



# ALIUS BULLETIN

exploring the diversity of consciousness

Martin Fortier, François Lemoine & Raphaël Millière (eds.)

ROBIN CARHART-HARRIS


JAKOB HOHWY

TANYA LUHRMANN

SIMON MCCARTHY-JONES

REBECCA SELIGMAN

JENNIFER WINDT



n°1

FEBRUARY 2017

[aliusresearch.org](http://aliusresearch.org)



Copyright: © 2017 Fortier, Lemoine & Millière. This is an open access publication distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

## Acknowledgments

The editors wish to express their gratitude to Cordelia Erickson-Davis, Daniel Friedman, Ella Letort and Kai Woolner-Pratt for their valuable proofreading of the interviews, and to all the contributors for accepting to participate to this issue.



# Foreword

## About ALIUS

ALIUS is an international and interdisciplinary research group dedicated to the investigation of all aspects of consciousness, with a specific focus on non-ordinary or understudied conscious states traditionally classified as *altered states of consciousness*.

In Latin, *alius* means “different”. This lexical choice reflects the group’s mission to study the diversity of consciousness in a systematic manner. ALIUS puts a particular stress on the need for a naturalistic approach to all aspects of consciousness, including states and experiences which have long been unduly associated to parapsychology and pseudoscientific hypotheses.

To this end, it fosters a unique interdisciplinary collaboration of researchers, involving neuroscientists, psychologists, philosophers of mind, psychiatrists and anthropologists, towards the development of a systematic and scientific model of consciousness supported by both theoretical work and experimental studies. This collaboration may take the form of joint articles, blog posts, editorial work on special issues, thematic workshops and international conferences.

Find out more about the group on the website: [aliusresearch.org](http://aliusresearch.org)

## About the Bulletin

The ALIUS Bulletin is an annual publication featuring in-depth interviews with prominent scholars working on consciousness and its altered states (ASCs). The goal of the Bulletin is to present a clear outline of current research on ASCs across a variety of disciplines, with an emphasis on empirical work. It also aims at dispelling the widespread stigma that still plagues the notion of ASC, while allowing a wider audience to discover rigorous scientific work on the topic presented by authors in their own words.



# Table of Contents

## **Psychedelics and consciousness**

Robin Carhart-Harris

Interviewed by Martin Fortier & Raphaël Millière

1

## **On different ways of being conscious: modes of consciousness and the predictive mind**

Jakob Hohwy

Interviewed by Matthieu Koroma

17

## **The anthropology of mind: exploring unusual sensations and spiritual experiences across cultures**

Tanya Luhrmann

Interviewed by Martin Fortier

25

## **The phenomenon of voice-hearing: an interdisciplinary approach**

Simon McCarthy-Jones

Interviewed by Mathieu Frerejouan

37

## **Towards a biocultural approach to dissociative consciousness**

Rebecca Seligman

Interviewed by Arnaud Halloy

47

## **Relocating dreams on the conceptual map: how the analysis of sleep and dreaming challenges our taxonomy of mental states**

Jennifer Windt

Interviewed by Alessio Bucci & Raphaël Millière

53





## Contributors

**Robin Carhart-Harris** is a research fellow at Imperial College London, where he is the head of psychedelic research at the Centre for Neuropsychopharmacology of the Division of Brain Sciences. He has been conducting pioneering brain imaging studies of psychedelic drugs over the past six years, focusing on psilocybin and more recently LSD. He has published over forty papers in neuropharmacology and psychology.

**Jakob Hohwy** is a professor of philosophy at Monash University in Melbourne, Australia. His research focuses on philosophy of mind, philosophy of cognitive neuroscience and philosophical psychopathology. He has also been involved in a number of experimental research projects with neuroscientists and psychiatrists. His book *The Predictive Mind* (2013) gives a unified account of conscious perception within the Bayesian framework of predictive coding.

**Tanya Luhrmann** is a professor of anthropology, and (by courtesy) of psychology, at Stanford University. Her work explores “how people know what is real—particularly how they know that God is present”. She has published a book, *When God Talks Back* (2012) on American evangelicals’ understanding of mind, God and reality. She had previously published an acclaimed book, *Of Two Minds* (2000), on the psychodynamic and the biomedical approaches to mental illness.

**Simon McCarthy-Jones** is an associate professor of psychiatry at Trinity College in Dublin. He works on the phenomenon of voice-hearing from a neuropsychological, phenomenological, and historical point of view. He is the author of *Hearing Voices* (2012), and his next book on voice-hearing, *Can’t You Hear Them? The Science and Significance of Hearing Voices*, will be published in April 2017.

**Rebecca Seligman** is an associate professor of anthropology at Northwestern University. Her research covers a broad range of subjects in medical and psychological anthropology as well as cross-cultural psychiatry. She has written a book on possession and dissociation in an Afro-Brazilian cult (*Possessing Spirits and Healing Selves*, 2014) and co-edited a groundbreaking volume on culture and neuroscience (*Oxford Handbook of Cultural Neuroscience*, 2015).

**Jennifer Windt** is a lecturer in philosophy at Monash University in Melbourne, Australia. Her research centers on dreaming, philosophy of cognitive science, philosophy of mind, self-consciousness, and sleep. Her book *Dreaming: A Conceptual Framework for Philosophy of Mind and Empirical Research* (2015) develops an empirically-informed account of dreaming, and has been universally acclaimed as the most comprehensive book ever written on the topic.



# Psychedelics and consciousness

## An interview with Robin Carhart-Harris

by Martin Fortier and Raphaël Millière

Citation: Carhart-Harris, R., Fortier, M., & Millière, R. (2017). Psychedelics and consciousness. An interview with Robin Carhart-Harris. *ALIUS Bulletin*, 1, 1-16.

**Robin Carhart-Harris**

[r.carhart-harris@imperial.ac.uk](mailto:r.carhart-harris@imperial.ac.uk)  
Centre for Neuropsychopharmacology  
Imperial College, London, UK

**Martin Fortier**

[martin.fortier@ens.fr](mailto:martin.fortier@ens.fr)  
Institut Jean Nicod  
ENS/EHESS, Paris, France

**Raphaël Millière**

[raphael.milliere@philosophy.ox.ac.uk](mailto:raphael.milliere@philosophy.ox.ac.uk)  
Faculty of Philosophy  
University of Oxford, UK

### What got you interested in the topic of psychedelics?

It was an interest in the mind. At the time (about 12 years ago), I felt that my curiosity about the mind could be best satisfied by a study of psychoanalysis, so I was studying psychoanalysis for a master's degree. Someway in, I began to realize that the main methods of psychoanalysis to access the 'unconscious mind' are unreliable. The existence of the unconscious mind is a cornerstone of psychoanalytic theory - yet its existence is largely denied or at least neglected by mainstream psychology and psychiatry. Dreaming is perhaps psychoanalysis' most valued window in on 'the unconscious' but dreaming happens during sleep and it is easy to forget or confabulate our dream material. Free association also seemed vague and unreliable, so I felt compelled to ask whether anything else had been considered as a way of altering the structure of the mind in such a way that access to 'the unconscious' could be facilitated. At the time, I was somewhat ambivalent about psychoanalysis. It felt like its main tenets had substance but there was also a lot of philosophical noise getting in the way of what I saw as a much more fundamental problem, namely testing and potentially demonstrating that the basic constructs of psychoanalysis, such as *repression*, *the ego*, and *the unconscious mind* 'exist' and have substance in both a psychological and biological sense. I felt psychoanalysis needed to get its act together and better demonstrate these things otherwise it could make no real scientific progress and would instead remain insular and somewhat disconnected from mainstream science. I entertained the idea that a drug might exist that could alter the topography of the mind to aid access to the unconscious and then I discovered Stanislav Grof's book, *Realms of the Human Unconscious*:

*Observations from LSD Research*, and I thought - here it is! The most logical thing from there seemed to be to take this to brain imaging and measure what's going on in the brain as the unconscious mind is accessed under LSD or a similar drug - and that's basically the idea I took to David Nutt, then Professor of Psychopharmacology at the University of Bristol, as a PhD proposal. Although I've never managed to work on that specific project, I'm a lot closer than I was 12 years ago.

In addition to neuroscience, you studied psychoanalysis. Do you think some of the hypotheses of psychoanalysis on the human mind are relevant for scientific research?

Yes.

It seems fair to say that you are one of the proponents of neuropsychanalysis (e.g., Carhart-Harris & Friston, 2010).

In a sense, yes.

Some authors (e.g., Ramus, 2013) have worried that the neuropsychanalytic research program might be a nonstarter, because it amounts to projecting neuroscientific discoveries onto psychoanalytic concepts (but not the other way around).

The relationship should be circular to be healthiest. Arguably, psychoanalysis hasn't made any discoveries for quite some time though. Perhaps what psychoanalysis had to tell us, has already been told, and mostly just by a very select number of people (e.g. Freud, Jung and Klein).

The worry, then, is that in neuropsychanalysis, neuroscience does inform psychoanalysis, but psychoanalysis hardly ever informs neuroscience.

It may inform or even replace cognitive explanations however. I suppose the point is, psychoanalysis isn't discovering in the same way as neuroscience is, because its 'discoveries' were made many years ago, and through observation but not in a controlled and systematic way (unlike neuroscience).

Sceptics might nonetheless argue that the two domains are not mutually enriching. Do you think this is true?

Not necessarily. I think psychoanalysis can provide enriching explanations for things that cognitive psychology and neuroscience struggle with.

In the same line of skepticism, do you agree that only neuroscience (owing to its scientific research methods) brings new findings on the table?

Yes, these days only neuroscience brings new findings to the table but because I believe psychoanalysis can offer a richer (arguably the richest) account of the human mind and condition, I think it has useful things to say about certain neuroscientific discoveries. I accept it doesn't bring anything new but what it brings that is 'old' is still deeply resonant and informative.

Let's take the example of the 'discovery' of the default-mode network (DMN) for example, around the turn of the millennia. Cognitive neuroscientists have struggled with this system as they can't easily identify what its main function is. It seems to be involved in either too much or nothing specific at all. Some have even challenged whether the approach of studying 'resting-state' brain activity (the approach that helped develop the notion of a 'default-mode') is really meaningful at all or even consistent with what cognitive neuroscience is 'meant to be doing' (Morcom and Fletcher, 2007). Terms like 'stimulus independent thought' been used in reference to the network and its functioning, which are telling more about the priorities of cognitive psychology and its behaviorist counterpart than about the phenomena in question. I think it's more interesting to speculate about and endeavor to sample the nature and content of spontaneous cognition, and tools such as 'experience sampling' (Csikszentmihalyi and Larson, 1987), psychedelics, and neuroimaging may prove to be a powerful combination in this respect. True, the DMN may serve some essential physiological functions (Leech et al., 2014), and granted, rudimentary prototypes of the adult human DMN may be observable in lower species (Mantini et al., 2011) and infants (Doria et al., 2010) but laying the emphasis here seems more about evading the problem than tackling it – and moreover, these things needn't be inconsistent with psychoanalytic conceptions of mind-function – such as the idea that the ego offers a reserve of energy that can be 'invested' in 'objects' of interest (Carhart-Harris et al., 2008).

“ If you want to look for 'the unconscious' in the brain, you probably needn't look far from the DMN and particularly its subcortical/paralimbic components, such as the extended hippocampus. ”

In 2008 (Carhart-Harris et al., 2008) and later in 2010 (Carhart-Harris & Friston, 2010) and more recently (Carhart-Harris et al., 2014), I proposed that the default-mode network may relate to 'the ego' system, as described by Sigmund Freud. That hypothesis seems to have been quite popular (e.g. the 2010 paper with Karl Friston has been cited over 250 times) and I still feel relatively comfortable about having proposed it. All I would want to do now to qualify it, would be to say that in Freudian theory, in health, 'the unconscious' is largely continuous with the

ego and the ego can contain much that is of ‘the unconscious’, just as the unconscious contains much of the ego, i.e. they are not so clearly differentiable from each other and exist instead, in relative harmony. It is in ill health that the systems are in conflict. The point I’d like to make here, is that if you want to look for ‘the unconscious’ in the brain, you probably needn’t look far from the DMN and particularly its subcortical/paralimbic component/s, such as the extended hippocampus. If you’d rather stick to cognitive or non-psychoanalytic terms when talking about these systems, e.g. proposing that the DMN is related to ‘the narrative self’ – then that’s fine, I wouldn’t want to argue with that – but at least know that psychoanalysis has talked about these systems in depth for over a century (Freud, 1927; Freud, Strachey, & Freud, 2001) and may offer a level of explanation that you might find surprisingly useful and informative. Also, if you’re going to judge psychoanalysis and particularly Freudian theory, judge it directly by going to his original work and not to interpretations of it.

### How do you think psychedelics research can contribute to the search for the neural correlates of consciousness?

Psychedelics alter ‘consciousness’ in a marked and novel way. Most studies of consciousness address the problem of consciousness via looking at states of reduced consciousness, e.g. anesthesia, sleep, brain injury and/or illness. Psychedelics do not really reduce the level of wakefulness or alertness but they do fundamentally alter the quality of consciousness. Psychedelics can also help motivate a more nuanced version of the question, ‘what are the neural correlates of consciousness?’ Because if you look at that question critically, it’s pretty vague, i.e. the definition of consciousness is pretty vague. The dominant one is ‘consciousness is that which is lost when we fall into dreamless sleep and returns when we wake’ but I would challenge that and say instead that what is lost and returns either side of dreamless sleep is actually more or less the whole of the mind, and the whole of the mind is not just consciousness.

### What is your general view on the topic of consciousness?

It needs to better define its terms. If the field of ‘consciousness research’ clarified that the reason we’re really interested in consciousness is because we’re wonderstruck about the complex nature of *human* consciousness and want to better understand its basis in the brain, then that would be a good start, because then we could focus on humans and how they differ from other animals and perhaps realize that the existence of ‘the ego’ is quite a fundamental difference. After that, we might realize that rather than aimlessly studying ‘consciousness’ and approaching it from the same old vantage of reduced states of consciousness, we might better focus on self-consciousness and its neural correlates. A stronger focus on self-consciousness

and the nature of the self as a system (i.e. ‘the ego’ in psychoanalytic terminology) would help along the initiative of trying to better understand the brain basis of human consciousness – and why human consciousness is special. I think the endeavor to better understand the brain basis of ‘consciousness’ is too vague a question/problem, and unless I’m missing something (which I may be), I don’t think it’s going to come to much. It may sound anthropocentric but I think we need to better admit/acknowledge that the real motivator driving enthusiasm for ‘consciousness research’ is the promise of understanding why and how the human mind is different to that of other animals – and why/how we can reflect on our own thought and behavior. I’m actually open to the idea that it is something relatively recently learned or acquired though, as was argued by Julian Jaynes (1990). Jaynes’ stuff about the breakdown of hemispheric separation, I’m less convinced about however, but the idea that (reflective) consciousness is a relatively recent acquisition of the human mind, is a fascinating one – and has some intuitive appeal.

Do you endorse any of the following approaches to consciousness: a topological approach (claiming that there is an area of consciousness), an electrophysiological approach (claiming that consciousness consists of a specific rhythm of neural activity), or a reticular approach (claiming that consciousness consists of a specific connectivity between and within networks)?

If any of those, it would be the last one, but like I said, I think you need to start with a clearer definition of what the problem is, i.e. ‘what is consciousness?’ Are we talking about the reflective capacity of healthy adult humans? If yes, then I vote for focusing on the DMN and how it interacts with the rest of the brain. DMN-parahippocampal interactions are of particular interest to me in this regard. I wouldn’t want to say that the DMN as an object in isolation is consciousness or indeed ‘the ego’, but I would advise that we study its behavior, and how it develops, how its constituent parts combine to form a gestalt – and how this gestalt interacts with others in the brain, and how the interaction of these gestalts relates to subjective phenomena that we feel and can (mostly) be conscious of, and report.

How do you interpret the seeming discrepancy between Vollenweider et al.’s study (1997) and your own study (Carhart-Harris et al., 2012) regarding the effect of psilocybin on frontal cortex activity?

This is covered in our more recent *PNAS* paper on LSD (Carhart-Harris et al., 2016).

In your seminal paper you pinpoint the difference in recording techniques (PET vs. fMRI) as possibly explaining why one observes such a discrepancy between your findings and Vollenweider’s: “this discrepancy relates to the fact that the radiotracer used to measure glucose metabolism (18F-fluorodeoxyglucose) has a long half-life (110

min). Thus, the effects of psilocybin, as measured by PET, are over much greater timescales than indexed by our fMRI measures” (Carhart-Harris et al. 2012, p. 2141). This hypothesis is corroborated by a recent fMRI study of ayahuasca which also detects a decrease in frontal activity (Palhano-Fontes et al., 2015).

Not really, as that mPFC decrease was task-related I think. A team working at Kings College London replicated our arterial spin labeled (ASL) intravenous psilocybin study and found decreased blood flow however. As I say in the 2016 *PNAS* paper, I think the discrepancy is due to the methods. Unless you are *au fait* with the methods, it’s easy to think brain imaging literally reads-out ‘brain activity’, but that isn’t at all the case. The process of recording to analysis to brain images, involves a number of assumptions, e.g. people often assume that the BOLD signal of fMRI is measuring brain activity and similarly that the ASL signal represents brain activity – but that’s not always the case. For example, a direct vascular action of a drug can interfere with the BOLD and ASL signal and give you a read-out that you misinterpret as being a change in ‘brain activity’. I think in the case of the cerebral blood flow (CBF) reductions that have been seen with intravenous psilocybin, it may represent an initial vasoconstriction in the brain but it could also represent an initial change (reduction) in brain activity (when I say ‘brain activity’, I really mean *neuronal activity*). The overlap between location of the CBF reductions and changes in oscillatory power seen with intravenous psilocybin and MEG was quite remarkable however (Muthukumaraswamy et al., 2013) – and importantly, the latter is a much more reliable and direct measure of brain activity. Basically, if you want this matter to be comprehensively resolved, you need to do a bit more work, perhaps in animals, looking at vascular action of intravenous psilocybin. But I’m not sure it’s a problem that is that important to worry about however, as the field has moved on, as have our techniques, and we’re now using better measures that more directly sample neuronal activity, and these are yielding reliable results across drugs and study teams. I would basically advise someone worried about this matter, not to worry, as consistent principles about the acute action of psychedelics on the brain are emerging and will be shown to be quite reliable. Our 2016 *PNAS* paper on LSD in the most comprehensive in the sense. In brief, look to the present and future for the answer/s.

“ Increased PFC metabolism is not a good explanation for how psychedelics work to produce their characteristic effects. ”



In addition to the recording technique, do you think that the mode of administration could also explain the discrepancy: i.e., oral administration in Vollenweider's study vs. intravenous administration in your own?

Possibly. I feel pretty confident that increased PFC metabolism is not a good explanation for how psychedelics work in the human brain to produce their characteristic psychological effects however.

On a related note, what are the potential limitations of fMRI imaging to study psychedelics? More generally, what are the benefits and downsides of each available monitoring techniques to study the activity of the brain on psychedelics?

fMRI doesn't measure brain activity directly. Psychedelics may well have a direct vascular action and this could confound interpretations of some fMRI findings with psychedelics, such as CBF measures. Fludeoxyglucose PET has very poor to no temporal resolution and so provides little/no information about brain dynamics. Simultaneous EEG-fMRI is an important way forward and we're embracing that. Other techniques will likely also emerge in time. For example, the temporal resolution of fMRI is improving – something we might also try to make the most of. Dynamic EEG-fMRI measures twinned with experience sampling may be a powerful way forward. Similarly, decoding methods could prove useful.

Psychedelics researchers have emphasized the importance of set and setting, that is the participant's state of mind and environment when the drug is administered. How do you think set and setting could be modeled and controlled with more precision within experimental studies, given the influence they have on the participant's experience?

It's a challenge because you want to test 'set and setting' as variables but you also want to maintain safety and certain ethical standards. Music offers a good means of modulating 'setting' and 'set' requires that we sample 'where people are at' psychologically prior to the trip itself. This hasn't really been done properly yet – at least by us. You could also try and manipulate expectations to manipulate set – and I suppose we've done that to an extent when looking at suggestibility – but far more work could and should be done here.

What are the difficulties of gathering reliable data about the subjective effects of psychedelics, and how can they be overcome?

Reports are given in retrospect and so can be unreliable and sensitive to biases. Subjective reports are difficult to obtain in real time however, as collecting them will affect the experience and also language skills may be impaired under a potent drug. One potential solution is 'experience sampling'. I'm quite keen to incorporate this into our work. I'm also aware of more sophisticated interviewing techniques

that are being developed - which may be particularly useful. Video and audio recording of sessions and interviews would also be useful.

In order to chart the neurophenomenology of psychedelic experience, the questionnaires you use resort to concepts such as “looking strange” or “having a supernatural quality” and correlations are subsequently made between subjective reports (the feeling that an experience was strange or supernatural) and neurophysiological data (a certain pattern of neural activity). Now, concepts such as that of strangeness or supernaturalness are notoriously ambiguous and likely to be interpreted in various ways by subjects. Thus, when correlations are made between subjective reports and neural patterns, what is being correlated may be extremely different in one case and in another.

To elaborate a little on the example of supernaturalness, it has been shown that this property is variably ascribed (depending on one’s personality, one’s culture and one’s level of expertise) to experiences characterized as: (i) being highly fluent, (ii) being highly disfluent, (iii) being sensorially vivid, (iv) being numinous, (v) being non-dual, (vi) featuring extraordinary beings, (vii) involving a loss of the sense of agency, etc. (see for example: Shanon, 2002; Taves, 2009; Laughlin, 2011; Luhrmann, 2012; Halloy, 2015). This diversity of meanings lying behind the concept of supernaturalness is highly problematic.

Do you think when they are dealing with such polysemous words, neuroscientists should finesse the concepts they use in their questionnaires by consulting anthropologists and phenomenologists who have extensively studied the underlying polysemy of these terms?

I do think we need to better define our terms, yes. Words are certainly vulnerable to interpretation and biases, but language is a difficult prison-house to escape from, however you might try. Words that are especially vulnerable to different interpretations can be problematic and if you want them to be interpreted in a particular way, then you could provide participants with some kind of briefing about what is meant by the term/s in question but then that could bias/prime people to see things ‘how you want them to see things’, thus causing a confirmation bias.

“Words are certainly vulnerable to interpretation and biases, but language is a difficult prison-house to escape from.”

You have investigated a phenomenon known as drug-induced ego dissolution (DIED), which is usually described as a breakdown of one's 'sense of self' and a feeling of unity with one's environment (Lebedev et al., 2015). You have suggested that DIED might be explained as a breakdown of the so-called narrative self, that is the network of beliefs, thoughts and autobiographical memories associated with being the particular person one self-identifies with. This is consistent with subjects reporting feeling as if they were not a person anymore when undergoing DIED during psychedelics use. However, there is room for debate about whether DIED is merely a disruption of the narrative self, and not also of lower-level self-specifying processes, related to the so-called « minimal » or « bodily » self. There is a long tradition in philosophy of discussing the idea of the minimal self (also called « sense of mineness ») as the pre-reflective, nonconceptual feature of consciousness in virtue of which my experiences feel mine (Zahavi, 2014). This notion has also come under investigation in cognitive science and psychiatry. According to neurocognitive accounts of the minimal self, especially in a predictive coding framework, it is crucially linked to multisensory integration (particularly of visuotactile and vestibular input), interoception and homeostatic regulation (Christoff et al., 2011; Limanowski & Blankenburg, 2013; Apps & Tsakiris, 2014). From a phenomenological point of view, it might be reducible to body ownership, self-location, and the experienced direction of the first-person perspective (Blanke & Metzinger, 2009). The minimal self does not depend on high level cognitive processes such as self-related beliefs and first-person thoughts, but rather on low-level bodily/perceptual processes. In this framework, the minimal self is a necessary condition for reflective self-consciousness (self-related thoughts), which in turns enables self-related beliefs (narrative self). If this is right, one could expect the narrative self to break down if the minimal self is disrupted, and this is indeed what appears to be the case to some extent in schizophrenia.

The idea that DIED is primarily a disruption of the minimal self seems consistent with the fact that classical psychedelics induce hallucinations and not delusions. In a predictive coding framework, given that drug-induced hallucinations presumably stem from impaired bottom-up processing coupled with relatively preserved top-down processing, it seems reasonable to expect excessive prediction error signaling to yield aberrant predictions and a disruption of multisensory integration, ultimately leading to the breakdown of the minimal self. In other words, DIED would primarily be about a perceptual/bodily anomalous processing that leads to the feeling to being selfless, rather than a cognitive anomaly leading to delusions impairing self-narratives.

Do you think that such an account of DIED has any plausibility? How might further studies bear on this debate?

The ego, in the Freudian sense, incorporates the narrative and bodily self. I understand the differences between these different aspects of ‘self’ but I’m not sure how important it is that we chop them up in order to understand DIED. I tend to think when people rate DIED they are recognizing that it is ‘the ego’ in the Freudian sense that is compromised, i.e. the self as a *system*. Generally speaking, the way people understand and use the term ‘the ego’ in everyday parlance, is more or less consistent with Freud’s account of it. Our ego-dissolution inventory puts emphasis on reduced ‘self-importance’ and a sense of connectedness or oneness (e.g. to self, others and nature), which naturally accompanies DIED. I suppose what’s happening is that boundaries necessary for the existence of the ego, breakdown and the sense of connectedness is the inevitable result. People can fight DIED but then they won’t feel the connectedness.

“ I tend to think when people rate DIED they are recognizing that it is ‘the ego’ in the Freudian sense that is compromised. ”

The DMN is usually associated with mind-wandering, self-reflection and introspection while the Task Positive Network (TPN) is associated with the orientation of the mind towards the external world. Elaborating on this distinction, you notice that these two networks are anti-correlated in order to guarantee the functionally important distinction between the internal world and the external one. Now, with your colleagues, you have shown that the DMN/TPN anti-correlation is significantly decreased after psilocybin intake. You have proposed that the diminution of the orthogonality between the DMN and the TPN was a very plausible explanation as to why psychedelic experiences are characterized by a “collapse of dualities” and by the “disturbance in one’s sense of self, and particularly one’s sense of existing apart from one’s environment” (Carhart-Harris et al., 2014, p. 16). Your proposal seems very intuitive indeed: if the DMN and the TPN stop being strongly anti-correlated this should result in a loss of neat distinction between the internal and the external.

However, there is good evidence suggesting that the loss of the internal/external dichotomy is to be explained by something different from the disturbance of the DMN/TPN anti-correlation. Indeed, the “collapse of dualities” and the “disturbance in one’s sense of self” seem to characterize the phenomenology of both psilocybin-induced experiences and ayahuasca-induced ones. Now, given that ayahuasca experiences also involve some confusion between the internal and the external, we would expect to find a clear decrease of the DMN/TPN anti-correlation after ayahuasca

intake. Remarkably enough, this doesn't seem to be the case (Palhano-Fontes et al., 2015). How do you interpret these findings?

The specifics of the data processing and statistical thresholds can explain Palhano-Fontes's alleged negative result. If you look closer at the method they used however, I think they did find this effect (i.e. reduced DMN-TPN orthogonality) – and I'm very confident others will in the future.

Do you think that it can still be maintained that the loss of a clear dichotomy between the internal and the external is to be explained by the decrease of the DMN/TPN anti-correlation?

Yes. Although we don't need to be too specific about the TPN, as there are a few "TPNs". "Task positive network" is a bit of a vague term for a network to be honest.

On a more methodological note, there has been some intense debate as to whether regressing out the global signal is a sound way of measuring anti-correlation between networks (Fox et al., 2009; Murphy et al., 2009). If regressing out the global signal turns out not to be valid, then the DMN/TPN anti-correlation could simply be an artefact. What is your take on this methodological issue?

We did our most recent analysis (*PNAS* 2016) with and without global signal regression and found the same result. This is discussed in our paper. You just have to be careful about the terms you use and, for example, go with "orthogonality" instead of "anti-correlation". I think this is discussed in Leor Roseman's *Human Brain Mapping* paper (2016).

In your 2014 *Frontiers* paper you sketch a general model of Altered States of Consciousness (ASCs) which aims at theorizing not only psychedelic states but more broadly any kind of altered or anomalous state (such as psychosis, coma, dreaming, etc.). Your proposal is that entropy is a spectrum which can take different values: for example, early psychosis is characterized by high entropy whereas coma or sedation are characterized by low entropy. Your model is also based on another dimension: that of criticality (i.e., the ability of the brain to reach certain critical thresholds beyond which new complex properties emerge, notably through cascade-like processes) (e.g., Beggs & Plenz, 2003; Chialvo, Balenzuela, & Fraiman, 2008). Interestingly, in your model, you defend the view that entropy and criticality are closely correlated: the highest the entropy, the highest the criticality; conversely, the lowest the entropy, the lowest the criticality.

Did I say that? Criticality is just one thing, i.e. it's a critical point, so you can't really have low or high criticality – but you can be above or below it. Being super-critical,



in the sense of being *above* a critical point, would be most consistent with a more entropic (random) state.

Don't you think that there are cases in which entropy can vary independently of criticality? For example, it has been proposed that certain altered states can involve plenty of prediction errors while the weight (or the accuracy) ascribed to these predictions errors remains abnormally low; conversely, predictions errors can be low while the accuracy ascribed to these limited predictions errors is very high. As Fletcher and Frith put it, "a relatively small prediction error might be given undue weight (if the uncertainty is underestimated), leading to a false inference. Alternatively, excessive noise might dilute the effects of even a large prediction-error signal, leading to a reluctance to accept an inference as adequately explaining the input." (Fletcher & Frith 2009, p. 55) Let's take the example of a state characterized by high prediction errors and low accuracy. In such a case, priors couldn't be revised and updated as they should and they would end up being abnormally steady (e.g., Fletcher & Frith, 2009; Adams, Brown, & Friston, 2015). The brain would thus combine both high entropy (there are plenty of prediction errors) and sub-criticality (high-level mental states are not malleable). This seems to contradict the idea that high entropy and high criticality always go hand in hand.

They don't, as you can't have high criticality, but you have "super-criticality" – and generally speaking, though perhaps not absolutely, super-criticality would be consistent with high entropy.

In your work, you have occasionally pondered upon what your findings suggest as to what the broad structure of the brain is. You have notably advanced that two models of the brain can shed very interesting light on psychedelics: the "free-energy principle" put forth by Karl Friston (2010; 2006) and the "reducing valve" model proposed by Aldous Huxley (1954). In your seminal study you seem to suggest that these two models are in fact almost identical (Carhart-Harris et al. 2012, p. 2142).

Huxley's model was explicitly inspired by Charlie D. Broad (1953, chap. 1), and Broad himself borrowed the "reducing valve" metaphor from Henri Bergson (1994 [1896]). Now, if we look at it closely enough, the model endorsed by Bergson, Broad and Huxley is arguably quite different from that defended by Friston. Bergson's key idea is that the default state of the brain is that of being overwhelmed by sensory data coming from the world; the brain's function is thus to diminish the amount of sensory data reaching consciousness. This view straightforwardly contradicts Kant's transcendental idealism: by and large, Kant says that the data coming from the world are poor and that the brain later enrich them; on the contrary, Huxley, Broad and Bergson claim that the data coming from the world are too rich to be processed by a normal brain and that the brain is precisely there to filter the massive surge of data coming from the world.

In this regard, Friston's model seems to fundamentally differ from that advanced by Bergson and his followers. The free-energy and predictive coding framework has it that prediction comes first and prediction errors (sensory data coming from the world) come next to rectify and update top-down predictions. While in Bergson's model the brain is simply a filter which cannot by itself generate any conscious representation (the richness of representations is to be found in the world), in Friston's model, the brain is notoriously able to generate conscious representation and it is even able to do so before getting any sensory feedback. In other words, the free-energy and predictive view contends that the role of the world is corrective and that the role of the brain is constitutive whereas Bergson, Broad and Huxley hold exactly the opposite view: for them, the world is constitutive and the brain is simply corrective.

Do you agree with the distinctions drawn here between Bergson's model and Friston's model? Five years after your seminal study, do you tend to think that the "reducing valve" model provides the best account of psychedelics or that the "predictive" model does a better job?

Friston is a leading contemporary neuroscientist whose free-energy principle is an elegant, empirically-informed model of how the brain works. Huxley was a brilliant author and philosopher but not a neuroscientist. The reducing valve idea is quite nice as a metaphor but I think people take it too literally and sometimes even want to use it in a sort of pseudo-scientific way, to suggest that there is a filter that stops us seeing what's really "*out-there*" in a matrix-esque kind of way.

If the metaphor is useful however, it's useful because it proposes that "the brain, in main, is eliminative rather than productive". That idea is consistent with the free-energy principle because top-down inferences work to explain bottom-up sensory information – so there's some functional suppression going on but suppression in the sense of predictive processing. Basically, I wouldn't worry too much about Bergson, Broad and Huxley or indeed Kant when it comes to a contemporary account of how the brain works. It's best to see where the field is now, and Friston's free-energy model is one of the best the field has.

In your opinion, what are the next steps for psychedelics research?

There's so much but better predicting response to the psychedelics is a good example of one potentially fruitful area.

---

## References

- Adams, R., Brown, H., & Friston, K. (2015). Bayesian inference, predictive coding and delusions. *Avant: Trends in Interdisciplinary Studies*, 5(3).
- Apps, M. A. J., & Tsakiris, M. (2014). The free-energy self: a predictive coding account of self-recognition. *Neuroscience and Biobehavioral Reviews*, 41, 85–97.
- Beggs, J. M., & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *The Journal of Neuroscience*, 23(35), 11167–11177.
- Bergson, H. (1994 [1896]). *Matter and memory*. New York: Zone Books.
- Blanke, O., & Metzinger, T. (2009). Full-body illusions and minimal phenomenal selfhood. *Trends in Cognitive Sciences*, 13(1), 7–13.
- Broad, C. D. (1953). *Religion, Philosophy and Psychological Research*. London: Routledge / Kegan Paul.
- Carhart-Harris, R., Erritzoe, D., Williams, T., Stone, J. M., Reed, L. J., Colasanti, A., ... Nutt, D. J. (2012). Neural correlates of the psychedelic state as determined by fMRI studies with psilocybin. *Proceedings of the National Academy of Sciences*, 109(6), 2138–2143.
- Carhart-Harris, R., & Friston, K. (2010). The default-mode, ego-functions and free-energy: a neurobiological account of Freudian ideas. *Brain: A Journal of Neurology*, 133, 1265–1283.
- Carhart-Harris, R., Leech, R., Hellyer, P., Shanahan, M., Feilding, A., Tagliazucchi, E., ... Nutt, D. (2014). The entropic brain: A theory of conscious states informed by neuroimaging research with psychedelic drugs. *Frontiers in Human Neuroscience*, 8.
- Carhart-Harris, R., Mayberg, H., Malizia, A., & Nutt, D. (2008). Mourning and melancholia revisited: correspondences between principles of Freudian metapsychology and empirical findings in neuropsychiatry. *Annals of General Psychiatry*, 7, 9–9.
- Carhart-Harris, R., Muthukumaraswamy, S., Roseman, L., Kaelen, M., Droog, W., Murphy, K., ... Nutt, D. (2016). Neural correlates of the LSD experience revealed by multimodal neuroimaging. *Proceedings of the National Academy of Sciences*, 113(17), 4853–4858.
- Chialvo, D., Balenzuela, P., & Fraiman, D. (2008). The brain: What is critical about it? *Conf. Proc. Am. Inst. Phys.*, 1028, 28–45.
- Christoff, K., Cosmelli, D., Legrand, D., & Thompson, E. (2011). Specifying the self for cognitive neuroscience. *Trends in Cognitive Sciences*, 15(3), 104–112.
- Csikszentmihalyi, M., & Larson, R. (1987). Validity and reliability of the experience-sampling method. *The Journal of Nervous and Mental Disease*, 175(9), 526–536.
- Doria, V., Beckmann, C., Arichi, T., Merchant, N., Groppo, M., Turkheimer, F., ... Edwards, D. (2010). Emergence of resting state networks in the preterm human brain.



- Proceedings of the National Academy of Sciences of the United States of America*, 107(46), 20015–20020.
- Fletcher, P., & Frith, C. (2009). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews Neuroscience*, 10(1), 48–58.
- Fox, M. D., Zhang, D., Snyder, A. Z., & Raichle, M. E. (2009). The global signal and observed anticorrelated resting state brain networks. *Journal of Neurophysiology*, 101(6), 3270–3283.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology - Paris*, 100(1-3), 70–87.
- Halloy, A. (2015). *Divinités incarnées: L'apprentissage de la possession dans un culte afro-brésilien*. Paris: PETRA.
- Huxley, A. (1954). *The Doors of Perception and Heaven and Hell*. London: Harper & Brothers.
- Jaynes, J. (1990). *The origin of consciousness in the breakdown of the bicameral mind*. New York: Houghton Mifflin.
- Laughlin, C. D. (2011). *Communing with the Gods: Consciousness, culture and the dreaming brain*. Brisbane: Daily Grail.
- Lebedev, A. V., Lövdén, M., Rosenthal, G., Feilding, A., Nutt, D. J., & Carhart-Harris, R. L. (2015). Finding the self by losing the self: Neural correlates of ego-dissolution under psilocybin. *Human Brain Mapping*, 36(8), 3137–3153.
- Leech, R., Scott, G., Carhart-Harris, R., Turkheimer, F., Taylor-Robinson, S., & Sharp, D. (2014). Spatial Dependencies between large-scale brain networks. *PLOS ONE*, 9(6).
- Limanowski, J., & Blankenburg, F. (2013). Minimal self-models and the free energy principle. *Frontiers in Human Neuroscience*, 7.
- Luhrmann, T. (2012). *When God talks back: Understanding the American evangelical relationship with God*. New York: Alfred Knopf.
- Mantini, D., Gerits, A., Nelissen, K., Durand, J.-B., Joly, O., Simone, L., ... Vanduffel, W. (2011). Default Mode of Brain Function in Monkeys. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 31(36), 12954–12962.
- Morcom, A., & Fletcher, P. (2007). Does the brain have a baseline? Why we should be resisting a rest. *NeuroImage*, 37(4), 1073–1082.
- Murphy, K., Birn, R. M., Handwerker, D. A., Jones, T. B., & Bandettini, P. A. (2009). The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced? *NeuroImage*, 44(3), 893–905.
- Muthukumaraswamy, S., Carhart-Harris, R., Moran, R., Brookes, M., Williams, T., Errtizoe, D., ... Nutt, D. (2013). Broadband cortical desynchronization underlies the human psychedelic state. *Journal of Neuroscience*, 33(38), 15171–15183.

- Palhano-Fontes, F., Andrade, K. C., Tofoli, L. F., Santos, A. C., Crippa, J. A. S., Hallak, J. E. C., ... Araujo, D. B. de. (2015). The psychedelic state induced by ayahuasca modulates the activity and connectivity of the default mode network. *PLOS ONE*, *10*(2), e0118143.
- Ramus, F. (2013). What's the point of neuropsychanalysis? *The British Journal of Psychiatry*, *203*, 70–71.
- Roseman, L., Sereno, M., Leech, R., Kaelen, M., Orban, C., McGonigle, J., ... Carhart-Harris, R. (2016). LSD alters eyes-closed functional connectivity within the early visual cortex in a retinotopic fashion. *Human Brain Mapping*, *37*(8), 3031–3040.
- Shanon, B. (2002). *The antipodes of the mind: Charting the phenomenology of the ayahuasca experience*. New York: Oxford University Press.
- Taves, A. (2009). *Religious experience reconsidered: A building-block approach to the study of religion and other special things*. Princeton NJ/Oxford: Princeton University Press.
- Vollenweider, F., Leenders, K., Scharfetter, C., Maguire, P., Stadelmann, O., & Angst, J. (1997). Positron Emission Tomography and Fluorodeoxyglucose Studies of Metabolic Hyperfrontality and Psychopathology in the Psilocybin Model of Psychosis. *Neuropsychopharmacology*, *16*(5), 357–372.
- Zahavi, D. (2014). *Self and Other: Exploring Subjectivity, Empathy, and Shame*. Oxford University Press.

# On different ways of being conscious

## Modes of consciousness and the predictive mind

An interview with  
Jakob Hohwy

By Matthieu Koroma

Citation: Hohwy, J. & Koroma, M. (2017). On different ways of being conscious: modes of consciousness and the predictive mind. An interview with Jakob Hohwy. *ALIUS Bulletin*, 1, 17-23.

**Jakob Hohwy**

[jakob.hohwy@monash.edu](mailto:jakob.hohwy@monash.edu)

Department of Philosophy

Monash University, Melbourne, Australia

**Matthieu Koroma**

[mkoroma@ens-paris-saclay.fr](mailto:mkoroma@ens-paris-saclay.fr)

Department of Cognitive Studies

École Normale Supérieure, Paris, France

Could you first explain how you have come in your academic career to adopt a strong empirically oriented philosophy – what triggered your interest in the study of the variety of conscious states?

It is hard to say what actually drives a particular career path – lots of coincidences, I suspect. Like most young philosophers back in the 1990's, I did not have much interest in, or understanding of, empirical science. I saw myself as engaged in a priori conceptual analysis. Most of my research was in philosophy of language. But I got to share an office with Ian Gold, who had done serious work in psychopathology. Ian Gold introduced me to theories of delusion formation and thereby to cognitive neuroscience. It immediately appealed to me and I think that deep down I must always have known that the philosophical problems that I wrestled with in fact do relate to science. In other words, these philosophical problems can't be addressed only with conceptual analysis (though conceptual analysis can be essential to science!). Even though I continued traditional philosophical work, I sought out more and more researchers from neuroscience and other disciplines. I talked a lot with Andreas Roepstorff at Aarhus who was busy building a research network around cognitive neuroscience and got intrigued by the clever experiments conducted there by Chris Frith and others. I think the overarching philosophical problem that interests me, and has been the focus of my work for a long time, is how we (or our brains) make sense of the sensory input that hits the senses. Perhaps I am interested in this because I am myself often perplexed at what the world around me means! Perception is, when you think about it, pretty amazing, and the varieties of conscious experience that accompany perception are so rich and multifaceted that it is hard to believe that a biological organ like the brain is responsible for it all. Of course, I would like to know what the ultimate nature of consciousness is – how the

brain gives rise to any kind of consciousness – but I am actually more interested in what shapes our conscious experience.

Could you please describe briefly the concept of a “mode of consciousness” and how it differs from other accounts of consciousness?

I wouldn't say that a “mode of consciousness” is an account of consciousness that is in competition with other accounts. It is more a way to point to an aspect of consciousness that hasn't received so much attention. Most researchers focus either on contents of consciousness (“I experience a red rose”) or levels or states of consciousness (“How conscious are you?”). What Tim Bayne and I did was simply highlight that there are global states of consciousness too – which we called “modes”, or ways, of being conscious. Modes are not just a matter of accumulating all the contents, they are a global way of having or characterizing those contents. Similarly, modes are not just a matter of having a level of consciousness, they are a way of characterizing what that level is (besides, we don't really think there are levels of consciousness as such; see our new paper in *Trends in Cognitive Sciences* (Bayne et al. 2016)). We argued that perhaps we can learn something new about consciousness from considering modes and different dimensions of consciousness rather than just levels and contents. A recent study in *Neuroimage Clinical* by Sergent and colleagues intriguingly supports this multidimensional approach to consciousness (Sergent et al., 2017).

“There are global states of consciousness too – which we called “modes”, or ways, of being conscious.”

You describe modes of consciousness as global states of consciousness being mutually exclusive (here I am referring specifically to when you say that a “creature cannot be in two distinct modes of consciousness at once”, p. 60 in “Modes of Consciousness” (Bayne & Hohwy, 2016)). Is this necessitated by unity being a fundamental property of conscious experience, as defended by Tim Bayne (2002), or does it add something else to the understanding of consciousness?

This is a good question. Perhaps Tim thinks there is a link between unity and modes – that is not a bad thought! But the main idea to begin with was that when we characterize modes, it seems we can't find cases where two modes overlap. I am quite open to hearing counterexamples to this though. A lot obviously hangs on how modes are defined. Considering a mode as a multidimensional construct (determined by dimensions of content and functionality, and perhaps arousal states), one can imagine two creatures being in the same mode for different reasons

(being placed in different positions on different dimensions might end up determining the same state overall). It is less clear how there could be two modes in the same creature, from this kind of perspective.

Following the previous question, functionally, it appears that for example some parts of the brain can enter sleep states while the rest of the brain remains in an awake state and affect properties of the conscious experience (Vyazovskiy et al., 2001; Hung et al., 2014). This phenomenon is experience dependent, and could affect differentially functional properties of the conscious experience (Sarasso et al., 2014). For example, “modes” have more of a familial resemblance rather than a shared core set of functional properties (or dispositional properties). I wonder if you would prefer to frame these transitional or mixed states as either “awake” or “asleep” with some gradation along the relevant dimensions and how you would accommodate these transitions in your model?

Again, this is a nice type of case to think about. One way to go would be to label that mixed case as its own overall, global mode. This is almost cheating, of course, since we naturally would want some taxonomy of modes (e.g. either “awake” or “asleep”). But perhaps what the multidimensional account teaches us is that modes can come in many different forms, and that as long as there is something it is like to be in that mode from a global perspective, then it is a mode, even if it is determined by an unusual constellation of positions on the various dimensions making up the current global conscious state. You might think that this move brings us in to the vicinity of dependence on a unity view, à la Bayne, as mentioned in the question above.

You describe a criterion for defining modes as “something to be like in, from a global perspective”. As you explain in “Modes of consciousness” (Bayne & Hohwy, 2016), modes are not to be understood as phenomenological constructs but rather characterizations of distinguishable states of consciousness. Nonetheless, it seems possible to develop some phenomenology or subjective knowledge about modes, in the forms of noetic feelings (e.g. feeling tired) or metacognition (e.g. becoming lucid during dreams). How do such metacognitive instances relate to modes and what could their function be in the conscious process?

This is exactly the kind of question we are hoping will arise once we begin focusing more on modes of consciousness. It opens up possibilities of new discussions. For noetic feelings, there is probably a spectrum from local to global (for discussion, see Dokic 2012). A feeling of knowing seems local to a particular content (e.g. “I feel I know what the capital of Latvia is but I can’t retrieve it”). A feeling of tiredness, or various moods, may be more global, coloring the entire global state of consciousness across both content and functional dimensions; this would make it closer to a mode. I think it is right that there is a distinct phenomenological feel to being tired but

apart from describing the physiological aspects (“I can’t keep my eyes open”) and offering metaphors (“feels like I am walking through treacle”) it is perhaps hard to pinpoint these feelings phenomenologically. Metacognitive states more generally are good candidates for teaching us about modes, I think. However, much here depends on what such states amount to. Feelings of uncertainty or confidence for example might or might not be important for the overall conscious process. Some people think that metacognition is essential to consciousness, others do not. But irrespective of their role in theories of consciousness, I am quite interested in the idea that such feelings in fact modulate the global state of consciousness. This also sits well with my pet theory of prediction error minimization: if the brain engages in precision-weighted prediction error minimization, then it would be natural to assume that it continuously assessed the overall rate of prediction error minimization and expected precisions, and brings this to bear on perceptual inference in a global sense. Developing these vague thoughts further is a project for the future (or for someone else)!

How would a multidimensional account tackle the apparition of entirely new abilities or functions in some states of consciousness (e.g. in some drug-induced or meditative practices)? Could we frame the emergence of new abilities in different conscious states as a hierarchical ordering of dimensions, where new dimensions appear to further classify subgroups of conscious' states, or as interaction between functional dimensions (by hierarchical, I am rather picturing a bush-like structure, just as in evolutionary science, rather than a single tree)?

Perhaps this is in fact an empirical question as well as a conceptual one – it may require extensive empirical work to somehow vary the dimensions of consciousness and figure out how they belong together and build on each other. It might be hierarchical in some sense (and I like the idea of wayward tree more than a linear hierarchy).

Following the previous question, I wonder then if all conscious states share the same multidimensional space. Would you identify some fundamental dimensions of consciousness on to which any state of consciousness can be projected?

We have certainly tried to argue that there are no levels of consciousness in any strict, straightforward sense (in the *Trends in Cognitive Sciences* paper mentioned earlier (Bayne et al., 2016)). Part of this conclusion comes from our suspicion that consciousness itself is not something that can be said to come in degrees but that rather the dimensions that determine global states of consciousness (in terms of content ranges, and functional properties such as attention) can come in degrees. It seems quite likely to me that these dimensions will interact with each other and thereby create fluctuations that are not easy to order in any hierarchical way.



“Consciousness itself is not something that can be said to come in degrees, but rather the dimensions that determine global states of consciousness can come in degrees.”

The diversity of modes suggests that consciousness could be realized in a multiple fashion rather than relying on a single mechanism. Could this consideration be relevant in the search of neuronal correlates of consciousness?

I very much think that the search for the neural correlates of consciousness needs to take modes or global states of consciousness into consideration. I am not sure multiple realization is the best way to describe this, mainly because that discussion pertains to metaphysical concerns about functionalism and identity theory, whereas the neural correlates is an empirical issue. But the thought is right: the same perceptual content might be experienced under different modes, such that the underlying neural mechanism differs. This is something that is dealt with nicely in multifactorial experimental designs where modes could be a factor alongside content. For example, one could look at binocular rivalry in different modes of consciousness (e.g. tired vs. alert). The idea is that contents might interact with modes such that the mechanism for consciousness differs for different modes and different contents. Note that the interaction may go both ways: we might perceive some object differently depending on the mode (a friendly face may be perceived as hostile when in a low mood), and some modes may be modulated by some contents (perceiving the tiger will make you feel less tired).

Do you think that a multidimensional account of consciousness would benefit from the identification of trade-offs linked to biological constraints? For example, using the externally-oriented vs. self-oriented axis to study wake vs. sleep, rather than just a unidimensional conception of connected vs. disconnected from the environment?

I think that is very likely. It seems that there many dimensions that can go into what determines a mode of consciousness. This will involve interoceptive states, aspects of action and decision-making, and perceptual inference concerning self vs. world. I agree it is important to move away from the kind of “environment connectedness” that typical bedside tests of consciousness are most focused on. This relates to the recent spate of studies of vestiges of consciousness in disorders of consciousness (e.g. work by Adrian Owen and Steven Laureys (Owen et al., 2009)). They show that consciousness does not need to be directly related to responsiveness to the environment but can be a more endogenous process. This invites the thought that

---

perhaps in some modes, there is a very strong focus on internal, endogenous, self-related aspects of consciousness.

The frameworks of Predictive Coding and Bayesianism have been successful in explaining several aspects of cognition in a neurocomputational framework (Hohwy, 2013). Some work has been done to link this to phenomenology (e.g. the work from Corlett's laboratory). How do predictive and Bayesian approaches shed light on the understanding of the different modes of consciousness and their phenomenology?

It is tempting to think of modes of consciousness as different overall ways of minimizing prediction error in a brain in a certain situation. That is, the brain may minimize prediction error at different rates, and may realize the same rate of prediction error minimization in different ways – and the overall rate of prediction error minimization may be reflected in the current mode. This perspective would matter because it ties different modes together under one theoretical umbrella, having the potential to explain why such different states (when explained in terms of content and functionality dimensions) all have in common that they are modes of consciousness. Of course, with this comes a rather speculative notion that the state of consciousness is determined by the overall rate of prediction error minimization: if you minimize too little, or too much, at any given time and relative to a model, then you lose consciousness. Don't ask me to justify this theory just yet though!



---

## References

- Bayne, T. (2012). *The Unity of Consciousness*. Oxford: Oxford University Press.
- Bayne, T., Hohwy, J. (2016). Modes of consciousness. In *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Sinnott-Armstrong, W. (Ed.) Oxford: Oxford University Press, 57–80.
- Bayne, T., Hohwy, J., & Owen, A. M. (2016). Are There Levels of Consciousness?. *Trends in Cognitive Sciences*, 20(6), 405–413.
- Dokic, J. (2012). Seeds of self-knowledge: noetic feelings and metacognition. In M. Beran, J. Brandl, J. Perner, & J. Proust (Eds.), *Foundations of Metacognition* (302–321). Oxford: Oxford University Press.
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- Hung, C. S., Sarasso, S., Ferrarelli, F., Riedner, B., Ghilardi, M. F., Cirelli, C., & Tononi, G. (2013). Local experience-dependent changes in the wake EEG after prolonged wakefulness. *Sleep*, 36(1), 59–72.
- Owen, A. M., Schiff, N. D., & Laureys, S. (2009). A new era of coma and consciousness science. *Progress in brain research*, 177, 399–411.
- Sarasso, S., Pigorini, A., Proserpio, P., Gibbs, S. A., Massimini, M., & Nobili, L. (2014). Fluid boundaries between wake and sleep: experimental evidence from Stereo-EEG recordings. *Arch Ital Biol*, 152(2e3), 169e77.
- Sergent, C., et al. (2017). Multidimensional cognitive evaluation of patients with disorders of consciousness using EEG: A proof of concept study. *NeuroImage: Clinical*, 13, 455–469.
- Vyazovskiy, V. V., Olcese, U., Hanlon, E. C., Nir, Y., Cirelli, C., & Tononi, G. (2011). Local sleep in awake rats. *Nature*, 472(7344), 443–447.



# The anthropology of mind

## Exploring unusual sensations and spiritual experiences across cultures

An interview with  
Tanya Luhrmann

By Martin Fortier

Citation: Luhrmann, T., & Fortier, M. (2017). The anthropology of mind: exploring unusual sensations and spiritual experiences across cultures. An interview with Tanya Luhrmann. *ALIUS Bulletin*, 1, 25-36.

**Tanya Luhrmann**

[luhrmann@stanford.edu](mailto:luhrmann@stanford.edu)

Department of Anthropology  
Stanford University, USA

**Martin Fortier**

[martin.fortier@ens.fr](mailto:martin.fortier@ens.fr)

Institut Jean Nicod  
ENS/EHESS, Paris, France

To begin, could you please tell us a few words as to how you position yourself within the academic field? In *Persuasions of the Witch's Craft* (Luhrmann, 1991) you advocated the psychological anthropology approach (1991, pp. 15 et sq.) while at the time acknowledging the importance of cognitive science (1991, pp. 13–14). Since then, you have constantly been reading and exchanging with scholars championing very different approaches—including psychological anthropology, phenomenological anthropology, cognitive anthropology, philosophy, evolutionary psychology, psychiatry, religious studies, developmental psychology, cultural psychology, neuroscience, etc. Besides, you are currently the Principal Investigator of a large interdisciplinary and cross-cultural research project where anthropologists are intensely collaborating with psychologists. How would you define your work today? What are the lines of research you find particularly inspiring? What are your likes and dislikes?

I define myself as an anthropologist with deep interests in psychology, and a commitment to using psychological methods to explore the questions ethnography cannot answer. The broadest question that interests me is how the world becomes real for people, particularly when what is real for them seems unreal to others. I'm interested in the texture of reality, and the way it changes for people, and in the quality of individual experience. I want to know about the differences between people, and the way that individual experience shifts and slides. The clearest way to see these moments of difference are in the edges of experience: in voices, visions, the world of the supernatural and the world of psychosis. What drives those experiences has a lot to do with the way people make judgements about what we call mental events. You might call me an anthropologist of mind.

While studying magicians in contemporary England (Luhrmann, 1991) or evangelicals in the US (Luhrmann, 2012), it seems that participant observation proved instrumental

for you. Unlike many anthropologists who prefer not to say much about their fieldwork experience, you have written quite openly about the anomalous experiences that you had during your fieldwork (Lurhmann, 1991, p. 319; Lurhmann, 2012, pp. 191–192). According to you, what status should be given to participant observation within anthropology? How native should anthropologists go in order to properly understand their object of investigation? Do you consider first-hand experience to be critical for the study and understanding of altered states of consciousness?

I don't think that first-hand experience is essential for understanding human experience—that leads us down the Winch/McIntyre rabbit hole and to the question of whether anything—God, for sure, being American, even owning cattle—can be understood from the outside. And yet there is no question that to have an insider's experience gives you a level of insight that you cannot have any other way. I remember one of my Cambridge supervisors, Stephen Hugh-Jones, telling me he'd gone off to do fieldwork in Amazonia at a time when there was much discussion about why people believed in their gods, and about the more abstract topic of social representation. He took ayahuasca—and he saw spirits. (First, to be clear, he saw London double decker buses.) That changed for him forever the answer to the question of why people believed in spirits. Even if you do not believe that what you see under the influence of ayahuasca is real, it changes your understanding of the event to have some grasp of its phenomenological quality. It adds subtlety to your explanation. In my own case, I had gone into the field anticipating that my explanation for why apparently rational people believed apparently irrational beliefs would be restricted to the realm of narrative and interpretation. When I had an anomalous experience myself—and when I realized that others had these experiences as well—I realized that talk of cognitive interpretation was not enough to capture what I had experienced. That changed the way I thought about religion and what a scholar should strive to explain.

“ I don't think that first-hand experience is essential for understanding human experience... yet there is no question that to have an insider's experience gives you a level of insight that you cannot have any other way. ”

In your work, you are especially interested in understanding the mechanisms through which imagination gradually becomes perception-like. In the case of evangelicals, for instance, you describe how they typically start by imagining God, and how, as they train themselves and pretend to be actually interacting with God, their experience becomes more and more real up to the point that they eventually hear God talking directly to

them. Understanding how God comes to be experienced as real is a very intricate question. In your research work, you explore several potential explanations:

(1) First, the shift from imagination to perception could be explained in metacognitive terms. For example, Richard Bentall has shown that schizophrenic hallucinations might be explained not by a change in the sensory content itself, but rather by a change in how the content is “tagged” by the mind: the same (more or less sensory-loaded) content can be tagged as “externally generated” – in which case it is experienced to be a real percept—or as “internally generated,” in which case it is experienced to be just imagination) (Bentall, 1990; Bentall, Baker, and Havers, 1991). Drawing upon this line of research, you suggest that the shift from imagination to perception could be mediated by a change in metacognition (Lurhmann, 2011a, pp. 72–73; Lurhmann, Nusbaum, and Thisted, 2013, pp. 171–172; Lurhmann et al., 2015b, p. 658).

(2) Second, you point out that the training evangelicals perform seems to affect the very content of their sensory experience, and not only how this content is being tagged. That is, the difference between the fictional God that is imagined and the real God that is encountered would be not just metacognitive but would also involve a properly sensory dimension. As you have explained, prayer practice “increases imagery vividness” and “lead[s] to reports of unusual sensory experiences and to reports of unusual sensory experiences associated with the religious ideas” (Lurhmann, Nusbaum, and Thisted, 2013, p. 172).

(3) Another proposal is that attention plays an important role in the shift from imagination to perception: “imagery rich practices may make what is imagined more real, not simply because increased attention leads to increased salience, but because the increased attention leads subjects to experience images as more “real”—more percept-like” (Lurhmann, Nusbaum, and Thisted, 2013, p. 161).

Now, it could be argued that these three strands of explanation are not as complementary as it may seem at first blush. If the shift from imagination to perception is explained by metacognitive processes (explanation (1)), this means that the experienced content is the same before and after spiritual training and what significantly changes is only how the content is tagged, as opposed to how the content becomes increasingly sensory-loaded. So, arguably, (1) and (2) are not intertwined. As regards explanation (3), attention could here be broadly interpreted as a kind of top-down process modulating access consciousness (Dehaene et al., 2006). One may then reason that through attentional training, a previously unconscious content could suddenly become conscious, which would explain why a rich sensory content is being reported (see (2)). Alternatively, the increase of sensory content could be construed as an instance of decreased sensory gating; as illustrated by the case of schizophrenic

patients, such a decrease typically leads to the experience of sensory overload (Micoulaud-Franchi et al., 2014). While (1) does not seem to fit so well with (2) or (3), on the other hand, it seems reasonable to surmise that (2) and (3) are possibly working complementarily.

What is your view on this debate? What do you think is instrumental in the shift from imagination to perception? Do you think that only one of the three aforementioned factors ((1), (2) and (3)) contributes to the shift or that all of them are required in order for the shift to occur? Finally, do you think that the shift from imagination to perception perfectly overlaps the shift from the sense of unreality to the sense of reality? Or do you think that these two pairs of concepts can be orthogonal, and that there is more to the sense of reality than simply experiencing something as being perception-like?

These are deep and complicated questions. I tend to assume that the metacognitive tagging changes the sensory qualia of the event post facto, through a micro-moment of attention. That is, the micro-moment decision to infer that the event (some string of words in the mind) is the memory of an event that took place in the world, rather than an event generated by the mind, shifts the experience of the event into a more sensory register. One remembers the event as more external and more sensory. That is more consistent with Richard Bentall's interpretation than Marcia Johnson's, but there is some suggestion in both sets of data that supports this interpretation. Meanwhile, Johnson's original work suggested that the more sensory content in the event under consideration, the more likely the event is to be interpreted as having a source in the external world and thus, more likely to have a sensory quality. So I see (3) as ultimate cause; and I do think that (1) and (2) work in concert. The increased sensory attention of prayer and absorption may lead people to infuse their events with more sensory information, and that in turn may lead to a greater likeliness of a judgment that the event had an external source and thus an experience with a richer sensory trace. I don't think that all three are necessary, but they can work together. I also do not assume that these are the only processes at work, and the only processes in play for any hallucination-like event.

At this point, the theoretical model I turn to explains how such phenomena might emerge out of ordinary cognitive process. Imagery and perception depend on many of the same neural structures, as Kosslyn among many others has shown. Increased attention to mental imagery should thus have some effects on a range of image-related cognitive processes: on perceptual processing, on the use of imagery, on unusual sensory experience, and on the vividness of imagery itself—as, indeed, my research has found. The individual trait of absorption (which seems to predispose people to having these unusual experiences may be capturing a similar attention to mental imagery, as many items seem to involve an interest in inner



imagery. Absorption is robustly and significantly correlated with the subjective experience of mental imagery vividness.

The puzzle in here, for me, is dissociation/hypnosis. There is a complex and poorly understood relationship between mental imagery vividness, absorption, hypnosis and dissociation. The absorption scale was developed as a pen-and-paper measure of hypnotizability, and while it correlates only modestly (if significantly) with the Stanford Hypnotic Susceptibility Scale, absorption is clearly related to hypnosis. Hypnosis practice increases imagery vividness, and intense spiritual practices can often be described as dissociation-inducing. There is already an active debate in which scholars have argued that most or all voice-hearing experiences are fundamentally related to dissociation due to past trauma. It may be that the pattern and pathway of voice-hearing for those with psychosis differs for those who dissociate and those who do not—regardless of a history of trauma. That should make us rethink some of our assumptions about psychotic hallucination.

I increasingly assume that the distinction between internal and external is a continuum rather than a binary. That is certainly what I hear from talking to people with psychosis and those who do not have psychosis but who have unusual experiences: people may say, I know it was not in my head, but I am not sure whether I heard it with my ears. I think certain kinds of events move up and down on the continuum with more ease than others. Indeed, that is the puzzle of the reality monitoring story (Marcia Johnson, Richard Bentall, Yoram Bilu and others)—that there are so few anomalies, as it were, in the way people judge the origin of events. I do not think we have a good account of the infrequency of hallucination-like events. My own sense is that we are constantly having somewhat chaotic events that we correct without being conscious of the corrections.

I do think that the sense of what is real is not merely to do with perception, but is another kind of process—related, of course, but not identical. One can see that simply through the fact that perception may be intact/normal while the sense of reality may be deeply disturbed, as in depersonalization or derealization. The sense of reality seems less a judgment about perception, and more a relationship to the act of perception. We look to Martin Fortier to explain our sense of reality in the years to come.

In your research, you have demonstrated that the shift from imagination to perception varies widely between individuals. More specifically, personal proclivity for absorption seems to be the driving force of an individual's propensity to have unusual experiences and to successfully transform imagination into percepts (Lurhmann, Nusbaum, and Thisted, 2010). This finding implies that religious experiences are underlain by training (for example, how much kataphatic praying one has performed) but that the

effectiveness of this training is modulated by idiosyncratic characteristics (i.e., one's score on the absorption scale).

According to you, what role has this dichotomy between cultural training (praying, meditation, fasting, etc.) and personal disposition (score of the absorption scale) played in the development and shaping of religion across cultures? For example, could it be that people rating high on the absorption scale were preferably selected in imagistic religions, as opposed to doctrinal ones (Whitehouse, 2000)? Do you think that, in some contexts, rituals or trainings are intense and strong enough to make the distinction between highly absorptive individuals and lowly absorptive ones less important (think, for instance, of shamanistic rituals involving the intake of a powerfully hallucinogenic substance)?

It has recently been posited that proclivity for absorption is largely underpinned by genetics (Ott et al., 2005). In your opinion, how important is this finding for scholars of religion?

Again, an excellent—and substantial—question. I used to think that membership in religious practice was unrelated to personal proclivity. After all, most religious practices have so many members (perhaps a quarter of all Americans are charismatic Christians) and people go to church for such varied reasons (proximity, spouse's preference, etc) that I assumed that differences in proclivity would wash out. I increasingly think proclivity plays some role in the choice of religion. I would expect that role to be more pronounced in smaller religions that require more effort to join.

“No great religion has been founded by someone without voices and visions of some sort, even though the predisposition to accept the plausibility of invisible others may be far more broadly distributed, part of the orientation of our evolved brains.”

Is proclivity inherently limiting? That is, should someone with low absorption give up an ambition to know God in a sensorially rich manner? Well, no. In my ethnographic work, I have seen people with low absorption, or perhaps better to say a low-absorption orientation to their world, develop the capacity to experience God vividly in a sensorial manner. I think that it is possible to train that style of attentiveness and engagement. At the same time: when I have seen this, I have also wondered about different kinds of proclivity which seem more akin to psychosis. The relationship between ordinary spiritual hallucination-like events and psychosis is fraught and contested (see the just published issue of *Schizophrenia Bulletin*). For that matter, the question of what processes are involved in psychosis is deeply



contested. We feel increasingly confident that psychosis is complex. I think there may be different processes involved, and that some of those processes may also be involved in spiritual responsiveness.

Should it matter to scholars of religion that absorption has a genetic component? No great religion has been founded by someone without voices and visions of some sort, even though the predisposition to accept the plausibility of invisible others may be far more broadly distributed, part of the orientation of our evolved brains. And no one, I think, has ever assumed that the people with visions and voices powerful enough to persuade others of their truth were run of the mill individuals. And yet: to begin to figure out what sets highly religious people apart is deeply interesting and, I think, important.

The question of cognitive penetrability (i.e., the question of knowing whether high-level beliefs can affect low-level perceptual processes) has sparked much discussions (e.g., Zeimbekis and Raftopoulos, 2015). In your own work—both past and present (Lurhmann, 2012; Lurhmann, 2011a; Lurhmann, 2011b; Cassaniti and Lurhmann, 2014; Lurhmann et al., 2015a; Lurhmann et al., 2015b)—you argue that local models of the mind affect how people experience things. It is not clear, however, whether in doing so you endorse a strong cognitive penetrability thesis, since when you speak of local models of the mind you are referring not only to reflective theories about the mind but also to more procedural processes such as attention.

To clarify things a bit, three levels could be distinguished: (1) explicit theories of the mind (e.g., claiming that the mind is made of three components: reason, spirit and soul); (2) habits of the mind (e.g., being good at controlling one's own attention as opposed to being constantly mind-wandering); (3) experience and perception (e.g., hearing someone talking or feeling pain in the back). To illustrate how these levels interact with each other, we could say, for example, that in Buddhism, theories about the mind (the philosophical theory of the five aggregates) does not directly affect perception, whereas what does affect perception is the daily training of the mind (what I have just defined as level (2)). And yet, the training of the mind (i.e., level (2)) seems to be improved and even enhanced by possessing and understanding the theory of the five aggregates (i.e., level (1)). So, in this specific case, it could be tempting to say that level (1) influences level (2) and that level (2) influences level (3), but that (1) does not directly influence (3).

What is your view on cognitive penetrability? What does your ethnographic and experimental work suggest about the interplay between explicit theories of the mind, habits of the mind and experience?

My own orientation is to presume that there is indeed downward influence, both from theological orientation and from local cultural expectations, and perhaps upward influence from habits of mind to explicit theories. We now have a large project focused on exploring exactly this topic. That project sets out to understand how cultural variation in ideas about the mind shapes the way people seek and experience the supernatural. We hypothesize that different cultural understandings of the mind—specifically, how separate the mind is from the world, how important inner experience is held to be, and how real the imagination is held to be—shape the way people pay attention to and interpret events they deem supernatural. We propose that although belief in supernatural agents may build upon psychological biases in human cognition, faith is culturally constituted. We are working in five different countries: Ghana, China, Thailand, Vanuatu/Oceania and the US, examining four populations per country: urban charismatic Christian; rural charismatic Christian; urban non-Christian; rural non-Christian. We compare these four populations not only to have comparable groups but also to investigate the impact of charismatic Christianity and industrialization on the way people think about thinking and their experience of the supernatural.

“ Anthropologists have found that different faith interpretations about the believer’s mental state have significant consequences for the believer. ”

In my previous work, I have seen that persons new to a church that taught them that God spoke inside the mind learned to experience thoughts which they might once have treated as self-generated as other-generated (God-generated). I and my colleagues found that prayer practice associated with inner sense cultivation (deliberate attention to inner experience) led those praying to experience what they called God as more person-like and more present, to feel that their inner sensory world became more vivid, and to increase the likelihood that they would experience what they identified as God’s voice, visions, and other unusual sensations. The ethnographic literature makes it clear that many religious practices involve training the mind (in particular, cultivating the inner senses) and that training alters the mental experiences of the person trained. The active discussions about embodiment in the psychological literature suggest that these ethnographic findings are an expression of the way that different practices shape subjective experience.

Anthropologists have found that different faith interpretations about the believer’s mental state have significant consequences for the believer. In particular, there is some evidence that Christianity may change the way that converts understand thoughts and thinking. In an ethnographic analysis of recent converts to

Pentecostalism in Melanesia, Joel Robbins detailed a sharply increased sense of the social importance of internal states by converts who were unable to rid their minds of thoughts they felt to be sinful. Webb Keane argues in a study of Indonesian Christianity that the expectation of sincerity and inner purity significantly shifted Sumbanese experience of thought and language. Ben Purzycki and Rich Sosis found that different understandings of God's mind also have consequences for the way that people think about human minds.

Meanwhile, philosophers, psychologists, and anthropologists have long argued that education and industrialization affects the way people think about minds and mental process (as in, for instance, the work of Karl Popper, Michael Cole and Sylvia Scribner, Ernest Gellner, Robin Horton, and Charles Taylor).

Our plan is to conduct both ethnographic and psychological work in order to tease out the relationship between explicit theology, cultural habits of mind, and perception.

I do agree that level (1) does not influence level (3) directly.

One of your present interests deals with participation, sensuous being-in-the-world, and the feeling of presence. This constellation of themes is evocatively illustrated by, among others, David Abram's (1997) book on the phenomenological approach of sensuous presence and Eduardo Kohn's (2013) book on iconic perception of the world among Amazonian indigenous people. How does this line of research relate to those that you have been exploring so far? Do you think researchers in religious studies and in the cognitive science of religion could gain valuable insights from the rehabilitation of Lévy-Bruhl's concept of participation?

I do think that researchers could learn something from Levy-Bruhl: I think he is describing a state of being in the world which is far more common, and far more important to religion, than we often realize. Levy-Bruhl saw that what it is to be religious is to experience the world as responsive and full of meaning. In *How Natives Think*, Levy-Bruhl argued that the distinctive feature of the "primitive" mind was that such primitives experienced themselves as participating in the external world, and the external world as participating in their minds and bodies. Levy-Bruhl called such an orientation "mystical" and he described it as governed by "the law of participation" in which objects are "both themselves and other than themselves." At the end of his life, in the posthumous *Notebooks*, Levy-Bruhl abandoned the claim that so-called primitive minds were fundamentally different from those of Europeans. He abandoned the term 'prelogical' (1975 [1949]: 99) and began to write of participation as common to all people, different modes of thought rather than different minds. The mystical mode of thought was both affective and conceptual,

and had those features which he had attributed to participation all along: independence from ordinary space and time, logical contradictions (an object is both here and there), identity between objects and their arbitrary features (like hair cuttings and the person from whom they came), “the feeling of a contact, most often unforeseen, with a reality other than the reality given in the surrounding milieu.” He thought that the mystical mode intermixed with everyday thought continually in our minds. For him, the puzzle became, “How does it happen that these “mental habits” make themselves felt in certain circumstances and not in others?”

I believe that what the absorption scale captures is an interest in feeling sensorially engaged with a responsive world, and that this interest facilitates a sense that the world is alive, aware, intelligent, interested—that it contains, in short, an invisible other. In recent years, I have been developing an alternate scale to the Tellegen absorption scale (the original scale is under copyright protection, which makes it irritatingly difficult to use as a research instrument). Here are some new items that strongly correlate with the absorption scale:

“Sometimes the world seems intensely present to me”

“When I walk through a forest, I like to think that the trees are murmuring words of wisdom for me”

“I have the distinct sense of a wise watchful presence”

“When I hear the wave lap against the shore, I sometimes think of how much those waves might know”

---

## References

- Abram, D. (1997). *The Spell of the Sensuous: Perception and Language in a More-Than-Human World*. New York: Vintage Books.
- Bentall, R. (1990). The illusion of reality. A review and integration of psychological research on hallucinations. *Psychological Bulletin*, 107(1), 82–95.
- Bentall, R., Baker, G., & Havers, S. (1991). Reality monitoring and psychotic hallucinations. *The British Journal of Clinical Psychology / the British Psychological Society*, 30(3), 213–222.
- Cassaniti, J., & Lurhmann, T. (2014). The cultural kindling of spiritual experiences. *Current Anthropology*, 55(S10), S333–S343.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, 10(5), 204–211.
- Kohn, E. (2013). *How forests think: Toward an anthropology beyond the human*. Berkeley/Los Angeles: University of California Press.
- Lurhmann, T. (1991). *Persuasions of the Witch's Craft: Ritual Magic in Contemporary England*. Cambridge MA: Harvard University Press.
- Lurhmann, T. (2011a). Hallucinations and sensory overrides. *Annual Review of Anthropology*, 40(1), 71–85.
- Lurhmann, T. (2011b). Toward an anthropological theory of mind. *Suomen Antropologi: Journal of the Finnish Anthropological Society*, 36(4), 5–69.
- Lurhmann, T. (2012). *When God talks back: Understanding the American evangelical relationship with God*. New York: Alfred Knopf.
- Lurhmann, T., Nusbaum, H., & Thisted, R. (2010). The Absorption Hypothesis: Learning to Hear God in Evangelical Christianity. *American Anthropologist*, 112(1), 66–78.
- Lurhmann, T., Nusbaum, H., & Thisted, R. (2013). Lord, Teach Us to Pray : Prayer Practice Affects Cognitive Processing. *Journal of Cognition and Culture*, 13(1-2), 159–177.
- Lurhmann, T., Padmavati, R., Tharoor, H., & Osei, A. (2015a). Differences in voice-hearing experiences of people with psychosis in the USA, India and Ghana: interview-based study. *The British Journal of Psychiatry*, 206(1), 41–44.
- Lurhmann, T., Padmavati, R., Tharoor, H., & Osei, A. (2015b). Hearing Voices in Different Cultures: A Social Kindling Hypothesis. *Topics in Cognitive Science*, 7(4), 646–663.
- Micoulaud-Franchi, J.-A., Hetrick, W., Aramaki, M., Bolbecker, A., Boyer, L., Ystad, S., ... Vion-Dury, J. (2014). Do schizophrenia patients with low P50-suppression report more perceptual anomalies with the sensory gating inventory? *Schizophrenia Research*, 157(1-3), 157–162.

- 
- Ott, U., Reuter, M., Hennig, J., & Vaitl, D. (2005). Evidence for a common biological basis of the absorption trait, hallucinogen effects, and positive symptoms: Epistasis between 5HT<sub>2a</sub> and COMT polymorphisms. *American Journal of Medical Genetics Part B (Neuropsychiatric Genetics)*, 137B(1), 29–32.
- Whitehouse, H. (2000). *Arguments and Icons : Divergent Modes of Religiosity*. Oxford: Oxford University Press.
- Zeimbekis, J., & Raftopoulos, A. (Eds.). (2015). *The cognitive penetrability of perception: New philosophical perspectives*. London/New York: Oxford University Press.

# The phenomenon of voice-hearing

## An interdisciplinary approach

An interview with  
Simon McCarthy-Jones

By Mathieu Frerejouan

Citation: McCarthy-Jones, S., & Frerejouan, M. (2017). The phenomenon of voice-hearing: an interdisciplinary approach. An interview with Simon McCarthy-Jones. *ALIUS Bulletin*, 1, 37-45

**Simon McCarthy-Jones**

[mccartsi@tcd.ie](mailto:mccartsi@tcd.ie)

Department of Psychiatry  
Trinity College, Dublin, Ireland

**Mathieu Frerejouan**

[mathieu.frerejouan-du-saint@univ-paris1.fr](mailto:mathieu.frerejouan-du-saint@univ-paris1.fr)

Department of Philosophy  
Pantheon-Sorbonne University, Paris, France

First of all, you usually define your object of study as “voice-hearing” rather than “auditory verbal hallucinations.” Can you tell us why this distinction is essential for you?

The distinction is important to me because it is important to a lot of people who hear voices. For someone new to the field, this might seem like a fairly dry terminological debate. It is not. Many have argued that it is part of a wider problem of dispossession, imposition, and colonization. This problem is set out well in a paper entitled “*Reclaiming Experience*” by Jacqui Dillon and Rufus May (2003), both of whom have been the recipients of psychiatric treatment. In this, they stress their right to “define ourselves” and to “find our own voice to describe our experiences and our lives”. You can see the potential benefit of this in something my colleague Amanda Waegeli (2013) wrote about her own voice-hearing experiences:

Psychiatry, professionals and academics... will all put their own interpretations on my experience and explain it in whatever way they like to, but the bottom line is it doesn't matter to me anymore now what they think. I accept my voice hearing experiences as being normal for me. I once wanted to know what they thought, and needed to know what they thought or diagnosed it as, because I thought they were the experts and I had something wrong with me and needed their knowledge to help me with my problem. Now... I am empowered on my recovery.

The term “auditory verbal hallucination” is often an entry point to these problems. It was created by people who didn't have the experience and then laid over the experiences of other people. In many cases this dispossessed voice-hearers of their experiences. For example, the term ‘hallucination’ automatically invalidates any spiritual explanation a person may have for their experience. It has a tendency to be the end of a conversation, rather than the start of one. Ontological implications aside, it also comes with a lot of other implicit baggage, including the intimation



that the experience is best seen through a biomedical lens and should be treated by psychiatric services. By using this term with people, it potentially imposes these ways of understanding. As such, the use of the term could be viewed as a colonial act.

We must take a balanced approach to this question though, and it should be stressed that some people find the term ‘auditory verbal hallucination’ a useful way to describe their experiences, and have no qualms about the term. Indeed, some people will actually find it empowering. For example, if someone has come to believe their voice really is the devil, it can be a relief to come to believe that it is not.

The term is also problematic for other reasons, such as when used to describe the experiences of historical figures (like Socrates or Joan of Arc) who would never have used this term. Yet the term can serve a useful communicative purpose. When you say ‘hallucination’ people know what you are talking about. Other terms, such as ‘locutions’ are less accessible. I briefly experimented with employing the strategy utilised by Jacques Derrida when a word was necessary, yet inaccurate. This involved placing it *sous rature* (under erasure). However, writing ~~hallucination~~ page after page isn’t the most aesthetically pleasing thing to do, and eventually you habituate to it anyway.

“ Professionals need to start from where the hearer is, not where the professional is, and there needs to be a process of negotiation, not imposition, when talking about these experiences. ”

As a result of such issues, service user-led organisations (such as the Hearing Voices Movement) have advocated the use of terms such as ‘hearing voices’, and ‘voice-hearing’. I am very happy to respect this nomenclature. My view is that (non voice-hearing) professionals need to start from where the hearer is, not where the professional is, and there needs to be a process of negotiation, not imposition, when talking about these experiences. Yet, the term ‘hearing voices’ still remains associated with madness in our culture. The challenge seems to be for voice-hearers to reclaim the term, in the way that the LGBTQI community have reclaimed words such as ‘gay’ and ‘queer’.

This phenomenon of voice-hearing has given birth to many different psychological explanations. However, since your first studies made in collaboration with Charles Fernyhough, you showed a clear interest for models of “inner-speech” connected to Vygotsky developmental psychology. How can these models enlighten the phenomenon of “voice-hearing”?



Charles is an expert on Vygotsky's writings and was an early proponent of the value of applying Vygotsky's work on inner speech to the experience of voice-hearing. Incidentally he has just published a great book on inner speech and voice-hearing (*The Voices Within*) which I would highly recommend. Some of my work with Charles attempted to further develop the inner speech model of voice-hearing. The basic idea underpinning this has been around for a long time, in part because it is a reasonably intuitive way to try and explain the experience. For example, back in the 1500s the Spanish mystic St John of the Cross railed against people who claimed to hear God's voice as, in his view, they were actually just "saying these things to themselves". In the 1800s, the French psychologist Eggers was arguing that voices were simply inner speech asserting itself with greater insistence than normal. Going beyond assertions, how can we test this hypothesis?

Some studies have found that when people hear voices there is a small but detectable amount of muscle activity in their throat and lips. Other studies have been able to amplify these signals so that external observers could hear what this silent speech was saying. Sure enough, what the experimenters heard were the same words that the voice-hearer reported their voices were saying. Of course, these are only small studies and we can't be sure that we can generalize the results to all people who hear voices. Going forward we might be able to make use of NASA's development of technologies that can decode your inner speech from your neural activity. Their interest in this area stemmed from wanting to allow astronauts to be able to control machinery, such as the Mars Rover, using only their inner speech: "turn left", "stop", "stay on target". Their technology has now advanced to the point that sensors attached to your throat can allow your inner speech to appear in front of you on a computer screen as you think it. This has a range of potential implications for voice-hearing. It could be used to get voices to write down what they say. This could allow the hearer to engage with what their voices say 'offline', giving them a bit more distance from their voices, and making their voices easier to understand and work with. From a linguist's point of view it would be a great source of material to analyse to try and better understand the voice-hearing experience. Some hearers I know have already been dictated stories by their voices. This technology could take things one step further; people's voices could directly write out poetry or stories. It could be a royal road to creativity.

If voices are based in people's own inner speech, then this can be used to create short- and long-term therapeutic interventions. Short term interventions focus on getting people to block their inner speech by doing things such as humming or opening their mouths wide. Obviously though, one can't go through one's life doing these things constantly. This raises the question as to why people who hear distressing voices should be speaking to themselves in inner speech so negatively. A

lot of the time this may relate to earlier traumas, although there are neurological models focusing on right Broca's area that offer alternative reasons. To me it seems that shame is a central emotion driving much voice-hearing. As such, getting people to challenge shaming inner dialogues or to create a more compassionate inner dialogue using techniques such as Compassion Focussed Therapy (which Charles Heriot-Maitland and Eleanor Longden are doing some great work with at the moment), could have the ability to make nasty voices into nicer voices, making life more manageable.

Your main work on “voice-hearing” includes a historical perspective, which is relatively unusual in cognitive sciences. How can a historical approach contribute to the study of “voice-hearing” in psychology?

I think there are three ways it can help. First, it can help inform our study of the phenomenology of voice-hearing. The writings of historical figures can help highlight aspects of the phenomenology of voice-hearing that we are not currently paying attention to. We then need to see if such experiences are occurring today and, if so, to update our models so they can account for these experiences too. The experience reported by many historical figures of ‘soundless voices’, where information is received, but not in explicit verbal form, is a good example of this.

Second, history can help us understand what aspects of voices are malleable in the hands of culture, and what features are essentially constant. For example, there was an interesting study that compared voice-hearing in patients admitted to a hospital in Texas the 1930s with voice-hearing in patients admitted to the same hospital in the 1980s. This found that both sets of patients heard commands, but the commands of the 1930s voices were mainly benign and religious (e.g., “live right”, “lean on the Lord”) whereas those of the 1980s were negative and destructive (e.g., “kill yourself”, “kill your mother”). The more negative commands of the later period could have reflected a more negative and hostile social environment. Unfortunately, we lack systematic data on how voices have changed over the decades in the West to be able to test this idea. The stability of commanding voice through the centuries demands some form of explanation. This is one of the strengths of the inner speech model we just mentioned. Vygotsky argued that one reason we develop inner speech is to be able to control ourselves. If we have a train track of silent instruction developmentally seared into our brains, should we be surprised that when voices occur they run on these rails?

The third way that history can help us is to give us insights into why we think about voice-hearing the way we do today. When we start to study voice-hearing, we don't start with a blank slate. There are centuries of culture pressing down on us, contorting us, and making us think about voice-hearing in a particular way. Once

we see how history has created a ‘common sense’ way of thinking about voice-hearing, we can step outside of this to consider whether this really is the best way to approach the problem. Probably the easiest example to discuss is why biological explanations for hallucinations dominate the literature today. There are obviously a range of factors causing this, but we can illustrate a few prominent ones.

The standard model of history is that everyone thought voice-hearing was due to spirits, until the birth of psychiatry put forward a biomedical model. This isn’t the case. Religion repeatedly promoted a biomedical model of voice-hearing for its own ends, long before psychiatry appeared on the scene. For example, during the English Civil War, when the controlling structures of society started to fall apart, a lot of voice-hearers popped up claiming to hear the voice of God. The Church of England was threatened by this because what ‘God’ said was often not what the Church wanted God to say. So, the Church pushed for the medicalization of the voice-hearing experience, allowing them to dismiss these voice-hearers as being ill. This medicalization was given further impetus at the birth of psychiatry, not mainly for scientific reasons (there was no good evidence of brain changes associated with voice-hearing at this time) but principally for political ones. Andrew Scull (2006) has argued that early psychiatrists were motivated to explain voice-hearing in a biological manner, because this was a good way to establish that they, as medical doctors, rather than priests were the best people to treat this. In the 20<sup>th</sup> century, the backlash against psychoanalysis, as well as the profitability of antipsychotic drugs, were further reasons pushing a biological understanding of voice-hearing. So we have prophet-bashing, profit-taking, psychoanalyst-denigration, and psychiatrist-promotion as contributors to the dominance of a medical model, none of which are good reasons for taking this view. Now, obviously, there are some very good reasons to take a biomedical view: antipsychotics do help some people who hear voices, and there are biological changes associated with voice-hearing. We need to make sure we adopt such views for these kinds of justified reasons, and even then we still need to be aware of what continued pressures may be on us.

To illustrate this last point, an awareness of the historical factors that have acted to downplay trauma-based models of voice-hearing (particularly the backlash against psychoanalysis and R.D. Laing), and to promote decontextualized medical models, can help us to reflexively ensure that we are fair-handed in our treatment of competing models of voice-hearing. For example, models of voice-hearing which foreground trauma are held to a much higher level of proof than those that foreground biology. Let me give you a quick example of what I am referring to here. The proposal that child abuse causes voice-hearing has been criticized on the grounds that it is possible that an evocative gene-environment correlation could lead to an illusory (non-causal) relation between such abuse and voice-hearing. The

argument runs that you would have genes for schizophrenia that lead to early developmental problems that in turn make you more vulnerable to child abuse. You then suffer such abuse and go on to develop the voice-hearing that your schizophrenia genes (not the abuse) had destined you to. This would then give the illusion that the child abuse had caused the voice-hearing, when both were to some extent actually evoked by genes (I would prefer to say that the genes made you more vulnerable to abuse, not that they evoked the abuse, so that responsibility and blame remains with the perpetrator). The point here is not the validity of this argument, or the details of it, but rather that this argument is being made at all. This is a pretty speculative argument to make against a trauma-based model of voice-hearing. Biological models of voice-hearing, such that voice-hearing is caused by altered neural connectivity, or that dopamine transmission abnormalities cause voice-hearing, are not subject to anything like this level of “let’s think of every possible reason why this theory may be wrong”-style of critique, despite it being pretty easy to do so based on the state of the literature in this area. This is not to say that trauma models should not be evaluated as hard as this. This is what science does, it is not personal. *All* models of voice-hearing should be as rigorously evaluated as possible.

“ History can help us understand why we think about voice-hearing as we do, and then allow us to hop off this ideological horse. ”

In such ways, history can help us understand why we think about voice-hearing as we do, and then allow us to hop off this ideological horse (whatever the horse be called; ‘trauma’, ‘biology’, etc.). We can then look to see if we have been systematically biased into looking at voice-hearing in a certain way, and then to ask if this is the best way to do things. We desperately need to try and eliminate the biases in our thinking, side-line professional interests, and see clearly.

As you just said, it’s necessary to reflect on the role of distal causes, such as trauma, in the explanation of psychological phenomena. However, do you consider distal causes as a simple addition to neuroscientific work, or as something that might deeply change the study of voice-hearing?

Trauma is no simple addition to neuroscientific work. The results of neuroscientific studies may be erroneously interpreted if they are not understood in the context of the person’s life. John Read and colleagues’ traumagenic model of psychosis makes this clear (Read et al., 2001). The decontextualised biological study of voice-hearing has the potential to confuse and harm as well as enlighten and help. For example, let’s say a study finds altered functional connectivity of the amygdala to be associated with voice-hearing. Suddenly, this becomes the core cause of voice-hearing, and

people will start to propose the use of neurostimulation techniques to attempt to rectify this aberrant connectivity. But what if this is the wrong level of analysis to understand things at? Imagine if hypervigilance for a genuine threat is at the core of an experience of voice-hearing. What we find at the neural level would simply reflect this. Thus, rather than see a genuinely threatening world as the source of the person's voice-hearing, we come to think that it originates from inside. This has practical implications. If, in this case, you were to go ahead and treat with neurostimulation, it may merely cause a brief and illusory period where the feeling of threat reduces. When the person goes back into their threatening environment, the problem will simply return. If the problem is in the world, we cannot get away with just treating brains. At some point, we need to change the world.

“ The decontextualised biological study of voice-hearing has the potential to confuse and harm as well as enlighten and help. ”

Again, as I am always very keen to seek balance, I should highlight that there are cases where biology is a very helpful level to understand someone's voice-hearing at. For example, metachromatic leukodystrophy is a rare white matter disease, most commonly caused by a mutation in the arylsulfatase A gene. This mainly results in damage to the myelin in the frontal and temporal lobes of the brain. Up to half of people with metachromatic leukodystrophy with adolescent or early-adult onset experience voice-hearing. This is a dysmyelination disease; it involves problems with the normal formation of myelin. Other white matter diseases are demyelination diseases. In these, white matter develops normally, but becomes damaged later in life. Multiple Sclerosis (MS) is probably the best known example of a demyelination disease, and has its average age of onset in people's 30's. Although MS patients can report voice-hearing, it is much rarer than in people with dysmyelination diseases. This suggests we should look for the causes of (at least some) voice-hearing in early adulthood changes to myelin in the fronto-temporal regions of the brain, and highlights the value of using a biological-level analysis.

Finally, your work is clearly guided by an ethical goal, as shown by your interest in the “Voice Hearing Movement”. In what ways do you think that scientific research can contribute to the lives of “voice-hearers”?

Scientific research should contribute to the lives of voice-hearers in whatever way people who hear voices want it to. A recent study I was involved with (led by Adele de Jager, 2015) found that people hearing voices had one of two basic stances towards recovery. The first was “turning away”. Here, people generally noticed a turning point when they were prescribed medication that helped. This group tended to

accept a medical model of the experience, viewed their voices as being symptoms of an illness, and had a strong sense of wanting to put the experience behind them and get on with their lives. For people who want to view their voice-hearing in this way, scientific research into the psychopharmacological treatment of voice-hearing (such as research into new anti-inflammatory and re-myelinating interventions) and new forms of biological treatment such as neurostimulation and neurofeedback (building on basic biological research into the experience), may be of value.

We also found another group of people whose recovery stories involved a “turning toward” approach. This group were characterised by a tendency to turn to face problems, to actively engage with voices, to be curious about what voice-hearing meant, to test their beliefs about voices, and to change their relationships with voices. Here, scientific research needs to focus on the meaning of the voice-hearing experience in the context of the hearer’s life, the relation the hearer has with their voice, and how/if tools such as the Maastricht Interview for Voice Hearing can be effective. There needs to be a lot more of this type of research funded.

We are lucky to live in a time when there is a great deal of wonderful collaborative research being undertaken across the world between people with and without lived experience of voices, each bringing specific skills to the table, shaping the questions asked, and deciding on the outcomes desired. Just to name a few, this includes work on the Maastricht Interview being led by people who hear voices themselves; the *Hearing the Voice* project based in Durham in the UK; work into peer-delivered interventions by Neil Thomas and his collaborators in Australia; and the multifaceted work of Nev Jones in the USA. It takes many voices to understand a voice, and I think we are finally getting somewhere together.

---

## References

- Dillon, J., & May, R. (2003). Reclaiming experience. *Clinical Psychology*, 17, 25-28
- Jager, A. de, Rhodes, P., Beavan, V., Holmes, D., McCabe, K., Thomas, N., ... Hayward, M. (2015). Investigating the lived experience of recovery in people who hear voices. *Qualitative Health Research*, 26(10).
- McCarthy-Jones, S., Waegeli, A., & Watkins, J. (2013). Spirituality and hearing voices: considering the relation. *Psychosis*, 5(3), 247–258.
- Read, J., Perry, B. D., Moskowitz, A., & Connolly, J. (2001). The contribution of early traumatic events to schizophrenia in some patients: a traumagenic neurodevelopmental model. *Psychiatry*, 64(4), 319–345.
- Scull, A. (2006). *The Insanity of Place / The Place of Insanity: Essays on the History of Psychiatry*. Routledge.





# Towards a biocultural approach of dissociative consciousness

An interview with  
Rebecca Seligman

By Arnaud Halloy

Citation: Seligman, R., & Halloy, A. (2017). Towards a biocultural approach of dissociative consciousness. An interview with Rebecca Seligman. *ALIUS Bulletin*, 1, 47-52.

**Rebecca Seligman**

[r-seligman@northwestern.edu](mailto:r-seligman@northwestern.edu)

Department of Anthropology  
Northwestern University, Evanston, USA

**Arnaud Halloy**

[halloy@unice.fr](mailto:halloy@unice.fr)

Department of Anthropology  
University of Nice Sophia Antipolis, France

For many anthropologists today, the Western category 'altered states of consciousness' (ASC) is still haunted by the ghost of pathologization. Some of them have instead proposed the use of culture-specific concepts such as 'hypnotic consciousness' (Obeyesekere, 1981), 'autonomous imagination' (Stephen, 1997) or 'shamanic state of consciousness' (Noll, 1983). They believe these terms may be more in sync with the experiential ground of possession, but at the same time, are more difficult to export as a general analytical category. Between these two attitudes, where would you drag your own epistemological cursor?

The problem with the term or category ASC is that it assumes a normative state that is widely shared, and from which certain states can be marked as alterations. The unmarked normative state that is implied is a self-conscious, reflexive form of awareness that is highly valued in Euro-American cultures. Of course, this state may not be the default state for all individuals in all places at all times, and this is where the element of pathologization comes from--states labeled as ASCs are understood to be non-normative even when they may not be. On the other hand, as you correctly point out, there is a utility in this label in the sense that it is widely known. Even Obeyesekere's attempt to replace this concept did not get a lot of traction, so I don't think that inventing my own term or using one of these lesser known labels would have gotten me very far. That's why in my own work, I have generally taken the approach that I will get more traction by using the term that is more widely recognized, and by defining it very carefully so that it is clear what I mean and don't mean. When I use the term ASC I try to be clear that I mean non self-conscious, non-reflexive states of consciousness. Such states may or may not be « altered » or divergent from the state that an individual or group understands to be normative – I think that is the crucial distinction that needs to be made. To be honest, in my

own work I more often use the term « dissociation, » which of course has strongly pathological implications as well. Much of my work has been devoted to understanding how and why dissociative states become pathologized/pathological.

You have spent many years studying possession in an Afro-Brazilian possession cult in Bahia. As an anthropologist, would you say that “possession” corresponds to the category of ASC? If not, how would you categorize such a phenomenon and experience? If so, how would you define the “building-blocks” (Taves, 2009) of the possession phenomenon and experience? In other words, what possession is made of?

If we use ASC to mean a non-reflexive state of awareness, as I suggested earlier, then I think what we call possession (phenomena in which individuals experience their self-consciousness to be displaced by that of a powerful “other”) can be said to fall into this category. But by definition the categories ASC and dissociation are centered around particular *states*, and the states they are concerned with are understood to consist of particular forms of awareness, attention, and physiology. As such, these categories may narrow our focus in a way that fails to capture the complexity of possession. My ethnographic work shows that possession is as much about particular ontological perspectives and social relationships as it is about states of consciousness. An ASC or dissociation focus may elide these other aspects of possession. For example, possession has sometimes been understood as a cultural explanation for, or interpretation of, trance, which is in turn understood as a physiologically based ASC. By this understanding, possession is a kind of post hoc cultural explanation or interpretation of a state that is explainable in non-sociocultural terms. This way of understanding possession follows a kind of dualism that treats biology and culture as separable from one another. But my research strongly supports the idea that possession itself is at once physiological and cultural. Being possessed by an entity other than the self is an experience that emerges out of a set of cultural and social meanings and commitments that powerfully influence the state of consciousness of the possessed. In other words, my findings suggest that the beliefs and practices associated with possession contribute to *producing* an altered state. In broader terms, this would mean that physiological states are fundamentally shaped by psychocultural processes rather than prior to them. I have made a concerted effort in my work to complicate understandings of ASC and dissociation to reflect this non-dualistic view.

Why do we need anthropologists to study ASC? To what extent are ASC such as possession culturally informed?

As I said earlier, altered states are not themselves inherently pathological or normative, and they are not universal physiological states that exist divorced from the sociocultural meanings and practices that surround them. The experiential

quality of these states is a direct product of their meaning for the individual and the social group. So this is why we need anthropologists to study ASCs -- because they are deeply socially and culturally conditioned. Researchers from other disciplines interested in altered states often view these states as the product of invariant physiological processes that can be understood in a decontextualized way. For example, psychiatrists tend to view dissociation as a neurophysiological mechanism that is triggered functionally. While my research and the research of a number of other scholars supports the idea that there is a physiological component to these states, I argue that patterns of physiological response associated with ASCs can be understood as embodiments of cultural practices and beliefs. Such cultural beliefs and practices take advantage of, shape, and enhance physiological potentials, but the physiology by itself does not constitute a particular kind of dissociative state (pathological or non-pathological), or a specific form of ASC (possession, shamanic trance, etc.). In my work I refer to the recursive mutual influences among cultural meanings and practices, physiological states, and social roles as « looping » and use the term « bio-looping » in particular to try to capture the role of physiology as both a condition for and a result of possession.

### What are the costs and benefits of incorporating psychological and psychophysiological methods into ethnography for studying possession?

I actually have a chapter in my book devoted to just this question! I can say for sure that I don't think that every scholar studying possession ought to be using psychophysiological methods in their research. I made a choice to incorporate these methods because I was interested in a particular set of questions for which these methods made a lot of sense. I wanted to be able to make empirically based arguments about how possession is embodied, and using psychophysiology measurement ultimately allowed me to make inferences about the bidirectional relationship between bodily states and social and cultural meanings. But certainly the use of these methods has its challenges and even a down side. The methods I used were challenging to apply in a field setting, for one thing, for a variety of technical and logistical reasons. In addition, these methods seem at odds with the ethos of ethnography in some ways – they take people out of the flow of their daily lives and put the ethnographer in the position of clinical data collector, as opposed to a participant observer. So I think that use of this kind of method is only really merited when the research questions call for it.

“Using psychophysiology measurement ultimately allowed me to make inferences about the bidirectional relationship between bodily states and social and cultural meanings.”

On the other hand, I think that the more people who ask questions about how body, meaning and practice go together, and who combine psychophysiology methods with ethnography in order to answer them, the more data we will accumulate that can help us escape from dualistic engagements with possession and with ASCs more broadly. And on the bright side, the technologies are definitely getting better – more portable, easier to use, and less invasive -- which means that in the future, these methods may fit more seamlessly with traditional ethnography.

For many decades, eminent scholars and psychiatrists have associated possession with mental illness and/or demonic influence. In your book *Possessing Spirits and Healing Selves. Embodiment and Transformation in a Afro-Brazilian Religion*, you affirm exactly the opposite by showing the self-transformative and healing potential of possession in Candomblé. In a few words, how would you explain that possession might be a healing practice or might have some healing effects?

Whew – that is a tough question to answer in just a few words – it took me a whole book! Yes, my book is ultimately about the ways in which spirit possession serves to transform and heal many of the individuals who participate as mediums. Many of the people I worked with came to their religious role with a sense of affliction and emerged from their initiations with a sense of well-being. So there is something important going on there. What I try to show in the book is that possession scaffolds a process of self-transformation, and that this self-transformation can result in healing. On the surface, this argument is not so far off from arguments by other scholars about the process of symbolic healing. But I tried to really dig deeper, to say “how does this work?” And what I found was that spirit possession works to transform people in a way that is complex and multifaceted. It transforms at once an individual’s social self, her subjectivity (meaning her cognitions, affects, and sense of self in the world), and her embodied self – that is, her way of being and experiencing via her body. These transformations are mutually dependent and inseparable, though for heuristic purposes we tend to carve things up into different domains (i.e. the social, the bodily, etc.).

“ These bodily, cognitive, and social aspects of self-transformation are linked together through a looping process, in which they continually shape and reinforce one another. ”

The case studies in my book show that what we think of as cognitive dimensions of self, as represented through self-narrative, are transformed through engagement with the cosmology of Candomblé. By taking on the ontological commitments of the religion, people come to understand and perform themselves

quite differently. Ontology emerges through practice in Candomblé spirit possession, and the enactment of the religious worldview also serves to transform people at a bodily, experiential level. My psychophysiology data support this argument in the sense that mediumship is associated with particular patterns of bodily response. At the same time, becoming a medium facilitates a transformation of the social self as well, as individuals become members of a close-knit religious community and come to be understood within this community as important ritual figures. A crucial part of my argument is that these bodily, cognitive, and social aspects of self-transformation are linked together through a looping process, in which they continually shape and reinforce one another. As a result, there is a kind of emergent property to the transformations people experience – the transformations seem to be more than the sum of their parts. The mediums I worked with not only came to feel differently, act differently, and think differently – they came to *be* different. I argue that this kind of transformation constitutes healing when people move from being afflicted, to being spiritually empowered.

---

## References

- Noll, R. (1985). Mental Imagery Cultivation as a Cultural Phenomenon: The Role of Visions in Shamanism ». *Current Anthropology* (26), 443-461.
- Obeyesekere, G. (1981). *Medusa's Hair: An Essay on Personal Symbols and Religious Experience*. Chicago: University of Chicago Press.
- Seligman, R. 2014. *Possessing Spirits and Healing Selves. Embodiment and Transformation in an Afro-Brazilian Religion*. Palgrave MacMillan: New-York.
- Stephen, M. (1997). Cargo cults, Cultural Creativity and Autonomous Imagination. *Ethos*, 25(3), 333-358.
- Taves, A. 2009. *Religious experience reconsidered: A Building-Block Approach to the Study of Religion and Other Special Things*. Princeton University Press, Princeton.

# Relocating dreams on the conceptual map

## How the analysis of sleep and dreaming challenges our taxonomy of mental states

An interview with  
Jennifer Windt

By Alessio Bucci and Raphaël Millière

Citation: Windt, J., Bucci, A. & Millière, R. (2017). Relocating dreams on the conceptual map: how the analysis of sleep and dreaming challenges our taxonomy of mental states. An interview with Jennifer Windt. *ALIUS Bulletin*, 1, 53-78.

**Jennifer Windt**

[jennifer.windt@monash.edu](mailto:jennifer.windt@monash.edu)

Department of Philosophy  
Monash University, Melbourne, Australia

**Alessio Bucci**

[alessio.bucci@unito.it](mailto:alessio.bucci@unito.it)

Department of Philosophy  
University of Turin, Italy

**Raphaël Millière**

[raphael.milliere@philosophy.ox.ac.uk](mailto:raphael.milliere@philosophy.ox.ac.uk)

Faculty of Philosophy  
University of Oxford, UK

### How did you become interested in the topic of dreaming?

I first became interested in dreaming as an undergraduate. Thomas Metzinger was teaching a series of lectures and seminars in philosophy of mind. A lot of it went right over my head, but what I did understand blew my mind, and I was hooked. For the seminar, each of us had to pick one week's readings to present to the class. When Thomas introduced the topic of dreaming, he asked if any of us had ever had a lucid dream—that is, if any of us had ever noticed while dreaming that we were now dreaming. I had had some nightmares in which I was vaguely aware that this wasn't really happening, that this was just a dream. And even though I had never really given much thought to dreaming, something about that question piqued my interest, so I put up my hand for the topic.

And then I started reading. The reading I had to present was Owen Flanagan's article on dreams as the spandrels of sleep (Flanagan, 1995). It was very empirically based, so I then started reading Hobson's *Dreaming Brain*, which is still one of my favorite books on dreaming (Hobson, 1988) and the 2000 BBS collection on dreaming, which had target articles by Allan Hobson, Antti Revonsuo, Tore Nielsen, and Mark Solms (*Brain and Behavioral Sciences*, vol. 23, issue 6, 2000). This ended up being way too much to fit into a presentation—but it eventually led to a term paper, and then to my MA and PhD theses, which I then turned into a book. There was always so much more to learn about dreaming, so much more to read and write and investigate that I never felt I had finished with the topic.



Now I am beginning to look beyond dreaming to think more about how dreaming connects to waking mind wandering, but also to experience in dreamless sleep, sleep disorders, and states of consciousness more generally. But essentially, these are just new ways of thinking about the same basic problem. What fascinates me about dreaming is that it is perhaps the clearest example of an utterly private and elusive conscious state in which conscious cognitive processes have become largely detached from outward behavior, ongoing tasks, and the environment. Dreams are so private that we can't even remember our own dreams most of the time—each night's dreams are elusive even to ourselves, slipping out of our conscious memories as we wake up every morning. That elusiveness puzzles me on a theoretical level, but it also gives the topic an air of mystery similar to uncharted territories that keeps me coming back to it. I think a similar elusiveness characterizes waking mind wandering and daydreaming. Again, the extent of mind wandering is surprising exactly because for the most part, it happens in the background, around the edges of awareness and retrospective recall (Smallwood & Schooler, 2015). There is a lot of uncharted territory in our waking mental lives as well.

“ Dreaming is the clearest example of an utterly private and elusive conscious state in which conscious cognitive processes are largely detached from outward behavior. ”

Another factor that has kept me interested in dreaming is how theoretical problems intersect with personal experience. For me, an interest in the philosophy and science of dreaming came first, and only later and after a lot of reading did I become more interested in my own dreams—in keeping a dream diary, trying (unfortunately without much success) to induce lucid dreams, and so on. I have kept this personal interest in my own dreams largely separate from my theoretical work, but it has made me more aware of the ubiquity and importance of dreams and sleep in our everyday lives. As elusive and mysterious as dreams are, the fact that we all dream every night and that virtually everyone has some dream story or other to tell makes the topic more tangible than, say, some purely metaphysical problem about the relationship between mind and body. This first-person familiarity then draws people into the more philosophical questions as well—for example almost everyone has an opinion as to whether they dream in color or in black-and-white, and as Eric Schwitzgebel has shown, the fact that these opinions change over time leads directly to the skeptical question of how well we know our own dreams, but also conscious experience more generally (Schwitzgebel, 2002, 2011).

There was one other crucial factor—and this one had more to do with mentoring and research environment than with dreaming. Thomas Metzinger had a group, called the MIND Group ([fias.uni-frankfurt.de/mindgroup](http://fias.uni-frankfurt.de/mindgroup)), for advanced undergraduate and graduate students interested in philosophy of mind, psychology, and cognitive neuroscience. The group met twice a year for conferences with invited international speakers in Frankfurt, and aside from the public lectures, the meetings were by invitation only. Just after I finished my MA thesis, Thomas organized a meeting with Olaf Blanke, Allan Hobson, and Antti Revonsuo—all leaders in the fields of OBE (out-of-body experience) and dream research. Thomas encouraged me to present my thesis at the meeting, and I was both thrilled and terrified—looking back, it seems I was a nervous wreck for weeks. But the presentation ended up being fine, and afterwards I got to have these long conversations over dinner with my intellectual heroes. And that just fascinated me—that this was actually a conversation I could be part of, and that once I got beyond the nervousness, I could take a peek at Allan’s dream diary or joke with Antti about dream bizarreness. This very much felt like a live area of research, and the people working in it were not just brilliant and genuinely passionate about their work, but also extremely friendly, laid back, and open enough to play around with different ideas with us students. For the first time, I had the feeling that this was a research area I could contribute to—that I could maybe go from being a student to doing actual research.

Part of this experience had, I think, to do with the intellectual atmosphere that characterizes dream research. Dream research, including the philosophy of dreaming, was a marginalized topic for a long time, and the group of people involved in it is still quite small. In addition, the air of fringyness that surrounds the topic ensures that most people working in this area are just genuinely interested—you wouldn’t pursue this line of research for reasons of prestige. (Especially for laboratory research, it’s also too draining—nights spent in the sleep lab exact an enormous toll on researchers.) So while there is indeed a lot of weird stuff out there on dreaming, many of the serious scholars who produce high-quality research have retained a certain playfulness, creativity, and openness to different positions and approaches. In consequence, this area is extremely pleasant and inspiring to work in, and I have made some great friends.

But the other factor had to do with mentoring—with having a teacher like Thomas who didn’t just get me hooked on philosophy of mind and cognitive science, but encouraged me, even as a student, to participate in an actual conversation with the philosophers and researchers whose work I had been studying. This was incredibly special about the MIND Group—and I know that for many other junior members, going to those meetings and presenting their work there was extremely motivating and a formative point in their research and career trajectories. I also saw this later when, as Thomas’ assistant and group manager of the MIND Group, I

encouraged my own students to participate in the meetings. I wish there were more opportunities to mentor and support students in that way—to create an environment where they don't feel that they are mere students, but where they can get a taste of what it's like to actually participate in an academic conversation with leading researchers, either at a conference or in print (as in the Open MIND collection, which I co-edited with Thomas: [open-mind.net](http://open-mind.net)). It really makes a huge difference, and it would be extremely beneficial to see more such opportunities out there.

How do you think the general conception of dreaming has changed over the course of the 20th century, in relation to physiological and neurological research on sleep mechanisms?

During the 20<sup>th</sup> century, dreaming went from being viewed as private and essentially unobservable to being established as a real phenomenon that could be investigated with the help of objective measures, such as polysomnographic data from the different sleep stages (Kroger, 2007). While the roots of scientific dream research reach into the 19<sup>th</sup> century (Schwartz, 2000), the pivotal moment came when William Dement and his colleagues discovered REM (rapid-eye-movement) sleep and its close association with dreaming in the 1950s (Dement & Kleitman, 1957). For the first time, the contrast between dreamful REM sleep and presumably dreamless NREM (or non-REM) sleep suggested that there might be an objective marker of dreaming over and above the subjective impression of having dreamt—and that this might make dreaming objectively diagnosable, turning it into a respectable and well-behaved target of scientific investigation.

The discovery of REM sleep also marked the beginning of scientific sleep research. Today, these two fields are largely separate, and dreaming plays at best a marginal role in important areas of sleep research—including research on memory consolidation in sleep and sleep disorders—but at least initially, both fields developed together. Importantly, the discovery of REM sleep, and more generally of a complex sleep architecture involving different stages of sleep associated with different levels and patterns of brain activity, profoundly changed scientific understanding of sleep. Sleep had traditionally been regarded as a period of uniform passivity and rest, but was now seen to require a more complex account. And at least initially, this changing conception of sleep was inextricably linked to the changing conception of dreaming. This raised not just empirical, but also profound conceptual questions. Was REM sleep/dreaming a third state of the brain, distinct both from sleep (now narrowed to NREM sleep) and wakefulness (Jouvet, 1999)? Or was it a state between sleep and wakefulness, of being half asleep, as it were? A similar conceptual uncertainty is implicit, today, in the question of whether lucid dreams are hybrid states, occurring on the border between REM sleep and

wakefulness, or genuine sleep states involving a substage of REM sleep (say, lucid REM sleep; Voss, Holzmann, Tuin, & Hobson, 2009; Windt & Voss, forthcoming).

“ How to reconcile the study of subjective experience with cognitive neuroscience remains one of the most pressing and unresolved challenges. ”

While researchers were initially enthusiastic about the prospects for a science of dreaming, in philosophy, these early studies on dreaming and REM sleep produced an almost allergic reaction. Norman Malcolm famously denied both that dreams are experiences and that something like a science of dreaming could exist, even in principle (Malcolm, 1962). His argument was based on purely conceptual considerations. Leaving all details to the side, it had two important parts. The first was that as dreams occur in sleep, and as sleep is by definition a state of unconsciousness, dreams cannot be experiences in anything like the sense in which waking experiences are. We use the same mental state terms to describe our waking experiences and our dreams, but this is a surface similarity only. If it seems to me, after awakening from a vivid nightmare, that I experienced intense fear in my dream, then I am just plain wrong: I did not, in sleep, experience anything at all. The second part of Malcolm's argument had to do with the impossibility of verifying dream reports and the absence of objective criteria for determining the occurrence of dreams. In Malcolm's view, the only way to determine whether and what someone has dreamt is the dream report given after awakening. Waking memory reports can, at least in principle, be checked and verified—but according to Malcolm, the very idea of verifying dream reports is absurd. If we introduced objective criteria for determining whether or not a person had dreamt, such as the presence vs absence of REM sleep, we would be changing the concept of dreaming. We would then no longer be talking about the same thing. So for Malcolm, dreams are neither subjective experiences nor targets for scientific research, whereas today, most researchers and philosophers think the opposite is true. Yet, how exactly to reconcile the study of dream experience, via the analysis of dream reports, with the study of sleep—and, more generally, the study of subjective experience with cognitive neuroscience—remains one of the most pressing and unresolved challenges in philosophy of mind and interdisciplinary consciousness research.

Today, the picture is much more complicated. In philosophy, the question of dream experience has been replaced by a number of more precise follow-up questions. For example, are dreams hallucinations occurring in sleep, or are they more like daydreaming and waking imagination? Does self-consciousness exist in dreams? Do dreams involve false beliefs and are they comparable to wake-state delusions? And so on (see Windt, 2015, 2016 for discussion). In this process, the

philosophical debate on dreaming has moved away from purely conceptual analysis and become much more interdisciplinary and empirically based. At the same time, in scientific dream research, it has become clear that not only do dreams occur outside of REM sleep (as well as REM sleep occasionally occurring without dreaming), but different kinds of conscious mentation distinct from dreaming occur even in the deep stages of NREM sleep (including slow-wave sleep: Windt, Nielsen, & Thompson, 2016). This last point is surprising because dreamless, deep sleep is often by definition thought to be unconscious. What is needed at this point is a taxonomy of both dreaming and dreamless sleep—including conscious mentation occurring in dreamless sleep—that is initially independent of sleep stages. Such a fine-grained taxonomy might then, in conjunction with more fine-grained sleep-stage scoring, allow for a more precise mapping of sleep stages. This type of project would be a huge step towards bringing sleep and dream research back together again. And it would necessitate yet another new shift in conceptions not just of sleep and dreaming, but of consciousness more generally.

### What are the trending topics in dream research at the moment?

To me, the most interesting trends in sleep and dream research converge on a similar theme. This is to develop a more fine-grained taxonomy for describing conscious experience in sleep alongside improved sleep-staging criteria.

Different lines of research are contributing to this development. In dream research, there is now increasing convergence, from different research groups and different disciplines, on simulation views (see Revonsuo, Tuominen, & Valli, 2015 for discussion and further references). The basic idea, here, is that dreaming is at core immersive: there is a here-and-now experience, a sense of being present in the dream world. And associated with this sense of presence is a representation of a self—the dream self, or the dream character the dreamer later identifies with—that is experienced at the center of the dream world. Simulation views allow for a lot of variance across different types of dreams (such as lucid dreams, nightmares, and so on). There are also different kinds of simulation views—for example, what exactly is involved in experiencing oneself as a self in dreams is open to debate. Still, in a field that was long characterized by considerable uncertainty and controversy as to how to define dreaming, increasing convergence on immersion and self- and world-simulation as defining features of dreaming is constructive and has the potential to unify different experimental and theoretical approaches.

Simulation views also suggest points of contact between dreaming and research on virtual reality and full-body illusions. In fact, dreaming is sometimes described as the gold standard of immersive virtual reality, and its investigation might help identify and empirically ground the conditions for presence and phenomenal

selfhood (Windt, 2015, chap. 12). Moreover, because simulation views offer a more precise definition of dreaming, they can also be used to develop a taxonomy for describing kinds of experience occurring during sleep that are distinct from dreaming. For example, instances of sleep thinking or even of visual or auditory imagery, where this occurs independently of an immersive hallucinatory context, would count as dreamless in this conception (Windt et al., 2016).

New evidence from sleep research is also chipping away at current sleep-stage scoring criteria. Findings suggest that sleep does not uniformly affect the whole brain, but can be unevenly distributed over the hemispheres (Tamaki, Bang, Watanabe, & Sasaki, 2016) or even occur in localized neuronal assemblies (Huber, Felice Ghilardi, Massimini, & Tononi, 2004). Similarly, research on sleep disorders suggests that parasomnias arising from NREM sleep may involve dissociations between sleep and wakefulness (Mahowald, Cramer Bornemann, & Schenck, 2011). These are just a few examples of how theoretical conceptions of sleep are becoming more complex and challenging existing classification systems. In this process, the distinction between wakefulness and sleep itself appears to be less clear-cut than often assumed.

Finally, there have also been important methodological advances involving neuroimaging during sleep, high-density EEG, non-invasive brain stimulation (for instance to induce lucidity during REM sleep; Voss & Hobson, 2014) and serial awakening paradigms (Noreika, Valli, Lahtela, & Revonsuo, 2009; Siclari, LaRocque, Bernardi, Postle, & Tononi, 2014; Siclari, Larocque, Postle, & Tononi, 2013), in which participants are awakened at very short intervals, thus maximizing the number of reports gathered per participant and per night.

These different developments can be combined to form powerful new research paradigms. By using novel methodologies to develop a more fine-grained taxonomy for describing the range of sleep-related experience alongside improved sleep-staging criteria, it might be possible to realign dream and sleep research. Much as was the case in the 1950s, the study of subjective experience in sleep might once more play a central role for sleep research as well, for instance for developing novel diagnostic and therapeutic measures for sleep disorders or helping to understand memory consolidation in sleep (Windt et al., 2016). This process will be inherently interdisciplinary, with theoretical-conceptual work from philosophy being informed by research findings and vice versa.

Another potential area of application is the search for the neural correlates of consciousness (NCC). To date, this research has largely been construed as a search for the neural correlates of specific contents of consciousness—such as seeing red or feeling pain. In this context, dream research can be used to investigate the neural correlates of specific contents of dream experience. For example, lucid dreamers can



use prearranged eye-movement patterns (e.g. looking left-right-left right in their dream) to signal that they have now become lucid and are engaging in a specific task, such as clenching a fist (Dresler et al., 2011). Because gaze shifts performed in lucid dreams correspond to the sleeping subject's actual eye movements, researchers in the lab can use these signals to investigate the associated pattern of brain activation. The data from lucid dreaming can then be compared to those from wakefulness, for instance, to actual fist clenching, but also to merely imagined fist clenching in waking participants.

By contrasting the presence vs absence of conscious experience in sleep, it might also be possible to move beyond the neural correlates of specific contents of experience to investigate the neural correlates of background states of consciousness (Noreika, 2014; Singer, 2014b, 2014a). Personally, I have come to think that dreamless sleep experience is the most interesting contrast condition in this context, because a subtype of dreamless sleep experience may involve a minimal form of phenomenal consciousness (Windt, 2015b; Windt et al., 2016). Investigating these kinds of dreamless sleep experience, in which the immersive, here-and-now structure that characterizes both dreaming and waking experience has been lost, can help identify and empirically ground the conditions for the simplest forms of conscious experience. If it turns out that minimal forms of dreamless sleep experience exist and can be systematically investigated in NREM and particularly slow-wave sleep, this would require a profound departure from current thinking both about the structure of phenomenal experience and its neural correlates. Standard characterizations of consciousness as what disappears in dreamless deep sleep would then require revision. I think a plausible argument can be made that this is indeed the case—and it will be very exciting to see where this research develops next.

**In your opinion, is dreaming an altered state of consciousness? If so, what defines it as such?**

Dreaming is clearly on the list of altered states of consciousness—at least according to folk psychology, dreaming is the most frequently occurring altered state, remarkable for its ubiquity and spontaneous occurrence as well as for its characteristic differences from standard wakefulness on the phenomenological, neuroscientific, and functional levels of description. But in trying to pinpoint how exactly dream experience is altered as compared to wakefulness and what defines altered states of consciousness in general, things get much more complicated.

The term altered states of consciousness, in my view, suggests an alteration not just in behavior and/or neural processing, but in phenomenal experience. But this means that dreaming, somewhat counterintuitively, does not clearly or necessarily



involve such an alteration. Recall that simulation views define dreaming by its immersive structure: the experience of a world centered on a self. On a purely phenomenological level of description, this here-and-now experience marks a deep similarity between dreaming and waking experience.

To be sure, there are differences on the neuroscientific and functional levels—neuroimaging studies suggest that brain activation patterns in REM sleep differ from wakefulness, and these differences are reflected, for instance, in the strongly visual and emotional character of dreaming and the frequency of movement sensations (Desseilles, Dang-Vu, Sterpenich, & Schwartz, 2011). The functional association between subjective experience on the one hand and environmental and real-body stimuli and real-body movements on the other hand is also much weaker and less predictable in dreaming than in wakefulness. Moreover, dreams typically misrepresent the sleeping subject's current location—only rarely, as in realistic false awakenings, do dreams mimic the sleeping subject's actual environment. These differences, however, don't necessarily show up on the phenomenological level of description. This comes back to the classical philosophical problem of dream skepticism: even in the face of bizarre events, and surroundings, dreaming quite often feels, subjectively, no different from being awake. And we can now see that this seeming resemblance between dreaming and wakefulness might have much to do with their common here-and-now structure. This, however, is just another way of saying that dreams are not altered states of consciousness with respect to what many now agree is their defining phenomenological feature.

To this, one might respond that there are a number of differences that typically distinguish the phenomenology of dreaming from waking experience. For instance, I think that we experience ourselves as embodied agents to a much weaker degree in dreams than in wakefulness, and this phenomenological difference is closely bound up with a weaker functional coupling between bodily experience and the physical body (Windt 2015, chap.s 7&8). But pointing to these ways in which dreams typically differ from standard waking experience will not give us a satisfying account of what it means to say that dreaming as such is an altered state of consciousness. To account for the exact kind of alteration involved, we would have to consider individual dreams on a case-by-case basis. And while I think this is the right way to go, this strategy cannot yield a general framework for distinguishing altered states of consciousness, including dreaming, from standard wakefulness.

A related problem is that speaking of altered states of consciousness involves an implicit comparison to a baseline. This alleged baseline of standard waking consciousness is, however, itself insufficiently understood and typically remains undefined. Both dreams and waking experience are heterogeneous and characterized by numerous fluctuations on the phenomenological, functional, and neuroscientific levels of description. Often, these fluctuations are subtle, hard to detect, and evade

any quick-and-easy, general characterization. I think that gaining a better understanding of these fluctuations in dreaming and wakefulness, along with a more fine-grained taxonomy, is an important goal for future research. But this also suggests that identifying a meaningful commonality that allows us to classify dreams as belonging to a broader category of altered states while also setting these apart from standard wakefulness may not be the most constructive way to go.

Rather than giving a categorical and necessarily coarse-grained account of what sets altered states, including dreaming, apart from baseline states of consciousness, I think research should move toward more fine-grained, multi-level classification systems able to capture fluctuations in experience across the sleep-wake cycle. This project will be applicable to many states that fall under the folk-psychological heading of altered states—including dreaming—but in this process, we might find deep continuities in experience across so-called altered and standard states, as well as, perhaps, genuine heterogeneity between and even within states commonly classified as either altered or baseline states. It might still be useful to develop a general account of what sets all of those states that in folk-psychology are commonly described as altered states apart from baseline or standard states of consciousness. I worry, however, that categorizing both wakefulness and dreams in the broad terms required for this type of project will result in an overly simplified and stereotyped view—and in many ways, this would be a move in the opposite direction from what I am proposing.

What are the main methods used in dream research to gather data about the first-person experience of dreaming? What role do dream reports play in this process, and what is their relation to objective, third-person data about sleep and the different sleep stages?

The study of so-called first-person data about dreaming is absolutely central to dream research, where this is understood in a very general sense as involving the study of conscious experience during sleep. In fact, I would go so far as to say that dream research is constrained, for methodological reasons, by data from the collection and analysis of dream reports. By this, I do not mean that dream research is exclusively about the study of dream reports—clearly, this would restrict the scope of the field quite drastically. While some areas of dream research are exclusively report-based, others try to relate data from dream reports to behavioral, polysomnographic, and/or neuroimaging data from the same sleep stage. Achieving this kind of one-one mapping of different types of data about the same experiential episode was at the center of the early laboratory studies of REM sleep/dreaming and continues to drive progress to this day.

“ Dream research is constrained, for methodological reasons, by data from the collection and analysis of dream reports. ”

In saying that dream research is constrained, for methodological reasons, by the space of reportable dreams—by what can and can't be reported about experience during sleep—some points need qualification. To begin with, the notion of first-person data is controversial. If we take first-person data to refer to introspective knowledge or inner observation of one's ongoing conscious states, it is unclear that we are still speaking of data in any interesting sense. Data are gathered with the help of measuring devices, and they are intersubjectively accessible and can in principle be verified and replicated (Metzinger, 2006). But none of this is true, it seems, for introspective knowledge (if it is knowledge) of ongoing conscious experience. When I use the term first-person data, I am using it in a different and fairly innocent sense to refer to data gathered from dream reports. Dream reports furnish the raw material for the science of dreaming—and dream reports, unlike the experiences they supposedly describe, are intersubjectively accessible. And while different research groups may disagree as to how best to analyze these data, there are more or less established scoring systems and statistical methods that can be applied to them. Through the collection and analysis of large sets of dream reports, researchers can then begin to investigate general questions about dreaming—for instance about the frequency of different types of emotions in dream reports as compared to waking reports (Sikka, Valli, Virta, & Revonsuo, 2014). The results of these studies are, in principle, replicable—even if the individual experiences are not.

When I speak of dream reports in this context, I do so in a very broad sense. Dream reports, in my view, are the results of behaviors conducted with the sincere intent of conveying or recording certain relevant information about a particular dream (Windt 2015, chap. 3). This can be done in many different ways—through written or spoken dream reports, by responding to specific questions in an interview with an experimenter or a questionnaire, or by using non-verbal media such as drawings or comparing one's dreams to photographs with different degrees of brightness or color saturation (Rechtschaffen & Buchignani, 1992). Dream reports may not even be necessarily retrospective. Signal-verified lucid dreams, in which lucid dreamers make prearranged patterns of eye movements to indicate that they are aware that they are now dreaming, are an example (Voss & Hobson, 2014). While the information conveyed by these eye movement signals is fairly coarse-grained, they can, I think, be described as a kind of concurrent behavioral report. Theoretically, this is extremely interesting, because it means that at least unidirectional communication from dreaming to wakefulness is possible. At the

same time, these concurrent reports cannot provide stand-alone evidence: to avoid false positives, it is crucial that the retrospective dream report later confirm that the dreamer was indeed lucid and made the eye movement signals deliberately.

The tricky question from a philosophical perspective, however, is whether dream reports as the primary source of data for scientific dream research are trustworthy with respect to the occurrence and phenomenal character of experience during sleep—whether dream reports reflect what it is actually like to dream. In philosophy, the trustworthiness of dream reports has long been doubted as a matter of principle. At its strongest, skepticism about dream reporting claims that no matter how much the methods for gathering and analyzing dream reports are improved, we still can't be sure that dream reports accurately describe whatever was or wasn't experienced during sleep (Dennett, 1976). My own view is that such a strong, principled kind of skepticism about the trustworthiness of dream reports, in which we can never be sure that dream reports provide evidence about conscious experience in sleep, is misguided (Windt, 2013, 2015a, chap.s 1&4). A more constructive and research-generating strategy is to start with the default assumption that at least a subset of dream reports are trustworthy—and then to use this subset as a baseline for further improving the conditions under which dreams are reported, along with training and more precise scoring systems, both for categorizing the range of sleep-related experience and scoring sleep stages.

What, then, are the ideal conditions for reporting dreams? Laboratory studies, in which timed awakenings can be used to gather reports from different sleep stages and minimize the temporal delay between the experience and the report, have long been considered the gold standard of dream research (see Windt 2015 for discussion and further references). By contrast, studies investigating spontaneous dream recall, in which participants sleep at home and write their dreams in a dream diary after waking up in the morning, offer less controlled conditions. At the same time, the laboratory situation itself may alter sleep quality as well as the content of dreams—it is not uncommon for participants to dream of the lab, the researchers, and so on. Further issues concern which participant groups to use (e.g. lucid vs nonlucid dreamers, participants with high vs low dream recall), to what extent training can improve dream recall or rather introduces bias, whether results from lucid dreams can be generalized to nonlucid ones, and so on.

All of these are methodological questions for dream research, and which type of report to use, as well as the optimal method and timing of awakening, the wording of questions etc., will depend on the research question asked in a given study. Shifts in methodologies can lead to interesting shifts in theoretical views. For example, changes in how best to report and rate the occurrence of dream emotions have led to different theoretical views on the frequency and kinds of emotions experienced in dreams (Sikka et al., 2014). For now, my main point is that if we start from a

default position in which at least some subset of dream reports are assumed to be trustworthy, it makes sense to say that scientific dream research, by developing new methods for gathering and analyzing dream reports and optimizing reporting conditions, can actually improve the trustworthiness of dream reports. By contrast, to deny that such an improvement is possible, even in principle, is to deny the possibility of scientific dream research in any meaningful sense. Dream research would then just be the study of dream reports, but it would not be clear that dream reports or the studies based on their analysis were at all informative about experience during sleep.

Might a future science of dreaming move beyond the study of dream reports entirely? Because of the elusive and unstable nature of dream recall, the dependency of dream research on the study of dream reports might seem to be a weakness. True progress, in this view, would involve moving beyond the study of dream reports to more objective and scientific kinds of data. For instance, we might envision future researchers predicting dream experience on the basis of polysomnographic and/or neuroimaging data alone. And while first steps are being taken in this direction—for example by using advanced machine-learning algorithms to decode patterns of brain activity that map onto the presence vs absence of dreaming (as in the dream catcher test proposed by Antti Revonsuo; Revonsuo, 2005), or onto different kinds of dream content (Horikawa, Tamaki, Miyawaki, & Kamitani, 2013)—it is important to keep in mind that the success of such predictions, ultimately, is measured by their correspondence to reported dreams. Even if it were the case that future researchers made predictions about the occurrence and content of dreams largely in the absence of data from dream reports, the attempt to match objective measures (such as neuroimaging data) to dream reports would have been instrumental in developing these methods in the first place. The strength of the resulting predictions would therefore still depend at least on their potential correspondence to reported dreams. This is why I think that the idea that dream research could move beyond the study of dream reports entirely is misleading (Windt, 2013). At least given the current state-of-the-art, studies that do not directly investigate dream reports but derive general claims about dreaming from purely behavioral, polysomnographic and/or neuroimaging data are best conceived of as sleep-only studies. They can identify meaningful and exciting future directions for dream research, but do not form part of scientific dream research proper.

Couldn't there be kinds of experience in sleep that are so subtle and fleeting that they are beyond the grasp of memory and below the threshold of reportability? I think that this is possible—perhaps even probable. Certain white dreams, in which participants say that they remember having dreamt but can't recall any details, might be an example of remembering a type of conscious experience that was so subtle and fleeting that nothing other than a vague impression of its occurrence can

be reported. In other white dream reports, the inability to give a more detailed description might be due to forgetting—and by investigating such reports in more detail it might, ultimately, be possible to tease apart different factors involved and distinguish different subtypes of white dream reports. Here we would have a case in which the potential limits of dream reporting would be investigated exactly through a careful analysis of reports (Windt, 2015b). Another example would be to use training and focus on specific participant groups. For instance, long-term meditators sometimes report witnessing dreamless sleep, and these reports might describe a minimal kind of dreamless sleep experience (Thompson, 2014). But again, this strategy would not move beyond dream reports, but would rather expand the space of reportable dreams and sleep-related experiences. By contrast, if experiences in sleep exist that cannot be rendered reportable even through the use of optimized methods and training, these experiences are beyond the reach of scientific dream research.

I think the relevance of this point is often under-appreciated: while the use of data from first-person reports—on dreaming, but also on conscious experience more generally—is often thought to hamper scientific progress and to be at odds with the requirements for an objective science of consciousness, I think that in fact, the systematic collection and analysis of report-based data is the condition of possibility of a science of dreaming, and of a science of consciousness more generally. A key challenge then becomes how to improve the methods for gathering and scoring these reports—for instance through questionnaires, training, and improved taxonomies for categorizing subtle and hard-to-describe differences in experience. This is an area, I think, where philosophy and cognitive neuroscience can constructively complement each other.

In your recent book (Windt, 2015a), you argue for dreaming as a “weakly embodied state”, against much of the previous philosophical literature that treated it as a form of skull-bound mentation. Do you see your position within the broader framework of embodied cognition? How do you think your hypothesis can be corroborated through scientific research?

My claim that dreams are weakly embodied states has two parts. The first is related to a phenomenological claim: in a majority of dreams, we do not experience ourselves as fully embodied agents. Dreams are characterized by frequent movement sensations, but other types of bodily experience, such as sensations of touch, pain, pleasure, or temperature, are only rarely experienced in dreams (Windt, 2015a, chap. 7). The pattern of bodily experience in dreams appears to be different from standard wakefulness. The dreams of subjects experiencing phantom limbs are a good example. Following amputation, many people report feeling that the absent limb is



still present; frequently the phantom limb will be frozen in an uncomfortable position, as if paralyzed, and associated with unpleasant tingling or even pain sensations. In dreams, phantom limbs can be represented in many different ways. But participants often report that unlike in wakefulness, they were able to see and move their limb, but that the unpleasant tingling and pain sensations that characterized the phantom in wakefulness had disappeared (Brugger, 2008; Mulder, Hochstenbach, Dijkstra, & Geertzen, 2008). This suggests not only that bodily experience in dreams typically differs from wakefulness, but also that body representation in dreams can't simply be described as a whole-body variant of waking phantom limbs. Bodily experience in dreams requires an account of its own.

Body parts can also be missing—in the sense of failing to be represented—in dreams. Sometimes, dreamers can have the feeling that an individual body part is absent, and this feeling can even extend to the whole body: dream reports occasionally describe the feeling of being a disembodied self or of lacking any kind of body, including the sense of being a spatially extended entity. Aside from disturbed multisensory integration, bodily experience in dreams therefore appears to be characterized by a disturbed integration of body-part/whole-body representations.

The second part of the claim that dreams are weakly embodied states is to give it a functional reading: bodily experience in dreams is not completely independent of the sleeping, physical body (Windt 2015a, chap. 8). Instead, real-body sensations, including sleeping position, REM-sleep related muscular paralysis, as well as subtle movements occurring throughout sleep (such as REM-related muscle twitching), shape dream experience, and bodily experience in dreams can often be described as involving illusory misperception of the sleeping body. While environmental stimuli such as sounds and light flashes are also occasionally incorporated into dreams without leading to awakening, the highest incorporation rates appear to occur for body stimulation. For instance, a blood-pressure cuff inflated on the leg leads to related dream content, as identified by independent raters, in 40-80% of cases (Nielsen, 1993; Nielsen, Ouellet, & Zadra, 1995). Moreover, many intense dreams, including nightmares, appear to have a strong bodily component. A good example is the dream of being unable to flee from a pursuer. The dream of being chased is at the top of typical dream themes (Nielsen et al., 2003; Schredl, Ciric, Götz, & Wittmann, 2004; Yu, 2008), or dreams that most people say they have had at some point in their lives. This does not mean that the chase dream is representative of the majority of dreams—but only that there is something particularly memorable about it. For now, the important point is that the dream of being unable to flee from a pursuer—of having incomplete control of one's legs or even feeling paralyzed—can be straightforwardly explained by appealing to illusory own-body perception: when one becomes aware of the comparative inactivity of the sleeping body and the REM-



related loss of muscle tone while dreaming, this may be experienced as inability to control one's legs—along with associated fear and the feeling of being chased.

By combining these two readings, we get the claim that dreams are *weakly phenomenally-functionally embodied* states: the distinctive pattern of bodily experience that characterizes a majority of dreams is closely related to the altered functional relationship to the sleeping body. Both readings mark a profound departure from standard theories of dreaming, where dreaming is considered a paradigm example of conscious experience unfolding independently of sensory input and motor output. In this view, the situation of the sleeping, dreaming brain is essentially that of a brain temporarily encased in a cranial vat. In philosophy, internalists have typically argued that because dreaming involves a rich form of conscious experience that on the phenomenological level of description is indistinguishable from waking experience, dreams show that conscious experience in general depends on brain activity alone (Revonsuo, 2005). Proponents of embodied, extended, and enactive accounts typically accept that dreaming is a state of functional disembodiment and cranial envatment, but deny that the same is true for perceptual experience (Noë, 2005). This is related to the phenomenological claim that conscious experience in dreams differs from standard perceptual experience precisely because in dreams, conscious experience is cut off from the world.

The disagreement between internalists and externalists is about the correct phenomenological description of dreaming; both sides agree that dreaming is a real-world example of spontaneously occurring cranial envatment. In my view, however, this is false: in a majority of dreams, the processing of peripheral and bodily stimuli is altered, but not completely suppressed. Moreover, this altered functional relationship between the body and the brain is reflected on the level of phenomenal experience itself. Given a better understanding of this functional relationship, we can work towards a more precise description and explanatory account of the phenomenology of embodied selfhood in dreams. This view is not just empirically plausible, but also has important theoretical and methodological consequences, suggesting that bodily experiences in dreams are best conceived of as illusory own-body perception and that any scientific explanation of dream experience will have to look beyond the brain and appeal to the sleeping body. I call this the *body-brain-body* problem: the problem of how the functional interaction between body and brain brings about bodily experience in dreams.

“ In a majority of dreams, the processing of peripheral and bodily stimuli is altered, but not completely suppressed. ”

The view and research strategy I propose are inspired by a classical theory about the sources of dreaming, which was popular in the late 19<sup>th</sup> century but has fallen into disfavor: *Leibreiztheorie*, or somatic source theory, which says that dreaming in general arises in response to bodily sensations (for a newer version, see Schönhammer, 2005). My view is weaker in a number of ways. For instance, I do not think that all dreams, or even all aspects of dreaming, can be explained in this way—instead, my claim is only that to understand the distinctive pattern of bodily experience, a purely brain-based account that assumes that dream activity unfolds completely independently of the sleeping body and environment is insufficient. I think most researchers would allow that external and peripheral stimuli can be incorporated in dreams. However, the extent to which this is the case, along with its theoretical implications, has not, I think, been sufficiently appreciated. This also has practical consequences for dream research: it suggests the importance of investigating not just the effects of body stimulation on dreams, but also of making more extensive use of EMG from the limbs (and not just, as standardly done, the chin), to investigate the association of movement sensations in dreams and muscle twitches. This would help investigate varying degrees of concordance and discordance between bodily experience in dreams and the sleeping body.

Which consequences does this view have for the debate between internalists and externalists? If the cranial envatment view of dreaming fails, and if there is evidence that most dreams do not replicate waking bodily experience, this means that internalists are deprived of their most important real-world example. Dreaming was, after all, supposed to be a clear and intuitive example of how wake-like experience, including bodily experience, depends on brain activity alone. We can now see that at least for a majority of dreams, this is false. A more differentiated phenomenological characterization, together with a better understanding and new paradigms for investigating the real-body basis of dreaming, is needed.

This does not, however, show that the underlying metaphysical point is false. For the internalist, the dream example can be merely illustrative—and to the extent that it is, I take issue merely with the adequacy of the example, not necessarily with the deeper metaphysical claim it is taken to support. The debate between internalists and externalists is about the constitutive supervenience base, or the minimal set of metaphysically sufficient conditions of conscious experience—roughly, whether anything other than brain activity is needed to bring about experience, or whether the vehicles of experience extend beyond the skull (Block, 2005). My point about dreaming does not directly speak to this issue. Even if it were the case that as a matter of fact, bodily experience in dreams never arose independently of sensory prompting and own-body perception, this would still show only that real-body sensations are causally enabling conditions for bodily experience in dreams to occur. For a

metaphysical claim about the constitutive supervenience base of experience, this is not enough. In defending the claim that dreams are weakly phenomenally-functionally embodied states, I am moving away from the metaphysical debate between internalists and externalists to a more empirically plausible account of dreaming and its place in our taxonomy of mental states. I think this is a constructive move to make, and hopefully one that can inspire new research—but it also involves changing the topic of conversation.

You advocate the *Immersive Spatiotemporal Hallucination (ISTH)* model of dreaming, according to which the invariant “phenomenal core” of dreaming across different sleep stages is a minimal sense of immersive spatiotemporal self-location. Why is this model superior to others in your opinion?

With the ISTH model of dreaming (Windt, 2010, 2015a, chap. 11), I try to address what I take to be two central challenges for a theory of dreaming. This is to find an account that is general enough to characterize the range of dream experiences while also helping to pick out what distinguishes dreams from wake states and experiences occurring during sleep-wake transitions and in sleep that do not qualify as full-fledged dreaming. My proposal is a version of simulation views, which focus on a structural feature of dreaming: the feeling of presence, or of being immersed in a world. I think the key commonality between my view and other versions of simulation views is that they all focus on the experience of a self in a world. To my mind, this convergence is really more important than potential differences—because taken together, simulation views have the power to unify different theoretical accounts of dreaming. They also suggest points of contact with other areas of research—such as virtual reality research or work on full-body illusions—and can be used to develop a more precise framework for describing dreamless sleep experiences as well.

Where ISTH differs from other versions of the simulation view (Revonsuo et al., 2015) is that it focuses on spatiotemporal self-location to offer a simplified account of what characterizes all kinds of dreaming—and with it a criterion for distinguishing minimal forms of dreaming from dreamless sleep experiences that no longer have this immersive, here-and-now structure. The key idea is that to identify the phenomenal core of dreaming, it is useful to look away from the characteristics of a majority of dreams—such as visual imagery and movement sensations, strong emotions, and so on—to experiences that are likely rare, but can still be characterized as dreamlike in some relevant sense.

I think the most interesting dream reports for this type of project are those describing a sense of phenomenal disembodiment. In these dreams, the experience of being or having a body has been lost entirely. Yet dreamers will often say that

they continued to have a self—they now experienced themselves as a disembodied entity, as an abstract mind, and the locus of identification had shrunk to an unextended point. This suggests that the experience of embodiment is not necessary for experiencing oneself *as* a self. There is still a here-and-now experience, or a sense of spatiotemporal self-location—but it is not tied to the experience of having a body or being an embodied agent. In fact, it can be associated with the experience of lacking a body. And yet, participants are still willing to describe this experience, after awakening, as having involved a phenomenal self: they still place themselves at the center of their dream.

“The experience of embodiment is not necessary for experiencing oneself *as* a self.”

I think these types of dream reports are highly informative for theories of self-consciousness and of dreaming. For theories of self-consciousness, they suggest that minimal phenomenal selfhood—the simplest form in which we can experience ourselves as being or having a self—is associated with spatiotemporal self-location. The analysis of dreaming thus extends existing work on minimal phenomenal selfhood (Blanke & Metzinger, 2009), but offers a simplified account. In particular, minimal phenomenal selfhood does not require experiencing oneself as a bodily self and cognitive agent, an entity able to direct their own thoughts and bodily actions. Even body ownership is not required. For theories of dreaming, the ISTH model is helpful because it suggests that spatiotemporal self-location and minimal phenomenal selfhood are a central point of transition between non-immersive and in my terminology dreamless experiences to richer forms of dreamful experience.

Concerning differences to other versions of simulation views, I think these are largely due to slightly different perspectives and research interests. For example, Antti Revonsuo uses the characterization of dreaming as the experience of a self-in-a-world to motivate the virtual reality metaphor of conscious experience (Revonsuo, 2005). His main point, here, is that a majority of dreams are actually very similar to standard waking experience—and this is in line with his aim of using the concept of inner presence, which is illustrated to the fullest extent by dreaming, as a metaphor for conscious experience in general. By contrast, with the ISTH model, I was more interested in focusing on a minimal characterization of dreaming—and this was in line with having a unified theory of dreaming first, along with a framework for describing different types of dream (and dreamless sleep) experience and accommodating the inherent variability of the target phenomenon. Only then, on this basis, does it make sense to determine the location of dreaming in the broader framework of concepts used to describe mental states in wakefulness. I think that

these different strategies—i.e. a primary interest in theories and metaphors of consciousness in general vs a theory and framework of dreaming in particular—led us to emphasize slightly different aspects of world- and self-simulation in dreams.

That said, I think one advantage of the ISTH model and its focus on minimal forms of dreaming, as opposed to the features of a majority of dreams, is that it allows for a high degree of variability within dreams—ranging from bodiless dreams to the experience of full, wake-like embodiment, to the experience of having two bodies at the same time or even of slipping back and forth between them (van Eeden, 1913). The model is also compatible with a large degree of variation in cognitive agency—ranging from certain nonlucid dreams in which thinking and attempts at rational reflection are completely absent to fully lucid dreams involving metacognitive insight into the fact that one is now dreaming plus dream control. I think any theory of dreaming will have to be able to accommodate this underlying variability. And while I think the ISTH model is a step in that direction, it is certainly open to further refinement. Different dimensions aside from spatiotemporal self-location can be distinguished; the spatial and temporal aspects of self-location can be dissociated; and all of these properties can vary by degree. The analysis of transitional states that are on the borders of immersive dreaming and either waking experience (as during sleep onset) or nonimmersive forms of dreamless sleep experience is particularly fruitful in this context, because it can help render the model more precise. For reasons of space I won't go into detail here (but see Windt 2015a, chap. 11); my main point is that as long as the ISTH model is useful to lend further precision to a theory of dreaming and formulate new research questions, it is successful.

### What, in your opinion, are the most important challenges for future research?

I would like to see progress in three main areas. The first is to move beyond dreams to a fuller characterization of the range of sleep-related experience. I think there are now good theoretical and empirical reasons for saying that a range of sleep-related experiences exists that can be characterized as dreamless. Some of these may involve a minimal form of phenomenal experience, and their investigation will yield a fuller account, I hope, of the transitions that take place in sleep from unconscious states via nonimmersive imagery and thoughts to fully immersive dreaming. This will require a change of focus to NREM and especially slow-wave-sleep, and it will hopefully bring together philosophy of mind with dream and sleep research, including work on memory consolidation in sleep and sleep disorders (Windt et al., 2016).

Second, I think that simulation views now offer a sufficient degree of unification to the field of dream research to move beyond sleep and sleep-related

experiences. Here, the challenge is how to integrate dreaming (and dreamless sleep experience) into a broader theory of consciousness and the self—to locate sleep-related experience, as it were, on the map of concepts used to describe standard and altered wake states, as well as disorders of consciousness. The most important contrast conditions, in my view, are full-body illusions, immersive virtual reality, and mind wandering, or spontaneous thoughts arising largely independently of ongoing tasks and environmental demands. I think important progress has already been made in all three areas. The next step will consist in investigating potential continuities that cut across the behavioral states of sleep and wakefulness as commonly defined.

For example, if dreams are, as I claim, weakly phenomenally-functionally embodied states, this places them on a continuum with full-body illusions (such as out-of-body experiences and cases in which participants identify with an avatar in immersive virtual reality) and body-part illusions (such as the rubber-hand illusion). The relative ease with which multisensory conflict (e.g. between visual and tactile cues, as in classical versions of the full-body and rubber-hand illusions) can be used to induce feelings of identification with and ownership for an artificial or virtual body or body part in waking healthy subjects suggests that standard bodily experience is surprisingly flimsy (Hohwy, 2010). By investigating the real-body basis of bodily experience in dreams and the extent to which bodily experience in dreams can be described as involving illusory own-body perception, we can now also chip away at the distinction between sleep and wakefulness from the other direction: in the case of dreaming, the link between bodily experience and the physical body appears to be stronger than commonly assumed. By investigating varying degrees of concordance between bodily experience and its real-body basis, it may then become possible to identify continuities across sleep and wakefulness.

“ Perhaps sleep and wakefulness themselves are not mutually exclusive, and a profound departure from the familiar sleep-wake dichotomy is needed. ”

Similarly, there are good reasons for thinking that conscious sleep mentation is closely related, on the phenomenological and neuroscientific levels of descriptions, to spontaneous thought in wakefulness. It has even been suggested that dreaming can be regarded as an intensified form of waking mind wandering (Fox, Nijeboer, Solomonova, Domhoff, & Christoff, 2013). Again, given a more precise taxonomy for describing the range of conscious experience in sleep—both dreamful and dreamless—we can now ask which features of experience and cognitive processing change in concert with sleep-wake transitions, and which ones are state-independent, remaining more or less stable across the sleep-wake cycle. I think this



is an extremely interesting and fruitful question to ask, both theoretically and empirically. The answer has the potential to undermine our understanding of what it is to be awake and asleep and to show that our understanding of the behavioral states of sleep and wakefulness is more poorly developed than we think. Perhaps, sleep and wakefulness themselves are not mutually exclusive and a profound departure from the familiar sleep-wake dichotomy is needed. If that were the case, we would need a new way of drawing even the most basic distinctions in classifying and experimentally investigating mental states.

Finally, the third area in which I would like to see progress is moving beyond predominantly theoretically and scientifically motivated questions about sleep and dreaming to clinical and practical issues related to sleep quality, sleep disorders, and their relation to mental health and emotional well-being. There has already been progress in moving towards interdisciplinary research bringing together philosophy with sleep and dream research, and dreaming is increasingly recognized as a topic for interdisciplinary consciousness science. The next step should be to think more closely about whether theories of sleep and dreaming, but also, for instance, of mind wandering, have practical and clinical consequences. For example, how can the analysis of sleep-related experience inform the diagnosis and therapy of sleep disorders, such as sleep behavior and insomnia? Can the study of sleep-related experience help make sense of differences between subjective and objective measures of sleep quality, as well as their impact on emotional well-being, attention, and performance in cognitive tasks in wakefulness? If yes, how can these insights be used to improve subjective sleep quality in patients experiencing disturbed sleep and in the general population? If some areas of dream research could have real-world impact on these or related issues, and if philosophical work, e.g. on different kinds of sleep-related experience, had some role in this, I think this would be an absolutely fantastic development.



## References

- Blanke, O., & Metzinger, T. (2009). Full-body illusions and minimal phenomenal selfhood. *Trends in Cognitive Sciences*, 13(1), 7–13.
- Block, N. (2005). Review of Alva Noe, *Action in Perception*. *Journal of Philosophy*, 102, 259–272.
- Brugger, P. (2008). The phantom limb in dreams. *Consciousness and Cognition*, 17(4), 1272–1278.
- Dement, W., & Kleitman, N. (1957). The relation of eye movements during sleep to dream activity: an objective method for the study of dreaming. *Journal of Experimental Psychology*, 53(5), 339–346.
- Dennett, D. C. (1976). Are dreams experiences? *Philosophical Review*, 85(2), 151.
- Desseilles, M., Dang-Vu, T. T., Sterpenich, V., & Schwartz, S. (2011). Cognitive and emotional processes during dreaming: a neuroimaging view. *Consciousness and Cognition*, 20(4), 998–1008.
- Dresler, M., Koch, S. P., Wehrle, R., Spoormaker, V. I., Holsboer, F., Steiger, A., ... Czisch, M. (2011). Dreamed movement elicits activation in the sensorimotor cortex. *Current Biology: CB*, 21(21), 1833–1837.
- Flanagan, O. (1995). Deconstructing dreams: the spandrels of sleep. *The Journal of Philosophy*, 92(1), 5–27.
- Fox, K. C. R., Nijboer, S., Solomonova, E., Domhoff, G. W., & Christoff, K. (2013). Dreaming as mind wandering: evidence from functional neuroimaging and first-person content reports. *Frontiers in Human Neuroscience*, 7.
- Hobson, J. A. (1988). *The Dreaming Brain* (1st edition). New York: Basic Books.
- Hohwy, J. (2010). The hypothesis testing brain: some philosophical applications (pp. 135–144). Macquarie Centre for Cognitive Science.
- Horikawa, T., Tamaki, M., Miyawaki, Y., & Kamitani, Y. (2013). Neural Decoding of Visual Imagery During Sleep. *Science*, 340(6132), 639–642.
- Huber, R., Felice Ghilardi, M., Massimini, M., & Tononi, G. (2004). Local sleep and learning. *Nature*, 430(6995), 78–81.
- Jouvet, M. (1999). *The Paradox of Sleep: The Story of Dreaming*. (L. Garey, Trans.). Cambridge, Mass.: The MIT Press.
- Krocker, K. (2007). *The Sleep of Others and the Transformation of Sleep Research* (1 edition). University of Toronto Press, Scholarly Publishing Division.
- Mahowald, M. W., Cramer Bornemann, M. A., & Schenck, C. H. (2011). State dissociation, human behavior, and consciousness. *Current Topics in Medicinal Chemistry*, 11(19), 2392–2402.
- Malcolm, N. (1962). *Dreaming*. Routledge and Kegan Paul.

- Metzinger, T. (2006). Reply to Hobson: Can there be a first-person science of consciousness? *Psyche*, 12(4). Retrieved from <http://www.theassc.org/files/assc/2648.pdf>
- Mulder, T., Hochstenbach, J., Dijkstra, P. U., & Geertzen, J. H. B. (2008). Born to adapt, but not in your dreams. *Consciousness and Cognition*, 17(4), 1266–1271.
- Nielsen, T. A. (1993). Changes in the kinesthetic content of dreams following somatosensory stimulation of leg muscles during REM sleep. *Dreaming*, 3(2), 99–113.
- Nielsen, T. A., Ouellet, L., & Zadra, A. L. (1995). Pressure stimulation during REM sleep alters dream limb activity and body bizarreness. *Sleep Research*, 24, 134.
- Nielsen, T. A., Zadra, A. L., Simard, V., Saucier, S., Stenstrom, P., Smith, C., & Kuiken, D. (2003). The typical dreams of canadian university students. *Dreaming*, 13(4), 211–235.
- Noë, A. (2005). *Action in Perception*. Cambridge, Mass.: The MIT Press.
- Noreika, V. (2014). It's not just about the contents: searching for a neural correlate of a state of consciousness. Open MIND. Frankfurt am Main: MIND Group. Retrieved from <http://open-mind.net/papers/it2019s-not-just-about-the-contents-searching-for-a-neural-correlate-of-a-state-of-consciousness-a-commentary-on-wolf-singer/getAbstract>
- Noreika, V., Valli, K., Lahtela, H., & Revonsuo, A. (2009). Early-night serial awakenings as a new paradigm for studies on NREM dreaming. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, 74(1), 14–18.
- Rechtschaffen, A., & Buchignani, C. (1992). The visual appearance of dreams. In J. S. Antrobus & M. Bertini (Eds.), *The neuropsychology of sleep and dreaming* (pp. 143–155). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.
- Revonsuo, A. (2005). *Inner Presence: Consciousness as a Biological Phenomenon*. Cambridge, Mass.: The MIT Press.
- Revonsuo, A., Tuominen, J., & Valli, K. (2015). The avatars in the machine: dreaming as a simulation of social reality. Open MIND. Frankfurt am Main: MIND Group. Retrieved from <http://open-mind.net/papers/the-avatars-in-the-machine-dreaming-as-a-simulation-of-social-reality/getAbstract>
- Schönhammer, R. (2005). “Typical Dreams”: Reflections of Arousal. *Journal of Consciousness Studies*, 12(4–5), 18–37.
- Schredl, M., Ciric, P., Götz, S., & Wittmann, L. (2004). Typical dreams: stability and gender differences. *The Journal of Psychology*, 138(6), 485–494.
- Schwartz, S. (2000). A Historical Loop of One Hundred Years: Similarities Between 19th Century and Contemporary Dream Research. *Dreaming*, 10(1), 55–66.
- Schwitzgebel, E. (2002). Why did we think we dreamed in black and white? *Studies in History and Philosophy of Science Part A*, 33(4), 649–660.

- Schwitzgebel, E. (2011). *Perplexities of Consciousness*. Cambridge, MA: Bradford Book.
- Siclari, F., LaRocque, J. J., Bernardi, G., Postle, B. R., & Tononi, G. (2014). The neural correlates of consciousness in sleep: a no-task, within-state paradigm. *bioRxiv*, 12443.
- Siclari, F., Larocque, J. J., Postle, B. R., & Tononi, G. (2013). Assessing sleep consciousness within subjects using a serial awakening paradigm. *Frontiers in Psychology*, 4, 542.
- Sikka, P., Valli, K., Virta, T., & Revonsuo, A. (2014). I know how you felt last night, or do I? Self- and external ratings of emotions in REM sleep dreams. *Consciousness and Cognition*, 25, 51–66.
- Singer, W. (2014a). State or content of consciousness? Open MIND. Frankfurt am Main: MIND Group. Retrieved from <http://open-mind.net/papers/state-or-content-of-consciousness-a-reply-to-valdas-noreika/getAbstract>
- Singer, W. (2014b). The ongoing search for the neuronal correlate of consciousness. Open MIND. Frankfurt am Main: MIND Group. Retrieved from <http://open-mind.net/papers/the-ongoing-search-for-the-neuronal-correlate-of-consciousness/getAbstract>
- Smallwood, J., & Schooler, J. W. (2015). The science of mind wandering: empirically navigating the stream of consciousness. *Annual Review of Psychology*, 66(1), 487–518.
- Tamaki, M., Bang, J. W., Watanabe, T., & Sasaki, Y. (2016). Night watch in one brain hemisphere during sleep associated with the first-night effect in humans. *Current Biology*, 26(9), 1190–1194.
- Thompson, E. (2014). Dreamless sleep, the embodied mind, and consciousness. Open MIND. Frankfurt am Main: MIND Group. Retrieved from <http://open-mind.net/papers/dreamless-sleep-the-embodied-mind-and-consciousness-the-relevance-of-a-classical-indian-debate-to-cognitive-science/getAbstract>
- van Eeden, F. (1913). A study of dreams. *Proceedings of the Society for Psychical Research*, 26, 431–461.
- Voss, U., & Hobson, A. (2014). What is the state-of-the-art on lucid dreaming? - recent advances and questions for future research. Open MIND. Frankfurt am Main: MIND Group. Retrieved from <http://open-mind.net/papers/what-is-the-state-of-the-art-on-lucid-dreaming-recent-advances-and-questions-for-future-research/getAbstract>
- Voss, U., Holzmann, R., Tuin, I., & Hobson, J. A. (2009). Lucid dreaming: a state of consciousness with features of both waking and non-lucid dreaming. *Sleep*, 32(9), 1191–1200.
- Windt, J. M. (2010). The immersive spatiotemporal hallucination model of dreaming. *Phenomenology and the Cognitive Sciences*, 9(2), 295–316.
- Windt, J. M. (2013). Reporting dream experience: why (not) to be skeptical about dream reports. *Frontiers in Human Neuroscience*, 7.
- Windt, J. M. (2015a). *Dreaming: A Conceptual Framework for Philosophy of Mind and Empirical*

- Research* (1 edition). Cambridge, Massachusetts ; London, England: MIT Press.
- Windt, J. M. (2015b). Just in time—dreamless sleep experience as pure subjective temporality. Open MIND. Frankfurt am Main: MIND Group. Retrieved from <http://open-mind.net/papers/just-in-time-dreamless-sleep-experience-as-pure-subjective-temporality-a-commentary-on-evan-thompson/getAbstract>
- Windt, J. M. (2016). Dreams and dreaming. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/win2016/entries/dreams-dreaming/>
- Windt, J. M., Nielsen, T., & Thompson, E. (2016). Does consciousness disappear in dreamless sleep? *Trends in Cognitive Sciences*, 20(12), 871–882.
- Windt, J. M., & Voss, U. (forthcoming). Spontaneous thought, insight, and control in lucid dreams. In K. Christoff & K. Fox (Eds.), *Oxford Handbook of Spontaneous Thought*. Oxford: Oxford University Press.
- Yu, C. K.-C. (2008). Typical dreams experienced by Chinese people. *Dreaming*, 18(1), 1–10.





# ALIUS BULLETIN

exploring the diversity of consciousness

n°1

FEBRUARY 2017

[aliusresearch.org](http://aliusresearch.org)