in Brun, Doguoglu, Kuenzle (eds.): 51-66.
- Jacoby, Larry L., Dallas, Mark, 1981, "On the relationship between autobiographical memory and perceptual learning", in *Journal of Experimental Psychology: General*, 110: 306-340.
- Kahneman, Daniel, 2011, *Thinking. Fast and Slow*. Straus and Giroux, New York.
- Kahneman, Daniel, Frederick, Shane, 2002, "Representativeness revisited: attribute substitution in intuitive judgment" in Thomas Gilovich, Dale Griffin, Daniel Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge Univ. Press, Cambridge, UK: 49-81.
- Kahneman, Daniel, Frederick, Shane, 2005, "A Model of Heuristic Judgment", in Keith J. Holyoak, Robert G. Morrison (eds.), *The Cambridge Handbook of Thinking and Reasoning*, Cambridge Univ. Press, Cambridge, UK: 267-94.
- Koriat, Asher, 1995, "Dissociating knowing and the feeling of knowing: Further evidence for the accessibility model", in *Journal of Experimental Psychology: General*, 124: 311-33.
- Koriat, Asher, 2000, "The feeling of knowing. Some metatheoretical implications for consciousness and control", in *Consciousness and Cognition* 9: 149-171.

- Koriat, Asher, 2006, "Metacognition and consciousness", in Philip David Zelazo, Morris Moscovitch, Evan Thompson (eds.), *Cambridge Handbook of Consciousness*, Cambridge University Press, New York: 289-326.
- Koriat, Asher, Ravit, Levy-Sadot, 2000, "Conscious and Unconscious Metacognition: A Rejoinder", in *Consciousness and Cognition*, 9: 93-202.
- Koriat, Asher, Hilit, Ma'ayan, Ravit, Nussinson, 2006, "The Intricate Relationships Between Monitoring and Control in Metacognition: Lessons for the Cause-and-Effect Relation Between Subjective Experience and Behavior", in *Journal of Experimental Psychology*, 135, 1: 36-69.
- Mamassian, Pascal, Simon, Bathelmé, 2011, "Internal uncertainty, rather than expected performance, determines visual confidence", in *Journal of Vision*, 11, 11.
- Mercier, Hugo, Dan, Sperber, 2008, "Intuitive and reflective inferences", in Evans, Frankish (eds.): 149-170.
- Metcalf, Janet, 2000, "Feelings and Judgments of Knowing: Is There a Special Noetic State?", in *Consciousness and Cognition*, 9: 178-186.
- Metcalf, Janet, Arthur P., Shimamura, 1994, *Metacognition: knowing about knowing*. MIT Press, Cambridge, MA.
- Michaelian, Kourken, 2012, "Metacognition and Endorsement", in *Mind & Language*, 27, 3: 284-307.
- Millikan, Ruth, 2004, *Varieties of Meaning*. MIT Press, Cambridge, MA.
- Oppenheimer, Daniel M., 2008, "The secret life of fluency", in *Trends in Cognitive Sciences*, 12, 6: 237-241.
- Peacocke, Christopher, 1986, *Thoughts. An Essay on Content*. Blackwell, Oxford.
- Povinelli, Daniel, 2003, "Chimpanzee minds: suspiciously human?", in *Trends in Cognitive Science*, 7, 4: 157-60.
- Price, Mark C., Elisabeth, Norman, 2008, "Intuitive decisions on the fringes of consciousness: Are they conscious and does it matter?", in *Judgment and Decision Making*, 3, 1: 28-41.
- Pritchard, Duncan, 2007, "Anti-Luck Epistemology", in *Synthese*, 158: 277-98.
- Proust, Joëlle, 2006, "Rationality and metacognition in non-human animals", in Susan Hurley, Matthew Nudds (eds.), *Rational Animals*, Oxford University Press, Oxford: 247-74.
- Proust, Joëlle, 2007, "Metacognition and Metarepresentation", in *Synthese*, 159: 271-295.
- Reber, Rolf, Schwarz, Winkielman Piotr, Norbert, 2004, "Processing Fluency and Aesthetic Pleasure: Is Beauty in the Perceiver's Processing Experience?" in *Personality and Social Psychology Review*, 8, 4: 364-382.
- Sainsbury, Richard M., 1995, "Vagueness, Ignorance, and Margin for Error", in *British Journal for the Philosophy of Science*, 46: 589-601.
- Smith, David J., Wendy E., Shields, David A., Washburn, 2003, "The comparative psychology of uncertainty monitoring and metacognition", in *Behavioral and Brain Science*, 26: 317-373.

- Smith, David J., 2005, "Studies of Uncertainty Monitoring and Metacognition in Animals and Humans", in Herbert S. Terrace, Janet Metcalfe (eds.), *The Missing Link in Cognition. Origins of Self-Reflective Consciousness*, Oxford University Press, Oxford: 242-271.
- Sosa, Ernest, 2009, *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Volume II, Oxford University Press, Oxford.
- Stanovich, Keith E., 1999, *Who is rational? Studies of individual differences in reasoning*, Lawrence Erlbaum Associates, Mahwah, NJ.
- Stanovich, Keith E., Richard F., West, 2000, "Individual differences in reasoning: Implications for the rationality debate", in *Behavioral and Brain Sciences*, 23: 645-726.
- Thompson, Valerie, 2008, "Dual-process theories. A metacognitive perspective", in Evans, Frankish (eds.): 171-195.
- Unkelbach Christian, Greifeneder, Rainer (eds.), 2013, *The Experience of Thinking: How the Fluency of Mental Processes Influences Cognition and Behaviour*, Psychology Press, Hove.
- Weatherson, Brian, 2004, "Luminous Margins", in *Australasian Journal of Philosophy*, 83: 373-383.
- Williamson, Timothy, 1994, *Vagueness*, Routledge, London and New York.
- Williamson, Timothy, 2000, *Knowledge and Its Limits*, Oxford University Press, Oxford.
- Wurtz, Pascal, Rolf, Reber, Thomas D., Zimmermann, 2007, "The feeling of fluent perception: A single experience from multiple asynchronous sources", in *Consciousness and Cognition*, 17: 171-184.