

# The Contents of Phenomenal Consciousness

## One Relation to Rule Them All and in the Unity Bind Them

Antti Revonsuo  
Department of Philosophy and Center for Cognitive Neuroscience  
University of Turku  
Turku FIN-20014  
FINLAND

revonsuo@utu.fi

Copyright (c) Antti Revonsuo 2003

PSYCHE, 9(08), February 2003  
<http://psyche.cs.monash.edu.au/v9/psyche-9-08-revonsuo.html>

**KEYWORDS:** phenomenal consciousness, unity of consciousness, binding, explanation, explanatory gap, phenomenal space.

**COMMENTARY ON:** Dainton, B. (2000). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.

**ABSTRACT:** *Stream of Consciousness* is a detailed and insightful analysis of the nature of phenomenal consciousness, especially its unity at a time and continuity over stretches of time. I find Dainton's approach to phenomenal consciousness in many ways sound but I also point out one major source of disagreement between us. Dainton believes that to explain phenomenal unity and continuity, no reference to anything outside experience is required. Thus, he postulates a fundamental experiential relation called co-consciousness which is supposed to do all the explanatory work. On the contrary, I hold that to truly explain features of consciousness such as phenomenal unity and continuity, reference to mechanisms outside the phenomenal realm are necessary.

## 1. Introduction

This is one of the best and most original books on consciousness that I have read for a long time. It is written in a beautifully clear style, the argumentation is entirely lucid and proceeds in a systematic manner. Best of all, and unlike most recent consciousness books

on the market, *Stream of Consciousness* mostly deals with the nature and internal structure of phenomenal consciousness itself, rather than with brain, computation, quantum physics, zombies, blindsight, representation, intentionality, and other topics whose connection to phenomenal consciousness is indirect at best and nonexistent at worst. And not only does it *deal with* the issue, but manages to make a significant contribution to the literature on consciousness.

The book is full of insights, arguments and important conceptual distinctions. It would take another book to cover all of it within the scope of a single review. Therefore, I will focus especially on the following questions:

What is the place and the significance of a purely phenomenological approach, such as Dainton's, in the science of consciousness?

When it comes to the nature and the explanation of consciousness, what kind of fundamental assumptions does Dainton commit himself to, and are the commitments plausible and fruitful for the empirical science of consciousness?

## **2. The Role of Phenomenology in the Science of Consciousness**

In the Preface, Dainton notes that although the recent interest in consciousness has led to a burgeoning literature, relatively little of it is concerned with phenomenological issues. Instead, most make an attempt to understand the *relationship* between consciousness and the brain (or the physical world). He points out that to understand the relationship between those two entities, we first need an understanding of each of them. Therefore it is premature to try to solve the problem of the explanatory gap before we understand what lies at *both* sides of the gap.

I can only agree with these observations. I would express the same concern in the following way: once we are committed to realism concerning the phenomenal realm, we need a science of consciousness that takes this commitment seriously (Revonsuo 2000, forthcoming). In traditional branches of science (e.g. biology), taking some domain of phenomena seriously means that what is first required is a detailed and systematic description of the phenomena in their own terms or at their own level of description. All science must start by the identification and systematic description of the *explanandum*. Only after we have an understanding of the nature of the *explanandum* can we start to construct hypotheses regarding the possible underlying *mechanisms* that might be invoked, in order to explain the phenomenon. To invoke explanatory mechanisms before the nature of the *explanandum* is clear will not lead to successful theoretical explanation, for in that case it is difficult to understand what the mechanisms are supposed to explain

and whether they bear any intelligible explanatory relationships to the actual *explanandum*.

Therefore, at this early stage of its development, the science of consciousness should direct considerable resources to the systematic study of phenomenological issues, in order to first construct a detailed map of the phenomenal level of description. The explanatory gap, or the question whether the phenomenal level can be intelligibly related to any physical or neurobiological levels in the brain, can be dealt with only after such a description is available. In any case, there is every reason to believe that an understanding of the inner structure and dynamics of the phenomenal level might render the gap less unbridgeable, for the features of the phenomenal level (how it is structured, how it dynamically changes across time, and so on) offer top-down constraints for the science of consciousness in the search for potential explanatory mechanisms in the brain.

It seems that Dainton realizes the role and value of the top-down constraints offered by phenomenological description; on p. 23 he writes that situations in which the phenomenological data conflict with accepted science may well provide valuable clues as to how the relevant science might be revised and improved. I agree with this: once we have a detailed description of the phenomenal level, it will suggest what sorts of lower-level neural phenomena might be closely associated with the higher phenomenal level, even if such phenomena would not yet have been discovered by neuroscience (and even if they were for the time being not discoverable due to current methodological limitations in neuroscience, see Revonsuo 2001). At least it gives us a clue what kind of entities neuroscience should be looking for if it wishes to explain consciousness.

However, I would like to add that even in the multidisciplinary science of consciousness, phenomenal descriptions should not be given any *absolute* authority over other sources of data. The advisable strategy would rather be along the lines of Owen Flanagan's (1992) "natural method": to treat phenomenology, empirical psychology and neuroscience all with equal respect, and in cases of conflict, any or all stories may have to be revised. It is conceivable that in some cases (e.g. confabulation, false memories) phenomenological descriptions are inaccurate or unreliable and therefore neuroscience does not need to accommodate them. However, such a situation is likely to be an exception rather than the rule.

A phenomenological approach such as presented in *The Stream of Consciousness* is not only useful, I believe it is absolutely necessary if the science of consciousness is ever going to make any progress at all. Without the phenomenological approach, the science of consciousness may become a fatally degenerating branch of science that is incapable of characterizing its core explanandum in any clear and systematic manner.

### **3. Dainton's Fundamental Commitments**

Dainton's starting point is, not surprisingly, "robust, full-blooded realism about consciousness". This commitment alone rules out eliminative and reductive physicalism. However, he is also committed to the reality of the physical world, which rules out phenomenalism and idealism, and he furthermore rejects substance dualism. Other than that, he remains neutral as to the precise nature of the matter-consciousness relationship. He calls his approach "moderate naturalism". Still, he clearly expresses sympathy for what he calls "phenomenalized materialism" or the view that the intrinsic character of at least some physical entities is phenomenal, but he remains uncommitted to this view.

On the issue of the nature of perception, Dainton commits himself to indirect realism according to which we are never directly in touch with external physical objects but our perception of them is mediated by phenomenal entities directly present for us. Perceptual experience comes in the form of "world-presenting"; perceptual contents are experienced as being "out there", but in fact the contents are projected mental entities, not the external stimulus objects themselves somehow directly perceived. He summarizes the implications of this view of perception (p. 18): "there is a sense in which we are all enclosed in spheres of virtual reality, phenomenal spheres somehow produced by activities within our own brains: all we are directly aware of are the contents within these spheres." He adds that this is not how it seems to us; we seem to be surrounded by ordinary material objects rather than by experience. However, critical phenomenology inescapably leads us to this counter-intuitive conclusion.

My role as a critic becomes increasingly difficult here, for I cannot detect any significant disagreement between Dainton's view and the Virtual Reality metaphor of consciousness that I have proposed as the starting point for the science of consciousness (Revonsuo 1995, 2000, forthcoming). Dainton is surely correct when he says that the view is in conflict with our naive intuitions. Our deeply rooted intuitions are probably the main reason why indirect or representative realism is so difficult to accept. In my own work I have tried to show that if we take the dreaming brain as a model system for consciousness science, then it becomes much easier to see that the internal workings of the brain must be fully sufficient for producing an entire sphere of externalized or projected experience. Dreaming is produced by the internal activation of the brain, yet experienced as "being-in-the-world" or as "world-presenting", not as a mental image inside our heads. What Dainton calls projection I have called the "Out-of-the-Brain" -illusion: even during dreaming, we do not conceptualize our experience as the experience of being trapped inside the brain or of seeing how the brain looks like from the inside. On the contrary, we experience direct presence in and full immersion into a phenomenal (dream) reality, which must however reside entirely within the confines of the brain. According to this view, veridical perception differs from dream experience only in one sense: whereas dream experience creates a virtual presence in an imaginary world, perception creates a kind of telepresence experience: the feeling of direct presence in the center of the external world physically surrounding the organism; yet the experience-generating mechanisms (and the experience too) remain literally within the confines of the brain.

## 4. The Basic Structure of Consciousness: Simple or Relational?

A further fundamental commitment that Dainton makes in the first chapters of the book is the rejection of a relational or bi-polar structure of consciousness. He calls the doctrine that consciousness has a relational "awareness-content" or "act-object" -structure by the name of "the A-thesis". He argues, to my mind quite convincingly, that there are no grounds to believe that the A-thesis is correct. It is difficult to see what sort of an entity the contentless awareness postulated by the A-thesis is supposed to be and why it would be needed. The wide acceptance of the A-thesis is probably due to the mistaken idea that experience should be modeled after perception, as some sort of internal perception. Dainton does an excellent job in showing the A-thesis untenable.

One currently popular version of the A-thesis seems to be widely accepted within the philosophy of mind. Those who advocate representationalism or some variety of an intentional analysis of consciousness typically defend the view that consciousness is a representational relation between the *vehicle* and the *content* (or object) of consciousness. This variety of bipolar analysis leads to the view that the contents of consciousness are not in the brain and that phenomenology "ain't in the head" (as Tye once put it). Dainton does not discuss representationalism, but it should be pointed out that also the representationalist doctrine is in conflict with the Simple Conception of experience adopted by Dainton. This is because representationalism implies a *relational* analysis of consciousness, if not in terms of "act-object", then in terms of "vehicle-content". In contrast to this, the Simple Conception is an entirely *non-relational* view of consciousness. Thus, it would seem to be difficult to accept both a representationalist analysis and the Simple Conception of consciousness at the same time, for any such view would have to accept two fundamentally conflicting views the relational and the non-relational of the basic underlying structure of consciousness.

According to the Simple Conception, consciousness is inseparable from phenomenal contents; contents are intrinsically conscious, or self-revealing. Dainton seems to be unaware that a similar view has very recently been defended by Stubenberg (1998). When phenomenal consciousness itself (instead of intentionality, representation, or organism- environment interaction) is considered of primary significance in a theory of consciousness, various authors, independently of each other, tend to end up with a simple, non-relational conception of experience. I have to confess that such a non-relational theory of experience seems to me the most promising avenue for the future science of consciousness. Therefore it should be adopted as the working hypothesis in the theoretical core of the science of consciousness (Revonsuo, forthcoming). Indeed, perhaps the time is ripe for the Simple Conception to seriously challenge or even overthrow functionalism and representationalism as the dominating philosophy of consciousness.

## 5. Phenomenal and Non-Phenomenal Space

Chapter 3 is devoted to phenomenal space. Dainton acknowledges that to a very large extent our consciousness is spatial, and that an understanding of the unity of consciousness requires an understanding of phenomenal space. The traditional view that consciousness is wholly non-spatial is based on questionable assumptions originating either in Cartesian dualism or in the act-object separation between consciousness (which is supposed to be non-spatial) and its (external perceptual) objects or contents which are spatial. Our phenomenology has an undeniable spatial character and no theory of consciousness committed to the primacy of phenomenology can brush this fact under the carpet.

But simply admitting the spatial character of the phenomenal level is only the first step. It must be followed by an analysis of the implications of such a position. Dainton provides us with a sensible distinction between two possible types of phenomenal space: P-space and V-space. A P-space is a substantial phenomenal space; space which in itself possesses some phenomenal reality, even in the absence of any other phenomenal properties or objects. A P-space is necessarily a phenomenally experienced expanse, even if it is totally empty. A V-space is not like that; by contrast, it is in itself a phenomenal vacuum, a phenomenal nothingness. However, other phenomenal properties or objects are experienced as spatially related within this in-itself non-phenomenal space. A V-space does not consist of any phenomenal substance. As Dainton points out, and I have no objections, our phenomenal space seems to be a V-space rather than a P-space.

Still, I am not fully convinced by Dainton's arguments and thought experiments to the effect that phenomeno-spatial unity would not be a necessary condition (and perhaps also the underlying explanation) for co-consciousness. Although he admits that spatiality is a very important factor in phenomenal unity, he suggests that there is reason to believe that experiences, although altogether dissociated phenomeno-spatially, could somehow be co-conscious. He describes a science-fiction scene in which a subject's audio-visual experience comes from sensory organs located at a mountain top and at the same time bodily experience (including voluntary movement) is fed into the subject's consciousness from a body diving and swimming at the bottom of the ocean. This bizarre arrangement would lead, according to Dainton, to a situation in which the two channels of experience would be co-conscious (i.e. being experienced together as synchronous component parts of a single stream of consciousness) but still entirely unconnected in any spatial manner. The subject of course would pay more attention either to one or the other channel of information at any one time but both streams would be conscious at the same time.

This argument is weakened by the well-known fact that our perceptual system (and our consciousness) abhors incoherent sources of input. For example, if two different objects or images are presented to the two eyes, either they are merged into one, if coherent enough (binocular fusion), or only one of them will be seen at any one time (binocular rivalry). Thus, a more plausible prediction as to Dainton's thought experiment would be that: 1) Either the two coherent streams of information compete for access to awareness and therefore only one of them is experienced at any one time or 2) They will be merged

together into a very unusual phenomenal world: You seem to be on a mountain-top as regards your audiovisual experience, but somehow you also seem to have an invisible body located there; furthermore the whole space around you seems to be filled with a cold, invisible liquid surrounding your body. There are also invisible objects and surfaces all around you. If you move your body, you feel that you can move relative to the fully transparent liquid and its invisible objects, but curiously enough you stay still in relation to the mountain-top. The entire invisible liquid world seems to be moving in relation to the visible world, rather than you moving inside the liquid world. Thus, in this case of the streams of information merging into one there would be experience of only one phenomenal space, albeit a very peculiar one: an audiovisual phenomenal space filled with transparent liquid moving in relation to your invisible body and the surrounding invisible objects.

There seem to be cases of altered states of consciousness in which people experience themselves in two places at once (e.g. out-of-the-body -experiences). Are they good evidence for fundamental spatial disunity in consciousness? The cases I am familiar with seem rather more like a dissociation between visual perspective and bodily awareness. The experience is usually described as one where "I could see myself from the outside", which indicates that the visual field is constructed from a disembodied perspective, but a human character identified as "me" or as "my body" appears as one distinct object amongst others in the visual field. Sometimes the two spatial frames of reference (the normal embodied and the abnormal disembodied perspectives) can alternate or one can rapidly turn into the other. Still, at any one time the world is seen either from the embodied or the disembodied perspective, not two perspectives at once. Thus, these cases do not seriously challenge the fundamental spatial unity of consciousness.

Of course one can always argue that spatially dissociated experiences might still be possible in some non-human consciousness. Perhaps, and perhaps not. I would be inclined to say that in such cases the alien creature possesses more than one single center of consciousness or that it houses more than one subject. Thus, I am inclined to treat phenomeno-spatiality as the basic unifying feature of human consciousness, and also as an explanation of why the relation of co-consciousness holds. Furthermore, I believe that the postulated phenomenal spatiality must in turn be explained by referring to even lower, non-phenomenal levels of organization. By contrast, Dainton takes co-consciousness to be the ultimate explanation of binding and the unity of consciousness, and says that there is no need to postulate any unifying principle over and above experience and inter-experiential relations.

## **6. Dainton's Explanatory Strategy Does Not Deliver Genuine Explanations**

There is one conspicuous difference between my view and Dainton's when it comes to the explanation of consciousness. I suppose we both agree that the systematic description

of phenomenal consciousness is a necessary first step towards an explanation of consciousness. But what should we do when we proceed from mere *description* to actual *explanation*? This is the point where we take different courses.

When seeking for an explanation of the unity of phenomenal consciousness, Dainton prefers to stay entirely within the phenomenal level and postulate explanations that refer to nothing outside the experiential realm: "[E]xperience is self-unifying, in that to understand the unity we find within experience, we do not need to look to anything above, beyond or external to experience itself" (p. 236).

I fear that this strategy of explanation renders the relation of co-consciousness inexplicable, after all. It is difficult to see how the phenomenal features of experiences (such as unity) can be made intelligible solely by referring to further phenomenal features of experiences, or by labelling the phenomenal features themselves with technical terms such as "co-consciousness". In order to truly *explain* the features of consciousness, one must leave the phenomenal level and refer to something outside of it.

That is exactly what the strategy of explanation that I favour does. I have proposed that consciousness should be treated as one specific biological level of organization in the brain, and to explain the phenomenal level its internal structure, unity, dynamics, and organization reference to the underlying non-phenomenal levels is appropriate, indeed necessary (Revonsuo 2000, 2001, forthcoming). This explanatory strategy is no different from the one used in the biological sciences to explain other complex biological phenomena at high levels of biological organization: by seeking the lower-level features that could actually constitute and support the features of the higher level phenomenon. Conversely, features of the higher-level phenomenon act as top-down constraints to what count as the possible underlying explanatory mechanisms at the lower level. This explanatory strategy can be called multi-level or cross-level mechanistic explanation.

The core idea of this strategy is that intelligible constitutive relationships must be identified between the higher and the lower levels of organization. To explain unity at a higher level (such as the phenomenal unity of consciousness) some lower-level non-phenomenal mechanism must be identified that makes the higher-level unity intelligible. Unless intelligible inter-level relationships are found, the relations between levels remain merely correlative and thus do not manage to explain the higher-level phenomena (e.g. the "neural correlates of consciousness" are not really explanations of consciousness).

## **7. When Unity of Consciousness Falls Apart**

Dainton explicitly states in the Preface that he has concentrated on getting a clear picture of *typical* streams of consciousness and of the most *general* features of experience. He has thus deliberately neglected altered states of consciousness (e.g. meditation, dreaming, drug states) and pathologically deteriorated consciousness (the data arising from neuropsychology). This restriction is understandable, for it would most likely have been



impossible to include all these diverse and plentiful sources of data within the scope of a single book. However, from the perspective of the science of consciousness, altered states and deteriorated or distorted consciousness constitute at least as important sources of data as our everyday normal consciousness does, if not more so. We should remember that in several other branches of science the anomalous cases have proved to be the theoretically most important ones. For example, when charting the cognitive architecture of mind, cognitive neuropsychologists concentrate on the patterns of *deficient* performance that emerge after selective brain injury. Those unusual patterns reveal dissociations between functions that normally go tightly together, thus revealing the organization of the underlying mechanisms, and the limits of possible patterns of performance. The theory of the normal cognitive architecture must be capable of accommodating and even predicting all the in-principle possible deficient patterns of performance.

In the case of consciousness, it seems to me that the unusual patterns of phenomenal experience occurring during dreaming, drug states, and after brain injury surely contain crucially important information as to the possible structure and the underlying mechanisms of consciousness. Anomalous distortions of consciousness (such as found in unilateral neglect, simultanagnosia, dream bizarreness etc.) in many ways violate our expectations of what phenomenal experience is like and in what kind of forms it can be manifested. Those anomalous cases have to be accommodated by any comprehensive account of the phenomenal realm, for they are actual even if not very typical forms of subjective experience. Furthermore, they may reveal the underlying mechanisms that are at work beneath the surface of the normal unity of consciousness, for in many of the anomalous cases it is some specific aspect of the normal phenomenal unity that is missing or malfunctioning (Revonsuo 1999).

By contrast, the anomalous neuropsychological cases and unusual phenomenological dissociations do not have all that much significance in Dainton's phenomenological approach to explanation. Dainton seems to believe that the phenomenal unity of consciousness does not *have* any underlying explanatory mechanisms; the unity of consciousness should not be explained by referring to underlying non-phenomenal mechanisms. In Chapter 4 Dainton does mention in passing some of the relevant neuropsychological syndromes, but hardly gives them the kind of treatment they would deserve.

It would be interesting to find out how Dainton's purely phenomenological approach treats something like the following cases as possible sources of data about the binding and unity of consciousness.

Epileptic absence seizures appear to involve a total cessation of phenomenal flow for a few seconds at a time, but do not necessarily have any influence on phenomenal continuity (Damasio 1999). Absence seizures would thus be relevant to the discussion of phenomenal time in Chapter 5.

Cerebral akinetopsia is a rare neuropsychological syndrome that involves an inability to see phenomenal continuity in visual objects:

The visual disorder complained of by the patient was a loss of movement vision in all three dimensions. She had difficulty, for example, in pouring tea or coffee into a cup because the fluid appeared to be frozen, like a glacier. In addition, she could not stop pouring at the right time since she was unable to perceive the movement in the cup (or a pot) when the fluid rose. ...In a room where more than two people were walking she felt very insecure and unwell, and usually left the room immediately, because 'people were suddenly here or there but I have not seen them moving'. .. She could not cross the street because of her inability to judge the speed of a car, but she could identify the car itself without difficulty. 'When I'm looking at the car first, it seems far away. But then, when I want to cross the road, suddenly the car is very near.' She gradually learned to 'estimate' the distance of moving vehicles by means of the sound becoming louder (Zihl, Von Cramon & Mai 1983, p. 315).

For this patient, there were only stationary visual objects which did not move, but appeared as suddenly jumping around. She could see the objects clearly at different distances and locations, but she had no idea what happened to them between these locations. For her, the objects did not move, they just jumped from one position to the next, but there was nothing in between (Heywood and Zihl 1999).

If there is an explanation as to why a bilateral injury of area V5 leads to loss of temporal continuity, I expect the explanation must be constructed in terms of underlying non-phenomenal mechanisms rather than in terms of experiential relations within consciousness. Thus, co-consciousness and the (dis)continuity of consciousness must be explained by referring to mechanisms outside consciousness, rather than declaring such relations fundamental and inexplicable.

## **8. Phenomenal Spatiality and Co-Consciousness**

Dainton treats co-consciousness as a more fundamental feature of experience than its spatiality. By contrast, I am inclined to treat (phenomenal) spatiality as the fundamental feature of the phenomenal level. However, I do agree with Dainton that the spatiality in question is not substantival phenomenal spatiality (P-space), but rather a V-space.

It seems to me that the idea of a V-space might pave the way for an intelligible explanation of phenomenal unity. In the framework of multilevel biological explanation, the fact that consciousness is based on a non-phenomenal V-space suggests that there must be some high-level neural phenomenon in the brain (perhaps a coherent bioelectrical or electrophysiological field) which is the physical substrate of the non-phenomenal V-space. This spatially organized but non-phenomenal system is exactly the kind of entity that could help us to bridge the explanatory gap: it is in itself wholly non-phenomenal (a phenomenal vacuum, as Dainton says), yet it allows all the phenomeno-spatial organization to be manifested at the higher phenomenal level. It has one foot in

the non-phenomenal realm, the other in the phenomenal realm (I elaborate these ideas in Revonsuo, forthcoming). It is rather intriguing that cognitive neuroscientists have recently discovered that some sort of content-independent, spatial representation appears to be necessary for information to be incorporated into consciousness (Kanwisher 2001). If consciousness is based on a non-phenomenal, neurally implemented V-space, as I suggest, then that is exactly the direction in which we should expect the empirical evidence to converge.

This naturalistic multilevel approach to the explanation of consciousness denies that the level of consciousness could be both described and explained by staying within the phenomenal realm. To *describe* consciousness, it is necessary to confine oneself to the phenomenal level; but to *explain* consciousness, it is necessary to relate the phenomenal level to the nonphenomenal.

## 9. Conclusion

*The Stream of Consciousness* is a definite contribution to the field of consciousness studies. It shows that the field is in urgent need of a clear and systematic account of the phenomenal level itself: the level of purely phenomenal description. No other branch of science (neuroscience, cognitive science etc.) involved in the study of consciousness is going to deliver such an account; the study of consciousness must do the job on its own. Although Dainton's account of the unity and continuity of consciousness will have to be integrated with (and constrained by) phenomenal descriptions from other sources of evidence (especially altered states of consciousness and neuropsychology), his analyses provide a plausible and well-founded starting point for further work at the phenomenal level of description.

However, anyone who expects to find genuine explanations for the unity and continuity of consciousness will be disappointed. Co-consciousness is taken as the fundamental experiential relation that binds all the distinct phenomenal contents together and provides the unity and continuity of consciousness. But co-consciousness itself cannot be explained. Dainton's explanatory strategy forces him to stay within the phenomenal level as he constructs the explanations. I believe that many empirically minded scientists share my view that a genuine explanation goes beyond the level of description where the *explanandum* itself is identified.

A book titled *The Stream of Consciousness* easily raises high expectations, since it immediately brings to mind the classical writings of William James. I must say that at first I did wonder whether the book (or indeed *any* book) could fulfil such expectations. Now, after reading the *Stream of Consciousness* I am pleased to say that its content fully justifies the grand title.

## Acknowledgement

Antti Revonsuo is supported by the Academy of Finland (project #45704).

## References

- Damasio, A.R. (1999). *The Feeling of What Happens*. New York: Harcourt.
- Flanagan, O. (1992). *Consciousness Reconsidered*. Cambridge, MA: MIT Press.
- Heywood, C. A., & Zihl, J. (1999). Motion blindness. In G. W. Humphreys (Ed.), *Case Studies in the Neuropsychology of Vision* (pp. 1-16). Hove: Psychology Press.
- Kanwisher, N. (2001). Neural events and perceptual awareness. *Cognition*, 79, 89-113.
- Revonsuo, A. (1995). Consciousness, dreams, and virtual realities. *Philosophical Psychology*, 8, 35-58.
- Revonsuo, A. (1999). Binding and the phenomenal unity of consciousness. *Consciousness and Cognition*, 8(2), 173-185.
- Revonsuo, A. (2000). Prospects for a scientific research program on consciousness. In: Metzinger T (Ed.) *Neural Correlates of Consciousness*, 57-75. Cambridge, MA: MIT Press.
- Revonsuo, A. (2001). Can functional brain imaging discover consciousness in the brain? *Journal of Consciousness Studies*, 8(3), 3-23.
- Revonsuo, A. (forthcoming). Inner Presence. A Biological Research Program on Consciousness.
- Stubenberg, L. (1998). *Consciousness and Qualia*. Amsterdam: John Benjamins.
- Zihl, J., von Cramon, D., & Mai, N. (1983). Selective disturbance of movement vision after bilateral brain damage. *Brain*, 106, 313-340.