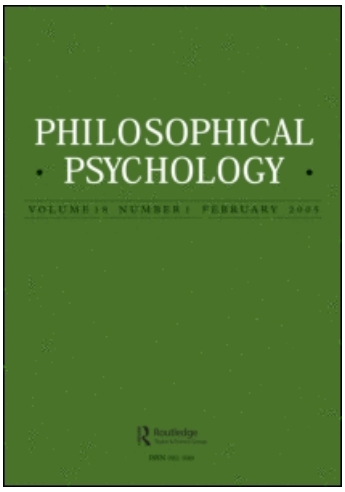

Philosophical Psychology
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713441835

Metacognition and consciousness: A review essay of Janet Metcalfe & Arthur P. Shimamura (Eds) metacognition: Knowing about knowing

Bernard W. Kobes [a]
[a] Department of Philosophy, Arizona State University, Tempe, Arizona, USA

Online Publication Date: 01 March 1997

To cite this Article Kobes, Bernard W.(1997)'Metacognition and consciousness: A review essay of Janet Metcalfe & Arthur P. Shimamura (Eds) metacognition: Knowing about knowing',Philosophical Psychology,10:1,93 — 102
To link to this Article: DOI: 10.1080/09515089708573208
URL: http://dx.doi.org/10.1080/09515089708573208

Metacognition and consciousness: a review essay of Janet Metcalfe & Arthur P. Shimamura (Eds) *Metacognition: knowing about knowing*

BERNARD W. KOBES

ABSTRACT *The field of metacognition, richly sampled in the book under review, is recognized as an important and growing branch of psychology. However, the field stands in need of a general theory that (1) provides a unified framework for understanding the variety of metacognitive processes, (2) articulates the relation between metacognition and consciousness, and (3) tells us something about the form of meta-level representations and their relations to object-level representations. It is argued that the higher-order thought theory of consciousness supplies us with the rudiments of a theory that meets these desiderata and integrates the principal findings reported in this collection.*

"I have never had an exaggerated interest in my own person, but ... that did not imply I could stop thinking about myself at will, from one moment to the next." So begins Cees Nooteboom's 1993 novella *The Following Story*. That we all in some sense think about ourselves from one moment to the next, and that this is central to human cognition, is a theme of much recent cognitive psychology. That it is the essence of consciousness is a theme of some recent philosophy.

J.T. Hart's (1965) doctoral dissertation reported that people have surprisingly good knowledge of which memory items they would be able to recognize, even if they cannot recall them. In Hart's recall-judgment-recognition (RJR) experimental paradigm, subjects were given common-knowledge questions. Where recall failed, they rated the likelihood that they would be able to recognize the item in a multiple choice task. Their judgments about this were largely accurate. Thus was born the field we now call *metacognition*—the study of our knowledge of our own mental states and processes. The book under review presents an excellent overview of the field as it stands today.

Thomas O. Nelson and Louis Narens make several criticisms of earlier research in cognitive psychology: it did not have a target outside the experimental situation; it overemphasized the non-reflective subject; and its experiments were designed to short-circuit the subject's control over his or her own cognitive processes. The

Bernard W. Kobes, Department of Philosophy, Arizona State University, Tempe, Arizona 85287-2004, USA. Email: kobes@asu.edu.

0951-5089/97/010093-10

cognitivist revolution made theorists more free than they had been to postulate causally efficacious inner states of the subject, but the picture was one in which inner representations were operated on by para-mechanical processes. It is a further stage of the cognitivist revolution to treat the subject as in some relevant sense *in charge* of his or her own mental life. The subject has knowledge of his or her own mental states, and that knowledge is put to use in inner action. A striking thing about the collection of essays under review is the extent to which it depicts the psychological subject as a cognitive *agent*, as one who plans and controls his or her inner mental life.

Nelson and Narens present a schematic model in which object-level cognitions are distinguished from meta-level cognitions. The meta-level monitors the object level; for example, it makes judgments of how easy an item will be to learn (ease-of-learning, or EOL), whether an item has been learned (judgment-of-learning, or JOL), whether an item not currently recalled is nevertheless known (feeling-of-knowing, or FOK), and how likely it is that a retrieved answer is correct. The meta-level also exerts *control* over the object level; for example, it selects strategies and allocates time for learning or for search.

On Nelson and Narens's model, the meta-level itself contains an "imperfect model" of the object level. This entails a point about what we might call the *ontology* of metacognition, namely, that object-level representations are wholly distinct from meta-level representations. Nelson and Narens do not say much in this article about the *form* of a particular meta-level representation that is wholly distinct from the relevant object-level representation. It would be implausibly redundant for the meta-level to reproduce the object-level representation completely, but it is difficult to imagine what any more abbreviated meta-level representation would be like (a gappy outline? a pale copy?).

But an ontologically distinct meta-level is by no means universally assumed. Of the writers in this collection, Asher Koriat is the most explicit about such matters. He presents evidence for an "accessibility account" of FOK, in which there is no distinct "trace" at the meta-level of the object-level representation. Instead, inferential mechanisms account for our FOKs. Even in what Koriat calls the "non-analytic" case, where FOK is a genuine *feeling*, rather than a conscious inference, the ease or difficulty of retrieving information related to the target accounts for that feeling. This theory apparently would not require that the metacognitive level be an ontologically distinct existence from object-level cognitive states. The theory has the advantage, therefore, of bypassing difficult questions about the form of a meta-level representation that is wholly distinct from the object-level representation.

Ann Miner and Lynne Reder argue that FOK is a rapid, automated, pre-retrieval state that triggers and guides *all* memory searches, not just those failed searches that result in an introspectively salient FOK. This provides an interesting contrast with Koriat's view that FOK is a function of the accessibility of related information. For Koriat, there is memory search from the word go, and search precedes and explains FOK. But for Miner and Reder this relation is inverted: FOK precedes and explains memory search. They would seem to agree, at least, in holding that the basic cognitive processes underlying *successful* search are the same

as those underlying *failed* search, but with the difference that for Miner and Reder FOK occurs in both cases, while for Koriat FOK, narrowly construed, occurs only with failed or not-yet-completed attempts, since in completed successful search the feeling of confidence in the retrieved answer pre-empts FOK.

Miner and Reder report a variety of manipulations in which simply making key words in the cue more familiar causes increased feeling of knowing without improving either recall or recognition of the target. Their theoretical account accordingly takes *cue familiarity* to be the key determinant of FOK. Koriat's accessibility account at least shares with this view the negative thesis that FOK does *not* consist in a partial access of the target or its trace. Tip-of-the-tongue (TOT) episodes, in which initial letter, number of syllables, and perhaps syllabic stress, are available though full recall remains maddeningly elusive, may be atypical in this respect, compared with the more general phenomenon of FOK.

That access to the target or to its trace plays little role in FOK is echoed by Louis Narens, Kimberly Jameson and V.A. Lee, who report that subthreshold target priming is largely ineffective in increasing FOK. There are tricky methodological issues here, for target priming will likely result in a higher rate of correct recall, thus removing those items from the base of potential FOKs. The article by Bennett Schwartz and Janet Metcalfe gives salutary warnings against a variety of pitfalls in designing and interpreting experiments in this field. In any case, Narens, Jameson and Lee report that only superthreshold *cue* priming has a robustly positive effect on FOK.

According to Stephen Stich, the pragmatic tasks of evaluating and criticizing our actual strategies of cognition, and the improvement of our cognitive skills, ought to be the primary goals of philosophical epistemology (see Stich, 1990). If so, there is much in the book to interest the epistemologist. Metacognition may be false or maladaptive, and a recognition of this could inform techniques and even technologies for cognitive improvement. Steven Smith, for example, gives evidence that many TOT episodes are illusory. Kenneth Weingardt, Jacob Leonesio and Elizabeth Loftus review evidence that eyewitness testimony is less reliable than we unreflectively take it to be. Robert Bjork presents some valuable and counterintuitive results on training. It turns out that many experimental manipulations that impede performance during training actually enhance long-term performance. Subjects who are satisfied with their training and confident of their mastery of the skill will perform less well in the post-training environment than subjects who feel frustrated and insecure about their mastery because difficulties have been introduced during training.

Christopher Hertzog and Roger Dixon review evidence that many memory problems associated with aging are metacognitive in nature. Many of the clinical cases described by Shimamura involve loss of metacognitive knowledge in those suffering from various kinds of brain damage. The work of Nelson and Narens suggests ways to improve study skills, and the work of Janet Davidson, Rebecca Deuser and Robert Sternberg suggests ways to improve problem-solving skills. So a variety of practical or therapeutic benefits are to be derived from a good understanding of metacognition.

Only the two articles by the editors discuss the neural bases of metacognition, and they are in my judgment the most interesting in the volume. Janet Metcalfe ascribes the feeling of knowing to a hypothesized system of novelty detection and habituation in which the frontal lobes play a key role. Arthur Shimamura offers an exceptionally lucid review of work in clinical neurology that bears on metacognition. He includes discussions of visual agnosia, blindsight, and memory disorders such as Korsakoff's amnesia. The idea is that these are all in some sense disorders of metacognition.

One impression the reader receives on reading these papers is that, despite the worthy efforts of Nelson and Narens, the field still lacks an adequately general theory of metacognition, or of its varieties, and of the relations between metacognition and consciousness. One reason why it would be valuable to see more theories of this sort is the threat of incipient fragmentation. It is often difficult to see the various kinds of work being done on metacognition as bearing on each other. In the absence of an adequate high-level theory, the impression is inevitably left that metacognitive phenomena are inherently disparate. Is there any single basic metacognitive structure or process that underlies the rest?

In Nelson and Narens's theory, the metacognitive level contains an "imperfect model" of the object level. But this is subject to the objection that it is difficult to see what kind of representation the metacognitive level could consist of, short of duplicating entirely the content of the object level. The field would profit from a general theory of metacognition that is more detailed than that on offer from Nelson and Narens.

An adequately general theory would go beyond that of Nelson and Narens in at least three ways. First, it would make clear how the kinds of phenomena that fall under the rubric of "metacognition" are related to each other. Second, it would have something specific to say about the relation between metacognition and consciousness (or kinds of consciousness). Third, it would say something about the nature of metacognitive representations, and their relation to object-level representations. For example, it would answer the question: are meta-level representations, i.e. representations of the self or of mental contents, wholly distinct from object-level representations, i.e. representations of the world?

Under the first heading would fall such questions as: is metacognition intrinsically a set of disparate phenomena? What, if anything, do all metacognitive states and processes have in common, in virtue of which they are metacognitive? Is there a central or focal kind of metacognition, of which others kinds are variations, and in terms of which those other kinds are to be understood?

It is not even clear, for many of the kinds of phenomena discussed in these papers, that they *are* genuinely metacognitive in nature. For example, Shimamura's review of apparently relevant literature in clinical neurology fails to show why a disorder such as associative visual agnosia, which involves impaired integration of atomistic visual perceptions into coherent perceptual wholes, should be considered metacognitive in nature.

Janet Davidson, Rebecca Deuser and Robert Sternberg discuss four component processes of problem solving that they call "metacognitive", viz., (1) identifying and

defining the problem, (2) mentally representing the problem, (3) planning how to succeed, and (4) evaluating what you know about your performance. But they fail to discuss the apparent possibility that some of these processes might go on without any psychologically real cognition *about* cognition. As Christopher Hertzog and Roger Dixon note (p. 230), the ongoing *use* of a strategy does not necessarily amount to *knowledge about* strategies and their potential effectiveness. Might not a pattern of adaptive behavior, even the "inner behavior" of selecting a problem-solving strategy, be uninformed by psychologically real metacognitive processing? Parsimony of levels has its attractions, and the field of metacognition needs an explicit response to such challenges.

Several authors take confidence judgments to be unproblematically metacognitive in nature. Miner and Reder (pp. 29–30) define the feeling of confidence as "the state of believing that a particular piece of information has been correctly retrieved from memory." But this presents a puzzle: is not a confidence judgment at bottom just another judgment *about the world*? If some cognitive process, e.g. memory search in eyewitness testimony, results in a report that p, and then the witness makes a retrospective confidence judgment about this memory report, is she not simply estimating, perhaps more carefully than before, how likely it really is that p? If so, then confidence judgments are basically object-level phenomena. A confidence judgment is at best superficially metacognitive if it arises through some inferential process such as: I have judged that p, but it is in fact only probable to degree .7 that p, therefore I have degree of confidence .7 in my judgment. The field of metacognition needs to deal directly and explicitly with such challenges, both because it needs to defuse skepticism about the robust psychological reality of the meta-level, and because our understanding of metacognition will be deepened by answering such challenges.

The second desideratum for a general, high-level theory of metacognition was that it should include a theory of the relation between metacognition and consciousness—or the varieties of each. Unification should be a goal of psychology, as it is of fundamental physics. Endel Tulving, in his foreword to this book, calls for a greater openness to the explicit investigation of consciousness, and expresses the widespread hunch that metacognition is intimately related to consciousness, since consciousness is a necessary condition of metacognitive experiences. Of course, proclaiming a relationship is one thing; articulating it is another. Just as there are many forms of metacognition, so there may be many forms of consciousness. We need a conceptual framework that makes these relations explicit.

A key strategic question for such a framework is whether it will posit essentially *one* notion of consciousness and identify some other parameter as the conceptual cause of the variety of types of metacognition, or whether on the other hand the various phenomena of consciousness and the various phenomena of metacognition map onto each other, so that for every kind of metacognition there corresponds a unique notion of consciousness.

As an example of what such a theory might do for us, consider the idea that consciousness consists in a species of *inner awareness* whose object is a first-order mental state. One way of working out this idea is the *higher-order thought* theory of

consciousness—the HOT theory, for short, if we may be permitted yet another three-letter acronym. The idea is that for a mental state or event to be conscious it must be accompanied by a suitable higher-order thought to the effect that one is in that very mental state. The higher-order thought, in virtue of which the *first-order* mental state or event is conscious, is itself typically unconscious [1].

Endel Tulving's foreword takes the connection between metacognition and consciousness to lie in the fact that metacognitive monitoring, judgment and experience are all conscious states and processes. But the HOT theory more sharply delineates this connection, and even inverts it in a crucial respect. According to the HOT theory the interesting connection lies not in the fact that *meta-level* states and processes are conscious, as Tulving's formulation would have it, but rather in the fact that an *object-level* state's being conscious consists in its being the target of a suitable metacognitive state. Consciousness of object-level states, not consciousness of meta-level states, is a necessary condition of metacognition. Though it may be true, as Tulving says (p. ix), that "conscious awareness is one of the defining attributes of the domain of metacognition", there is more theoretical power in the converse, that metacognition is a defining attribute of the domain of consciousness.

According to the HOT theory, higher-order thoughts are not themselves conscious, in the ordinary case. However, they are psychologically real, and play a causal role in enabling adaptive patterns of cognition. Higher-order thoughts may *become* conscious through introspection, and thus we may come to have conscious metacognitive experiences, e.g. a conscious feeling of knowing. But conscious metacognition is derivative, and draws on pre-existing real metacognitive mental structures. Introspection is not the typical case; metacognition ordinarily occurs, and contributes to effective cognition, even when metacognitive states are not themselves conscious. The surprisingly many limitations and errors of metacognition are to be understood against a background or default assumption of meta-level cognition of the contents of current, conscious object-level mental states and processes—that is precisely why the limitations and errors of metacognition are *surprising*. This, then, would be the outline of a theoretical framework for defending the robust psychological reality of adaptive metacognition against levels-parsimony, even in cases of subjects who do not report any conscious metacognitive states and processes.

The third desideratum for a general theory of metacognition is that it should tell us something about what metacognitive representations are like. What is their form, and what is their relation to first-order representations? Are metacognitive representations distinct from the first-order representations that they are about? Here too the HOT theory provides a framework for fruitful articulation of competing accounts.

If the higher-order thought were *wholly distinct* from the first-order state, then a metacognitive state such as FOK could consist in accessing a distinct trace of the target, and this could be identified with the re-activation of the relevant higher-order thought—the higher-order thought in virtue of which the target was once conscious. However, on either Koriat's account of FOK as a function of accessibility of target-related information, or on Miner and Reder's cue familiarity account, FOK does *not* consist in the access of a wholly distinct trace of the target. If activation of

a higher-order thought wholly distinct from the target corresponded to the FOK, then we would not expect FOK to be as sensitive to cue familiarity as it in fact is. Moreover, we would face embarrassing questions about the form of the meta-level representation.

An attractive alternative would be to say that the higher-order thought includes the target, the first-order state itself, as a *constituent*. I do not mean merely that the *content* of the first-order state, construed as a type that has instances, is a constituent of the content of the higher-order thought, but rather that the *particular occurrence* of the first-order state is a constituent of the relevant particular occurrence of the higher-order thought (see Kobes, 1995).

Besides having an advantage in economy of theoretical apparatus, making the first-order state a constituent of the higher-order thought allows us to identify the feeling of *familiarity* as the feeling characteristic of the re-activation of an old higher-order thought. This is consistent with the plausible idea that only items of which one was once conscious can seem familiar, since only if one was once conscious of the item could there be a stored higher-order thought. (Of course this is not to say that every item a subject was once conscious of can, in principle, seem familiar, or be recognized or recalled. A higher-order thought of consciousness is not *necessarily* preserved.)

Familiarity of a retrieved item ought to be as interesting to psychologists as FOK, for it is equally metacognitive, having roughly the content "I've thought/experienced this before." It ought to be an object of investigation rather than a theoretical primitive. We should ask not what familiarity can do for metacognition, but what metacognition can do for familiarity. Janet Metcalfe is the only writer in this volume to recognize this explicitly. (See especially p. 139 for a parallel point about habituation.) By identifying familiarity with the re-activation of an old higher-order thought, the HOT theory may contribute to empirically testable articulation of the phenomenon.

If the higher-order thought contains the first-order state as a constituent, how can we understand FOK? FOK cannot consist in full activation of the higher-order thought, as it did on the view that the higher-order thought is wholly distinct from the target, for it would be highly implausible to suppose that FOK always consists in access, or partial access, of the target itself. I do not think, however, that this is a difficulty for the "constituent" version of the HOT theory. For FOK without accessing the target itself can be understood as the *partial* re-activation of an old higher-order thought, or as the re-activation of an old higher-order thought with a *gap* in the target position.

In fact, given the importance of cue familiarity as a determinant of FOK, and the understanding of familiarity as the re-activation of an old higher-order thought, we may derive the result that FOK can be understood as a phenomenon of spreading activation. The old higher-order thoughts associated with the cue, together with the old higher-order thought for the target, constitute a *region* of the meta-level. When the old higher-order thoughts for the cue are re-activated, i.e. when the cue is *familiar*, the higher-order thought for the target is activated too, by spreading activation. In many cases the target will be thereby retrieved,

and the metacognitive feeling of confidence in the retrieved answer can be identified with the re-activation of an old higher-order thought containing the target. But if the target itself is missing or not currently retrievable, the subject will be left with only a *feeling* of knowing, without the target itself as a component. Spreading activation would result in the partial re-activation of an old higher-order thought, or in the re-activation of an old higher-order thought with a gap in the target position.

This bare-bones account may be elaborated in various ways. Re-activation of higher-order thoughts for quickly accessed target-related information may contribute to the re-activation of the higher-order thought for the target; thus the evidence for Koriat's account may be assimilated into the current theory. Wholly illusory FOK may be understood as a *filling-in* phenomenon in the meta-level region, on analogy with the retinal blindspot.

A striking recurrent theme of this volume is the extent to which metacognition is *active*. Allocation of study time, selection of study strategy, selection of memory search strategy, selection of problem-solving strategy, and deciding when to cease study or memory search—even the allocation of attention—are construed as metacognitive control processes. Cognition is not something that a subject *undergoes*; he or she is typically *in charge*.

Shimamura (pp. 269–270) provides a fascinating elaboration on this theme in comparing the phenomenon of blindsight with that of implicit memory in certain kinds of organic amnesia such as Korsakoff's syndrome. Blindsight is perception without conscious awareness, while implicit memory in amnesics, as exhibited in priming effects for example, is memory without conscious awareness. Blindsight is an impairment in perceptual *activity*. (Is associative visual agnosia construed as metacognitive for the same reason—that it is a disorder of perceptual control processes?) In organic amnesia, memories may "pop into the mind", but cannot be *sought* (pp. 266–267). Both cases seem to involve an impairment in the evaluation or integration of cortical activity, and both leave their victims cognitively *passive*.

Korsakoff's patients, as well as patients with frontal lobe pathology, exhibit impaired FOK, and this is manifested in impaired metacognitive *judgments* or *decisions* (pp. 274–275). And Metcalfe suggests that impaired ability to detect novelty and habituation, such as characterizes frontal lobe pathology, may manifest itself as an attentional deficit, as apathy and distractibility, as failure of drive (pp. 153–154).

That the meta-level exerts control over and can effect changes in the object level is something that the Nelson and Narens model rightly stresses. Even so, it would be salutary to articulate a stronger sense of how the meta-level can be active than their model supplies. For if the meta-level is ontologically distinct from the object level, as it is in the Nelson and Narens model, we face embarrassing questions about *how* the meta-level acts upon the object level. Out of what materials does it create object-level thoughts, and what kind of *grip* does the meta-level have on the object-level material? These seem bad and fruitless questions, perhaps because the relation between meta- and object levels is in fact more direct than it is in the Nelson and Narens model.

By embedding object-level states and processes in meta-level acts, as in the "constituent" version of the HOT theory, the embarrassing questions dissolve. An object-level state can be literally a *part* of a meta-level cognitive act. Metacognitive control processes change the contents of one's conscious awareness, not by changing sensory inputs, and not even by manipulating wholly distinct object-level materials, but *directly*. Consciousness—an arena of first-order states accompanied by higher-order thoughts—is in fact an arena of affordances for inner action.

Assertoric speech acts are sometimes said to have *word-to-world* direction of fit, while imperatives have *world-to-word* direction of fit. (It may help to think of these as short for "word-to-fit-world" and "world-to-fit-word". Direction of fit is typically the opposite of the direction of causal flow.) That is to say, we conceive of assertions as successful or not by holding fixed the world, and holding the words, or processes that produce them, as in some sense *responsible* for effecting a fit; but we conceive of imperatives as successful or not by holding fixed the words, and holding the world, or processes that act on it, as responsible for effecting a fit.

A similar distinction may be made among higher-order thoughts. Some have assertoric force; they are successful or not according to whether they accurately reflect object-level states, where the object level is conceived as held fixed. They have the direction, as it were, *meta-to-(fit)-object*. On the other hand, some higher-order thoughts have a force analogous to imperatives, or even acts as such. They have the direction, as it were, *object-to-(fit)-meta*. They not only *are about* the first-order states they contain, they *bring about* the first-order states they contain. A subject's knowledge of her object-level processes may therefore relevantly resemble her knowledge of her current, conscious *acts*. So the "constituent" version of the HOT theory has the resources to articulate a very basic sense in which metacognition can be active. A variety of metacognitive control processes may be brought together into a unified and coherent account by postulating that a higher-order thought may have the force of a mental act.

The main goal of these reflections has been to provide the rudiments of an attractive example of the *kind* of theory that we need, namely, a unified theory of consciousness and metacognition. Though I think the HOT theory promising, other approaches to a unified theory need to be explored—e.g. approaches that see *broadcasting* as the chief explanatory primitive for a theory of consciousness. Competing theories may then be compared for adequacy, though the possibility should be kept in mind that apparently competing theories may in fact differ only in superficial vocabulary, while being equivalent in empirical import.

A theory of metacognition and consciousness will be judged ultimately not only by empirical tests, though of course its susceptibility to such tests and its success in meeting them is of the highest importance. The theory will also be judged by its power to bring about explanatory unification, that is, its power to bring together into a coherent whole, with a single small set of explanatory primitives, a variety of phenomena whose inter-relations we now see dimly at best. This form of thinking about ourselves will require, I am convinced, both philosophical acumen and a knowledge of the psychological work so richly sampled in Metcalfe and Shimamura's collection.

Note

[1] Armstrong (1981) defends a view of conciousness as perception of one's own mental states. Rosenthal (1997) presents an account of conciousness in terms of non-perceptional higher-order thoughts. See also the select bibliography on pp. 522–523 of Metzinger (1995).

References

ARMSTRONG, D.M. (1981). What is consciousness. In D.M. ARMSTRONG, *The Nature of Mind and Other Essays* (pp. 55–67). Ithaca, NY: Cornell University Press.

HART, J.T. (1965) *Recall, recognition, and the memory monitoring project.* Doctoral dissertation, Stanford University (University microfilms no. 66–2565).

KOBES, B.W. (1995). Telic higher-order thoughts and Moore's paradox. In TOMBERLIN (Ed.) *Philosophical perspectives 9: AI, connectionism, and philosophical psychology* (pp. 291–312). Atascadero, CA: Ridgeview Publishing Company.

METCALF, J.E. & SHIMAMURA, A.P. (Eds) (1994). *Metacognition: knowing about knowing.* Cambridge, MA: MIT Press.

METZINGER, T. (Ed.) (1995). *Conscious experience.* Lawrence, KS: Imprint Academic.

ROSENTHAL, D.M. (1997). A theory of consciousness. In N. BLOCK, O. FLANAGAN & G. GUZELDERE, (Eds) *The nature of consciousness.* Cambridge, MA: MIT Press.

STICH, S. (1990). *The fragmentation of reason.* Cambridge, MA: MIT Press.