

The Philosophy of Psychedelic Transformation

by

Christopher Edward Ross Letheby B.A., M. Phil.

Journal of Consciousness Studies

*Department of Philosophy, School of Humanities The University of Adelaide
A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy*

December 2016

original text: <https://digital.library.adelaide.edu.au/dspace/bitstream/2440/104676/2/02whole.pdf>

backup source: <http://www.psilosophy.info/resources/02whole.pdf>

Table of Contents:

Abstract

Declaration

Acknowledgements

Introduction

Psychedelic Phenomenology

Psychedelic Science

The Philosophical Background

Philosophy of Psychedelics

The Philosophy of Psychedelic Transformation

Statement of Authorship

Introduction and Methodological Preliminaries

Psychedelic Transformation: An Overview

Evidence for the Causal Relevance of the ASC

Epistemic Aspects of Psychedelic Transformation

Conclusion

The Epistemic Innocence of Psychedelic States

Statement of Authorship

Introduction

Psychedelic Therapy: An Overview

The Concept of Epistemic Innocence

Epistemic Benefit and Epistemic Risk

The No Alternatives Condition

Conclusion

Acknowledgements

Naturalizing Psychedelic Spirituality

Statement of Authorship

Introduction

The Neuroexistentialist Predicament

An Existential Medicine?

Naturalistic Entheogenics

Mechanisms of Mysticism

Conclusion

Anatomy of an Avatar: Ego Dissolution in Psychedelic Experience

Statement of Authorship

Introduction

Binding

Binding and Predictive Processing

The Architecture of Subjectivity

- Self Binding
- Psychedelic Ego Dissolution
- Ego Dissolution as Unbinding
- Conclusion
- Conclusion**
- Summary of Arguments
- Directions for Future Research
- References**

Abstract

Recent scientific research arguably confirms the existence of a remarkable phenomenon: durable psychological benefit to an individual resulting from a single ingestion of a psychedelic drug. In this thesis by publication I ask what exactly is going on in such cases of 'psychedelic transformation'. The thesis is situated in the context of a resurgence of interest in psychedelics within neuroscience and psychiatry, and motivated by the need for philosophical examination of the foundations and results of this research program.

Two common claims in the literature on therapeutic and transformative uses of psychedelics are: (a) psychedelic experiences are a reliable means of knowledge acquisition or spiritual growth, and (b) such epistemic or spiritual benefits are centrally involved in psychedelics' psychological benefits. My aim is to show how such a conception of psychedelic transformation as an epistemic or spiritual process may be reconciled with philosophical naturalism. Naturalism denies the existence of non-natural or supernatural realities such as gods, souls, and immaterial minds. Naturalism is a very widespread view in philosophy today, and is supported by strong arguments.

However, there is a tension between naturalism and the epistemic or spiritual conception of psychedelic transformation, because many psychedelic users claim drug-facilitated knowledge of non-natural realities, such as a mystical 'universal consciousness'. One possible naturalistic response is to dismiss psychedelic users' claims of epistemic or spiritual benefit as mistaken. Here I offer an alternative, showing that we need not throw the epistemic/spiritual baby out with the non-naturalistic bathwater. I show that some kinds of psychedelic-induced epistemic and spiritual benefits are compatible with naturalism and plausibly involved in the drugs' transformative effects.

In 'The Philosophy of Psychedelic Transformation' I review evidence for psychological benefits of psychedelics and defend the claim that psychedelic transformation is unlike standard pharmacotherapies in centrally involving meaningful conscious experiences. I give arguments for three kinds of epistemic benefits: knowledge by acquaintance of the mind's potential, knowledge by acquaintance of the contingency of the sense of self, and revitalised capacities for the acquisition of modal knowledge. In 'The Epistemic Innocence of Psychedelic States' I extend this work, arguing that whatever psychedelics' epistemic demerits, they offer reliable and sometimes unique access to substantial epistemic benefits, including indirect epistemic benefits resulting from psychological benefits. I argue that a balanced picture of the drugs' epistemic merits and demerits is essential to policy discussions about their uses.

In 'Naturalizing Psychedelic Spirituality' I argue that by disrupting mechanisms of self-representation in the brain, psychedelics engender transformative experiences of self-transcendence and mind-expansion which amount to a naturalistic form of spirituality. I propose that such naturalistic spirituality constitutes a viable response to existential anxiety resulting from a naturalistic worldview. Finally, in 'Anatomy of an Avatar: Ego Dissolution in Psychedelic Experiences', I (with Philip Gerrans) argue that such self-transcendent experiences result specifically from disruption to cognitive binding processes implemented by predictive models in the brain. This provides a mechanistic basis for some of the claims of epistemic and spiritual benefit defended earlier.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint award of this degree.

I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

I acknowledge that copyright of published works contained within this thesis resides with the copyright holder(s) of those works.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Christopher Letheby

Signature:

Date:

Publications:

Letheby, C. (2015). The philosophy of psychedelic transformation. *Journal of Consciousness Studies* 22, 170-193.

Letheby, C. (2016). The epistemic innocence of psychedelic states. *Consciousness and Cognition* 39, 28-37.

Letheby, C. (under review). Naturalizing psychedelic spirituality.

Letheby, C. and P. Gerrans (under review). Anatomy of an avatar: ego dissolution in psychedelic experience.

Acknowledgements

With standard but nonetheless sincere apologies to anyone I've forgotten, I thank the following people who have been instrumental in bringing this thesis to completion: Gerard O'Brien especially, and Jon Opie as well, for first-rate supervision, and for taking a chance on an unusual and controversial-sounding topic; Philip Gerrans, for inspiration, collaboration, and ongoing casual employment; everyone in the Department of Philosophy at the University of Adelaide, but especially the participants in our cognitive neuroscience reading group: Glenn Carruthers, Simon Eddy, Rob Farquharson, Anastasiya Kravchuk, Michael Lopresto, Andy McKilliam, Matt Nestor, Greg O'Hair, Laura Ruggles, Liz Schier, and Dook Shepherd; Miri Albahari, for timely encouragement; all the courageous researchers and experimental subjects who have played a part in the psychedelic renaissance; all my friends who have helped and supported me over the past four years; and finally, Mum and Dad, for arguing with me about God and raising me on Douglas Adams and Star Trek, both of which I'm pretty sure helped make me the aspiring philosopher that I am.

Introduction

Recent scientific research has arguably confirmed the existence of a remarkable and long-attested phenomenon: durable and beneficial change to an individual's psychological make-up resulting from a single ingestion of a psychedelic drug. What exactly is going on in such cases of 'psychedelic transformation'? That is the central question motivating this thesis.

The renaissance of scientific research into classic, serotonergic psychedelics, such as lysergic acid

diethylamide, psilocybin, mescaline, and dimethyltryptamine, has provided evidence relevant to this question, but has not thereby answered all aspects of it directly. In particular, psychedelic subjects frequently claim that their transformation is a direct result of some kind of learning, insight, or spiritual growth. Determining whether psychedelic transformation is essentially an epistemic or spiritual process requires not just empirical research but philosophical analysis of the results of such research. This is so because what counts as 'epistemic' or 'spiritual' is a partly conceptual, philosophical question anyway, and because there are no straightforward entailments of philosophical conclusions by scientific findings; such inferences require synthesis and interpretation of empirical results.

Despite repeated calls for philosophical analysis of psychedelic phenomena, there has as yet been relatively little attention paid in academic philosophy to the renaissance of psychedelic research. The four papers collected here comprise an effort to initiate philosophical discussion of psychedelic transformation. Specifically, they collectively amount to an attempt to reconcile two independently plausible but seemingly incompatible claims. The first of these is philosophical naturalism: the view that the natural world is all that exists, and that there are no non-natural or supernatural entities such as gods, souls, or immaterial minds. The second is the above-mentioned claim that psychedelics are effective agents of knowledge acquisition or spiritual growth.

I call this second claim the 'entheogenic conception' of psychedelics, from the neologism *entheogen* ('generating the divine within') which has been coined to refer to psychoactive substances when used specifically for spiritual purposes (Ruck *et al.* 1979). The tension between naturalism and the entheogenic conception of psychedelics arises from the fact that claims of knowledge acquisition and spiritual benefit consequent upon psychedelic use are often closely tied to non-naturalistic metaphysical claims (e.g. Vaughan 1983). In these four papers I attempt to show that the entheogenic conception of psychedelics is independent of metaphysical nonnaturalism - that there is a plausible and robust conception of psychedelics as agents of epistemic and spiritual benefit that is perfectly consistent with naturalism. I apply a broadly neurophilosophical approach to this issue, examining the findings from psychedelic science and then showing how they support this naturalized entheogenic conception.

In this introduction I begin by describing the phenomenon to be analysed: therapeutic and transformative effects occasioned by psychedelics. I give a relatively detailed description of the phenomenology of psychedelic states, since the psychedelic experience is the fundamental background to my project, and its phenomenology is not described in detail in any of the papers. I then very briefly summarise results from recent psychedelic science, which are described in more or less detail in each of the four papers, before describing the philosophical background to my project. This consists, firstly, of several non-psychedelic currents in recent philosophy which provide thematic and methodological inspiration and context for my arguments, and, secondly, of the few formal philosophical discussions of psychedelic epistemology and spirituality in recent times. Finally, I state the specific aims of each of the four papers and their intended contributions to the overall project.

I think my arguments show that an entheogenic conception of psychedelics is consistent with naturalism and plausible in light of current scientific knowledge. But regardless of the success of these specific arguments, the most general aim of the thesis is to put psychedelic science on the philosophical map. My project functions as a sort of 'proof-of-concept', demonstrating that rigorous philosophical analysis of psychedelic phenomena is both possible and desirable. The four papers collected here show beyond doubt that these controversial and long-neglected substances raise fascinating questions meriting substantial and sustained philosophical attention.

A note on terminology: my main focus in this thesis is 'classic' psychedelics such as lysergic acid diethylamide (LSD), psilocybin (found in 'magic mushrooms'), mescaline (from peyote and other psychoactive cacti), and dimethyltryptamine (DMT, found in various plants and animals.) All of these drugs exert their psychoactive effects primarily by serotonin-2a receptor agonism (Halberstadt 2015). In most cases (though there is some variation between papers) I reserve the word 'psychedelic' for substances of this class, and not other substances such as ketamine, *Salvia divinorum*, and ibogaine, all of which act by different neuropharmacological mechanisms, despite being classified as 'psychedelic' on phenomenological grounds (Sessa 2012).

Psychedelic Phenomenology

What is it like to be intoxicated by a psychedelic? An economical but unhelpful answer would be that psychedelic consciousness is too heterogeneous and ineffable to admit of useful description. It is true that each psychedelic experience is unique, being shaped not just by substance, dosage, and route of administration, but by 'set and setting' - the individual's psychological make-up and the circumstances in which they take the drug. Moreover, psychedelic subjects often say that the intoxication is unlike anything in their prior experience and words cannot do it justice.

However, even if we cannot give necessary and sufficient phenomenological conditions for a psychedelic experience, we can describe some of the more typical effects; certain themes recur in both narrative accounts and more formal quantitative (Studerus *et al.* 2011, Schmid *et al.* 2015) and qualitative (Turton *et al.* 2014, Gasser *et al.* 2015) studies of psychedelic phenomenology. These include alterations of various kinds to sensory experience in all modalities, as well as to thinking, the experience of one's body, space, and time, affect, and the sense of self. It is worth appreciating, though, that detailed and poetic narrative accounts probably come closest to conveying a true sense of what psychedelic experience is like (cf. Durr 1970).

Alterations to vision are a characteristic effect which fall into at least three categories: intensification of stimulus-bound percepts; alteration of stimulus-bound visual experience; and stimulus-independent percepts. The latter can happen with eyes open or eyes closed. In the most dramatic case, visual experience becomes entirely constituted by stimulus-independent imagery, amounting to the apparent experience of a completely different reality.

Aldous Huxley describes the intensification of colour in his famous account of mescaline intoxication: "Mescaline [sic] raises all colours to a higher power and makes the percipient aware of innumerable fine shades of difference, to which, at ordinary times, he is completely blind" (Huxley 1954, 27). Similar intensification can occur to other aspects of vision: shapes, edges, depth and so forth may all seem more clear, vivid and forceful than usual.

Alteration or distortion of stimulus-bound percepts introduces an element of misrepresentation. One common variety of this is objects which are actually stationary being seen as moving or morphing (as in the famous phenomenon of 'walls breathing'); another is for objects to be seen as having shapes or proportions other than their actual ones. Objects may be overlaid with geometrical patterns, stationary or moving, with no apparent basis in the environment (Strassman 2005).

Overlain geometric patterns occupy a middle ground between distortion of stimulus-bound percepts the experience of novel, stimulus-independent percepts. A more unambiguous example of this third category is furnished by a peyote subject who described seeing a "ball of red fire about the size of a golf ball" moving towards him in the room:

It drifted, swaying a little from side to side, while moving toward me... when the ball of fire had come close enough I poked at its centre with my finger. It then exploded, a lavish shower of multicoloured sparks cascading and dropping on the rug at my feet.
(Masters and Houston 1966, 7-8).

This fireball and its behaviour were part of the subject's visual experience but seemed to have no basis at all in the objects and events of his environment.

Vivid endogenous visual experience with eyes closed is also very common. The content of such 'closed eye visuals' ranges from simple moving or transforming geometric patterns through fleeting visions of faces, motifs and scene snapshots, all the way to elaborate, immersive and in some cases interactive visual narratives. At high doses, many subjects describe watching and participating in cinematic dramas which seem to constitute allegories of situations in their lives (Masters and Houston 1966). The most extreme case of such visionary experience is when a subject is seemingly transported to a completely different reality or location which is seen in graphic detail with eyes open or closed. In such 'breakthrough' experiences awareness of the actual environment in all sensory modalities is completely lost (Strassman 2001).

Alterations to visual experience are probably more common than changes in other sensory modalities (Strassman 2005). However, with the exception of closed eye visuals (which do not really have analogues) all of these effects can occur in other modalities. Auditory percepts may be intensified, seeming louder (Fischman 1983, 75), clearer, or possessed of more fine-grained detail; appreciation of music is often heightened considerably (Kaelen *et al.* 2015). One subject on psilocybin described hearing his own breathing as though it were the sound of a waterfall (Malitz *et al.* 1960, 10), which seems like an alteration of an existing percept. Novel auditory percepts are sometimes experienced; the peyote subject quoted earlier reported hearing the music that was being played by characters in a vision he observed (Masters and Houston 1966, 9).

Tactile, gustatory and olfactory perception are less notable sites of perceptual changes than vision and audition, although dramatic intensification can occur in these modalities:

Sensations were acute. I heard, saw, felt, smelled and tasted more fully than ever before (or since). A peanut butter sandwich was a delicacy not even a god could deserve ... To touch a fabric was one's fingertip was ... to experience intense touch-pleasure.

(Masters and Houston 1966, 10).

Reports of perceptual distortion or novelty in these three modalities are rare. Intensification or apparent enhancement seem to be the main effects of psychedelics on touch, taste, and smell.

Finally, synaesthetic interactions between sensory modalities often occur in the psychedelic state (Luke and Terhune 2013). The most common form seems to be auditory-visual synaesthesia; psychedelic subjects listening to music often find that their closed-eye visuals vary with, and seem to embody, such features of the music as rhythm and tone (Pahnke and Richards 1966, 185).

One's own body can be a significant locus of abnormal experience on psychedelics. Some such experience occurs in one of the five modalities already discussed; for instance, visual perception of one's own body can be intensified, distorted, or embellished with novel percepts. In particular, psychedelic subjects looking at themselves in a mirror often report unusual experiences such as seeing themselves as caricatures emphasising physical features or personality traits, or seeing themselves in the guise of archetypal or mythical figures. Such experiences can also occur in visual perception of other people (Masters and Houston 1966).

Other abnormal bodily experiences also occur. Some subjects report feeling as though their body is heavier or lighter than usual, or more or less near to their conscious mind (which is sometimes felt to leave the body altogether.) Subjects may also report experiencing the transformation of their body into the form of an animal, robot, or other non-human creature, or its shrinkage to microscopic or expansion to cosmic proportions. The body may feel as though it has been transmuted into a different substance, such as glass, or wood (Masters and Houston 1966).

It is also possible to feel as though one is directly experiencing internal physiological processes, such as the circulation of blood, the functions of organs, or the operations of the brain:

I felt that I could count the beats, the throbbing of my heart, feel the blood moving through my veins, feel the passage of the breath as it entered and left the body, the nerves as they hummed with their myriad messages.

(Masters and Houston 1966, 10).

As well as alterations to sensory and bodily experience, changes to spatial and temporal awareness are distinctive features of psychedelic states. Huxley claimed his perceptions of space and time were unaltered but his attitudes toward them had changed: he no longer had much interest in either of them (Huxley 1994). Space and time seem to have diminished in salience for him.

More unambiguous alterations to spatial awareness involve changes to perceived distance and relationships: objects appear nearer or further away than they actually are. Subjects who report 'mystical' experiences typically claim to have inhabited a state without any experience of space or time (Pahnke and Richards 1966).

Distortions of temporal perception are very common. Our much-quoted peyote subject described feeling as though he had been smoking a cigarette for hours, then looking down at it to see that the first ash was still on

it, implying he had been smoking it for less than a minute (Masters and Houston 1966). Time dilation - the sense that more time has passed than actually has - is far more common than the reverse; which, however, is not unknown. Once again, mystical experiences are a limiting case of alterations to temporal experience: one of the criteria for classifying an experience as 'mystical' is that it involves some sense of eternity or transcendence of time (Griffiths *et al.* 2006).

It is common for the experience of thinking to change dramatically under the influence of psychedelics. Interestingly, many people comparing psychedelics with other drugs remark on how *unimpaired* they find their mental faculties. Huxley again: "The ability to remember and to 'think straight' is little if at all reduced. (Listening to the recordings of my conversation under the influence of the drug, I cannot discover that I was then any stupider than I am at ordinary times.)" (Huxley 1994, 13).

Nonetheless, many psychedelic experiences involve unusual thinking of some kind. Apparent psychological insight is a common example: psychedelic subjects often feel that they comprehend their own personalities, motivations, behavioural patterns and relationships with an unprecedented clarity (Masters and Houston 1966, 185). Thoughts about significant philosophical and scientific questions such as the mind-body problem, the function of the brain, and the meaning of life are also common, even in subjects not predisposed to contemplating such matters (Shanon 2002). Subjects sometimes feel that their thinking is faster or slower, or more or less clear, than usual, and often feel more prone to creative and associative than logical or rational thought (Grof 1975).

Changes to affective experience are a core aspect of the psychedelic state. They mark the fundamental difference between 'good trips' and 'bad trips'. Subjects enjoying mystical experiences experience unparalleled heights of rapture, serenity, and joy, albeit often combined with or preceded by awe or fear (Wasson 1957, Griffiths *et al.* 2006). More intuitively 'psychotomimetic' (psychosis-mimicking) experiences involve equally unparalleled depths of terror, helplessness, and despair (Masters and Houston 1966, 15-17).

Katz *et al.* (1968) made much of their observation that low-dose LSD subjects often displayed an affective ambivalence, consisting of the alternation between or simultaneous experience of mutually antagonistic, if not contradictory, emotions or moods. They suggested that this ambivalence is a central and theoretically important aspect of the psychedelic experience which may explain some characteristics of the more dramatic experiences possible at high doses. The idea that ambivalence or uncertainty is a key aspect of the psychedelic state features in recent theoretical proposals discussed further below (Carhart-Harris *et al.* 2014).

In general, a typical psychedelic experience involving neither a profound mystical rapture nor a terrifying temporary psychosis involves changes to affect which can be seen as analogous to the changes in sense experience: intensification of feelings can occur, as can affective lability (Strassman 1984) - similar to the distortion of existing sensory percepts - and novel, apparently endogenous affective states can occur too: changes in emotion or mood with no obvious environmental trigger or object (Katz *et al.* 1968).

Changes to the sense of self are another core feature of psychedelic experience. Subjects often feel that their sense of individual identity, or the boundaries between self and world, have been altered, loosened, or diminished. As with changes to spatial and temporal experience, this 'ego dissolution' phenomenon reaches its limit in mystical experiences in which the sense of separate individual identity is said to be altogether abolished, replaced by a much larger sense of identity.

Of all the varieties of psychedelic experience, the 'mystical' variety has provoked some of the greatest interest and controversy. Obviously the first question is what exactly a mystical experience is. The category is vacuous if it includes any experience which a subject calls 'mystical' or 'religious', so it needs to be given a more specific definition (Masters and Houston 1966).

The wellspring for academic discussion of psychedelic-induced mystical experience is Walter Pahnke's (1963) doctoral dissertation, which reports the results of the famous 'Good Friday experiment' in which divinity students were given psilocybin in a placebo-controlled study and many more psilocybin than control subjects satisfied criteria for having had a mystical experience. Pahnke based his criteria on the work of the philosopher W. T. Stace (1960) who endorsed the 'common core thesis' (Hood 2006): the claim that mystical experience, despite differences in description and interpretation, has a temporally and culturally invariant

phenomenological essence. Pahnke follows Stace in endorsing the common core thesis and offers a phenomenological typology according to which the central features of a mystical experience are: a sense of unity; transcendence of time and space; deeply felt positive mood; sense of sacredness; feeling of objectivity and reality; paradoxicality; and alleged ineffability.

A mystical experience, then, is an experience of a sense of unity transcending time and space, which is felt by the experiencer to be sacred and objectively real, albeit paradoxical and ineffable, and is accompanied by a deeply felt positive mood. Descriptions of such experiences abound in the literature of the world's religions (Huxley 1945). Apart from being induced by drugs, such experiences can be induced by a variety of religious or spiritual practices, or can occur spontaneously.

There has been considerable debate about the relations between psychedelic-induced and non-psychedelic-induced mystical states; are psychedelic mystical experiences 'authentic', or an artificial facsimile of the genuine article? (Zaehner 1958, Watts 1962, Smith 1964). Pahnke's study was an attempt to apply an empirical methodology to these issues. Subjects wrote descriptions of their experiences, and independent observers equipped with Pahnke's typology rated the extent to which each description exemplified each of the elements of mystical experience. Reports of psilocybin subjects, but not of control subjects, overwhelmingly satisfied the criteria, leading to the conclusion that at least some drug-induced states are indistinguishable from spontaneous mystical experiences (Pahnke 1963, Smith 1964).

Pahnke's study had several methodological problems (Doblin 1991). One of the most celebrated results from the new wave of psychedelic research has been a far more rigorous replication of Pahnke's findings by a research group at Johns Hopkins University (Griffiths *et al.* 2006, 2008, 2011, MacLean *et al.* 2011). In these studies, most high-dose psilocybin experiences satisfied criteria for a 'complete' mystical experience. Subjects who had a mystical experience invariably rated their psilocybin session among the five most "personally meaningful" and "spiritually significant" events of their lives, comparing the meaningfulness of these experiences to such things as the birth of a child or the death of a parent (Griffiths *et al.* 2006).

Thus, as well as inducing a characteristic constellation of remarkable changes to sense perception, affect, and the senses of space, time, body, and self, psychedelics are capable of inducing states of consciousness phenomenologically indistinguishable from non-drug-induced mystical experiences. Indeed, Yaden *et al.* (2016) found that respondents to a large-scale survey rated psychedelic-induced mystical experiences as significantly *more* mystical, intense, and life-changing than non-drug-induced mystical experiences.

It is all of these changes to consciousness, and their sometimes durable psychological consequences, that the recent renaissance of psychedelic science seeks to understand.

Psychedelic Science

The new wave of psychedelic research has sought, among other things, to reassess earlier claims made in the 1950s and 60s for the therapeutic and transformative effects of this class of drugs, but with higher standards of methodological rigour. Importantly, supervised psychedelic administration in controlled research conditions has been shown to be quite safe in both healthy volunteers and patient populations (Johnson *et al.* 2008). Various small pilot studies have found preliminary evidence for the alleviation of symptoms of depression, addiction, and anxiety lasting several months after just one or two psychedelic experiences (dos Santos *et al.* 2016a, Garcia-Romeu *et al.* 2016). More recently, two larger, double-blind randomized controlled trials have found further evidence for the efficacy of psilocybin in alleviating psychological distress relating to terminal illness (Griffiths *et al.* 2016, Ross *et al.* 2016).

Studies in non-patient populations have also found that psilocybin can induce mystical experiences causing positive personality change lasting over a year in a majority of healthy volunteers (Griffiths *et al.* 2006, 2008, 2011, MacLean *et al.* 2011). Subjects who had mystical experiences rated their psilocybin sessions as among the most personally meaningful and spiritually significant events of their lives. More recently, a smaller study found that a single LSD session caused similar personality changes in healthy volunteers lasting at least a fortnight (Lebedev *et al.* 2016).

Besides studies of therapeutic and transformative effects, much basic research has been carried out aimed at

identifying the precise nature and mechanisms of psychedelic effects. Numerous animal and human studies have probed the molecular mechanisms and effects of psychedelics on intracellular signalling, on behaviour, and on measures of neuropsychological function and cognitive processes such as attention, working memory, visual and emotional perception, and creativity (Halberstadt 2015, Nichols 2004, 2016).

One of the most interesting lines of research in the psychedelic renaissance has been the use of modern neuroimaging technologies to probe neural correlates of psychedelic consciousness alteration in humans (reviewed in dos Santos *et al.* 2016b). Positron emission tomography (PET) studies of mescaline and psilocybin in the 1990s consistently found increases in glucose metabolism concentrated in frontal brain regions, suggestive of a pattern of metabolic 'hyperfrontality' in the psychedelic state (Hermlle *et al.* 1992, Vollenweider *et al.* 1997).

This finding has since been challenged by functional magnetic resonance imaging (fMRI) studies of the psychedelically intoxicated brain. Carhart-Harris *et al.* (2012) reported that acute intravenous psilocybin administration caused decreases, rather than increases, in neural activity as measured by blood-oxygen-level-dependent (BOLD) signal. Moreover, these decreases were localised to nodes of the Default Mode Network (DMN), a functionally coherent network of brain regions repeatedly implicated in mind-wandering, autobiographical memory, and various aspects of self-reference. Subsequent fMRI studies of psychedelics including LSD and ayahuasca have not replicated all the findings of Carhart-Harris *et al.*, but decreased BOLD signal in DMN hubs is a frequent result.

The discrepancy between the fMRI and PET results remains to be explained. Various solutions have been proposed. Carhart-Harris *et al.* (2012) suggested that it might be due to the different time courses of oral vs. intravenous psilocybin intoxication, as well as differences in the temporal resolution of the imaging technologies used - but fMRI studies of the DMT-containing psychedelic beverage ayahuasca, which has a similar time course to oral psilocybin, also found decreases rather than increases in signatures of neural activity. Another possibility is that the BOLD signal measured by fMRI tracks neural synchrony, rather than activity, and the psychedelic state is characterised by decreased synchrony but increased activity in key hub regions (Halberstadt 2015). It is safe to say that more research in this area is required.

Such unresolved issues notwithstanding, fruitful and intriguing theoretical proposals have been made on the basis of the recent spate of fMRI results. Citing their findings that psychedelics increase the unpredictability and diversity of patterns of functional connectivity in the brain, as quantified by the notion of entropy, Carhart-Harris *et al.* (2014) proposed that psychedelics induce a highly entropic state of consciousness by destabilizing key connector hub regions whose activity ordinarily functions to constrain the quality of cognition, orchestrating and coordinating the activity of many large-scale networks to ensure that cognitive resources are directed towards adaptive goals. Speculations about the mechanisms of psychedelic therapy have focused on the idea that by this kind of method, psychedelics can break down entrenched and maladaptive patterns of thought and behaviour, creating a period of greater flexibility from which the cognitive system can re-form into more adaptive patterns (Nichols *et al.* forthcoming).

Despite many open questions concerning functional neuroanatomical details, this broad perspective according to which psychedelics disrupt systems involved in self-reference, leading to temporarily 'unconstrained cognition' with potential therapeutic benefits, sits comfortably with the phenomenology as described earlier, and has formed the basis for many of my arguments in this thesis. I turn now to the non-psychedelic philosophical background to my project.

The Philosophical Background

In recent decades philosophy of science has taken a turn towards recognition of the heterogeneity of science and analysis of the specific and distinctive methods and explanatory strategies deployed in the various special sciences (e.g. Craver 2007, Ross 2014, Sober 2000). Despite the extensive literature on ethical issues relating to drug use, especially for non-medical reasons, psychopharmacology as a discipline has received relatively little attention from a philosophy of science perspective. The notable exception is Stein's (2012) pioneering volume, which discusses conceptual, explanatory, and methodological questions relating to how we classify and categorise drugs and how best to understand their effects on cognition, consciousness and the self (as well

as discussing relevant ethical issues.) However, Stein only mentions psychedelics in passing.

A few recent publications suggest that interest in philosophy of psychopharmacology is increasing (e.g. Aragona 2013, Cavanna 2015). The renaissance of psychedelic research constitutes a golden opportunity for this nascent field. All sorts of interesting conceptual and methodological questions are raised by psychedelic science. For instance, given the dramatic acute and long-term effects of psychedelics on mood and behaviour, psychedelic therapy seems like an ideal case study for recently influential ideas about multi-level mechanistic explanation in the life sciences - an issue touched on in the first paper of this thesis. Moreover, there are methodological questions about how to test for therapeutic effects of an intervention for which adequate blinding is difficult if not impossible (Langlitz 2012). In general, recent trends in philosophy of science suggest a need for more philosophical analysis of (psychedelic) psychopharmacology.

Another part of the philosophical context in which my project is situated is the recent spate of interest in re-evaluating the epistemic status of intuitively suboptimal cognitive conditions and processes. There is a long history of philosophical distrust of non-ordinary states of cognition and consciousness, and a tendency to see such conditions not only as highly dubious sources of knowledge, but as positive threats to knowledge (Windt 2011). This tendency is understandably more pronounced in the case of conditions which clearly are attended by obvious psychological or epistemic costs, such as delusion, confabulation, and implicit bias. And the undeniable detrimental psychological effects which can sometimes result from psychedelic experience, combined with the conception of psychedelics as 'psychotomimetic' or 'hallucinogenic' - that is, essentially epistemically detrimental - substances would seem to place psychedelic states squarely in the camp of epistemically bankrupt consciousness alterations.

However, recently a number of philosophers have been arguing that the epistemic status of such 'imperfect cognitions' merits re-evaluation. The literature on 'epistemic innocence' proposes that even cognitive conditions with undeniable epistemic costs sometimes lead to significant and otherwise unavailable epistemic benefits (Bortolotti 2015a, 2015b, Bortolotti and Miyazono 2016, Sullivan-Bissett 2015). In many cases, the argument is that these cognitive conditions can lead to psycho-emotional benefits, and epistemic benefits result from those, since knowledge acquisition is often dependent upon adaptive psycho-emotional functioning. The scientific reassessment of the potential of psychedelics to cause lasting psychoemotional benefits suggests that psychedelic states ought to be subjected to the same kind of epistemological re-evaluation. Moreover, the many distinctive aspects of psychedelic therapy, including frequent claims of epistemic benefit by subjects, mean that the analysis of psychedelic phenomena may have implications for this broader project.

The growing interest in how to tackle the existential anxiety and 'disenchantment' resulting from a naturalistic worldview is another recent philosophical current which is relevant to my concerns here. It has long been appreciated that the transition to a post-religious naturalist metaphysics often evokes anxiety about the absence of traditional foundational sources of meaning. Recently, it has been suggested that this existential crisis is filtering into the broader culture and becoming more acute due to the revolution in the self-image of humankind brought about by advances in contemporary neuroscience (Flanagan and Caruso forthcoming, Metzinger 2009). How to individually and collectively come to terms with this 'neuroscientific turn' in the image of humankind is arguably one of the most urgent philosophical challenges of our age, connecting with a classical Socratic vision of philosophy as concerned with the art of living.

The capacity of psychedelics to induce intense and meaningful transformative experiences immediately suggests their potential relevance to this issue. What is more, there has recently been a spate of philosophical interest in the related project of 'naturalizing spirituality' (Stone 2012). This literature is based on a recognition that the term 'spirituality' can legitimately be used to refer to a kind of personal quest for meaning and transcendence that is in principle independent of traditional religious belief. The question whether such a quest can be undertaken in the spirit of naturalism is interesting and important. Plausibly, something like naturalizing spirituality may form part of a viable response to the 'neuroexistentialist' anxiety resulting from the naturalistic disenchantment of the world. Psychedelic subjects' frequent tendency to describe their experiences as 'spiritual' suggests that an analysis of psychedelic phenomena may help to address some of these issues.

Finally, recent philosophy of mind has taken a distinctly empirical orientation, looking to data from neuroscience and psychiatry to inform its theorizing. This trend is most clearly exemplified in the

interdisciplinary enterprise of 'philosophical psychopathology' (Graham and Stephens 1994), which treats observations from the study of psychological disorder as a crucial source of information about the structure of the mind in both normal and abnormal conditions. One of the key concerns of scientifically and psychopathologically informed philosophy of mind has been the nature of self-awareness, a central and ubiquitous mental phenomenon that undergoes sometimes bizarre transformations in altered conditions of psychological functioning (Stephens and Graham 2007, Metzinger 2013). Moreover, such conditions may even offer useful clues to the nature of phenomenal consciousness itself (Brogaard and Gatzia 2016).

The neurophilosophy of self-awareness and consciousness has hitherto drawn on whatever sources have been available, including, to some extent, the study of psychedelics (Metzinger 2003). But the ever-increasing body of new knowledge about psychedelic-induced alterations to consciousness and self-awareness has perhaps not yet been fully philosophically appreciated or exploited. The empirical turn in philosophy of mind is yet another recent development which justifies a detailed philosophical investigation of psychedelic phenomena.

Philosophy of Psychedelics

As I have mentioned, relatively little explicitly philosophical attention has been paid to the psychedelic renaissance. However, that is not to say that none has. I will now describe the few philosophical discussions which are directly relevant to my project of naturalizing the entheogenic conception of psychedelics, before stating the aims and intended contributions of each of the four papers in the thesis.

Benny Shanon is a cognitive psychologist with a philosophical background who has conducted extensive research on ayahuasca. His monograph *The Antipodes of the Mind* (2002) is a landmark work of psychedelic phenomenology. But Shanon has also explored some of the philosophical issues raised by the psychedelic experience.

In his paper 'The Epistemics of Ayahuasca Visions' Shanon (2010) assesses psychedelic subjects' claims of knowledge acquisition from an avowedly naturalistic perspective. He argues that psychedelics (ayahuasca, in particular) can allow a subject to attain therapeutically valuable psychological insight into their own and others' minds and the human condition. This claim is echoed by Metzinger (2003) whose primary concern is the use of psychedelic data to inform theories of self-awareness. Nonetheless, he notes that the apparent therapeutic efficacy of the drugs is readily explained by the hypothesis that they can in fact facilitate valuable insights.

Shanon proposes that other putatively epistemic benefits are accessible by means of psychedelic experience. One of these is the apprehension or appreciation of truths about nature in new and deeper ways that are comparable to the kinds of distinctive and insightful conceptions of things communicated in great art.

Other kinds of epistemic benefits mentioned by Shanon include increased well-being and stamina; more harmonious social interaction; novel solutions to problems in one's field of specialisation; and enhanced abilities of artistic performance and creativity. It is an interesting conceptual question to what extent each of these proposed benefits merits the appellation 'epistemic'.

Shanon says a little about the means by which he thinks ayahuasca-induced visions impart these kinds of knowledge. One is the direct impact of the visions themselves, which are extremely beautiful; merely experiencing such beauty causes a profound affective response leading to changes in attitude and outlook. Sometimes subjects experience visions of entities who instruct them directly; for example, about their field of specialisation. In other cases, vivid visions of past (actual) and future (hypothetical) situations in a subject's life enables them to see abstract features and commonalities of these, particularly when sequences of such scenes are presented juxtaposed in thematic rather than chronological ways. Further, interactive visions can function as a kind of virtual reality, by acting in which a subject can learn problem solving strategies and other things learnable by acting in the real world.

Albeit not explicitly cast in terms of knowledge or spirituality, it is also worth mentioning the work of Kenneth Tupper, an education theorist who analyses the cognitive enhancement potential of psychedelics in a way that seems consistent with naturalism. Drawing on the work of philosophers and education theorists, Tupper (2002) proposes that psychedelic experiences can be a means of accessing "mythic and somatic forms of

understanding" that are neglected in contemporary educational practice, and also that they can be a means of increasing a hypothesised "existential intelligence" (Tupper 2003). Thomas Roberts (2013) is another education theorist who has extolled the potential cognitive benefits of psychedelics.

On the other side of the debate, G. T. Roche (2010) critiques the claim that psychedelics can induce epistemically or spiritually beneficial experiences. His two basic grounds for scepticism are that psychedelics disrupt the functioning of the cognitive mechanisms which allow the brain to represent reality accurately, and that drug-inspired claims need independent justification to count as knowledge.

Many of Roche's arguments specifically target putative non-naturalistic knowledge gained from psychedelic states and are therefore irrelevant to my concerns. But he argues that even the potentially naturalistically palatable claim that psychedelics can induce veridical apprehensions of oneness with the universe is untenable because the idea of experience without a subject is contradictory. He also argues that psychedelics cannot reliably yield moral or existential insight, because many putative such insights seem flippant, amoral, or otherwise objectionable, and says that there is no good evidence that psychedelic experiences can lead to scientific insight. Such experiences, for Roche, involve radical distortions of cognition and perception which are obviously inimical to such insight. He also argues against the claim that psychedelics can induce mystical experiences equivalent to non-drug-induced ones.

Against the background of these discussions, here are the aims and intended contributions to the overall project of the four papers in this thesis. In the first paper, 'The Philosophy of Psychedelic Transformation', I begin by arguing for the importance of a rigorously naturalistic philosophical analysis of psychedelic phenomena. I then combine scientific evidence concerning the benefits of psychedelics with ideas from philosophy of science to argue that psychedelic transformation is a distinctively humanistic drug therapy because its mechanism of action involves conscious mental representations. I further suggest that this meaning-involving therapeutic mechanism has epistemic dimensions. Establishing the consciousness- and meaning-involving nature of psychedelic transformation is a prerequisite for naturalizing the entheogenic conception of the drugs, and arguing for epistemic dimensions in psychedelic therapy contributes to this project directly by specifying plausible naturalistic epistemic benefits which are consistent with scientific knowledge about transformative mechanisms.

In the second paper, 'The Epistemic Innocence of Psychedelic States', I extend these efforts by arguing that psychedelics can induce additional epistemic benefits as a consequence of their psychological benefits, and moreover that their epistemic benefits - both direct and indirect - are sometimes not available by any other means. The aim here is to further develop the naturalistic grounds for seeing psychedelics as agents of epistemic benefit, and also to connect these ideas to broader discussions about the epistemic innocence of 'imperfect cognitions'. I emphasize the relevance of these conclusions to policy discussions about psychedelic use.

In the third paper, 'Naturalizing Psychedelic Spirituality', I argue that the beneficial effects of psychedelics on existential distress in terminal illness suggest that these drugs may offer the way to a solution to the neuroexistential anxiety accompanying the naturalistic turn in the image of humankind. I connect this proposal to recent philosophical discussions of naturalized spirituality and argue that many of the transformative aspects of psychedelic states which lead subjects to call them 'spiritual' are independent of non-naturalistic metaphysical apprehensions. This amounts to a naturalistic vindication of the specifically spiritual aspect of the entheogenic conception.

Finally, in the fourth paper, 'Anatomy of an Avatar: Ego Dissolution in Psychedelic Experience', I (with Philip Gerrans) argue that the ego dissolution experiences occasioned by psychedelics can best be explained as a disruption to mechanisms of cognitive binding (i.e. representational integration) which give rise to the experience of self-awareness in ordinary waking consciousness. This contributes to naturalizing the entheogenic conception by vindicating the idea that ego dissolution, a core element of spiritual psychedelic experiences, amounts to a veridical apprehension of the non-existence of a durable substantial self. It also provides a plausible mechanism for many of the pro-epistemic and pro-spiritual qualities of the psychedelic experience in the form of detaching attention and salience from their usual coupling to personal concerns, goals, and agendas.

Many questions remain, and I will mention these in the conclusion section at the end. But in these four papers I will show that despite arguments to the contrary, the idea of psychedelics as agents of epistemic and spiritual benefit is alive and well in light of the recent renaissance of research.

The Philosophy of Psychedelic Transformation

Statement of Authorship

Title of paper:	The Philosophy of Psychedelic Transformation
Publication status:	Published
Publication details:	Letheby, C. (2015). The philosophy of psychedelic transformation. <i>Journal of Consciousness Studies</i> 22, 170-193.
<i>Principal Author</i>	
Name of principal author (candidate):	Christopher Letheby
Contribution to the paper:	Devised the arguments and wrote, proofread, polished, formatted, and submitted the paper.
Overall percentage (%):	100%
Certification:	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.
Signature:	
Date:	

Introduction and Methodological Preliminaries

Psychedelic drugs are remarkable substances which have been hailed as indispensable epistemic instruments for the sciences of mind, as unparalleled psychotherapeutic interventions, as unique sources of insight into the nature and genesis of psychosis and religion, and as keys to the survival and flourishing of the human species (Osmond 1957, Sessa 2012). After a politically-driven decades-long hiatus, scientific study of these drugs in humans has resumed with impressive results. Given the magnitude and variety of significance ascribed to the substances, it is surprising that philosophers have not shown much interest in this 'psychedelic renaissance'

Here I aim to remedy this deficit. Recent philosophical work on psychopharmacology focuses on bioethical questions of authenticity and autonomy with respect to enhancement (Parens 2005, Glannon 2008, Stein 2012). As such, this is a natural place to begin philosophical analysis of psychedelic phenomena. The psychopharmacological enhancement literature is driven by a concern that certain psychopharmacological interventions may be dehumanising or may compromise the authenticity or identity of patients. Without committing myself to the view that other drugs do compromise authenticity, I want to explore the notion that psychedelics respect authenticity in a unique way - by involving the person in a transformative process which is somewhat¹ transparent, rational, and meaning-involving.

This is similar in spirit to contrasting utopian and dystopian views of psychopharmacology which the philosopher M. H. N. Schermer (2007) finds in the work of Aldous Huxley. In Huxley's novels, fictional psychedelic-like drugs are depicted as humanising and empowering while fictional non-psychedelic-like drugs are depicted as dehumanising and disempowering (Huxley 1932, 1962). I am not committing myself to the dystopian view of non-psychedelic drugs, but exploring the possibility that psychedelic drugs have uniquely utopian credentials. My conjecture is that this difference arises from the fact that psychedelics engage the self in a humanistic transformative process which is (somewhat) transparent and meaning-respecting, rather than performing sub-personal surgery on the constituent parts of a passive self.

I begin by briefly reviewing the history and phenomenology of psychedelics, and recent evidence for therapeutic and transformative efficacy. Next, I discuss my conjecture about the meaning-involving nature of psychedelic transformation. This conjecture depends on the empirical claim that the altered state of consciousness (ASC) induced by psychedelics is causally relevant to the long-term benefits caused by the drugs. I discuss four lines of evidence for this claim - three briefly, before spending longer on the fourth, which draws on recent neuroscientific studies of psilocybin. This research has been claimed to support a theory of psychedelic transformation which implicates the ASC. I describe this theory and argue that it not only

implicates the ASC in the transformative process, but supports the idea that psychedelic transformation is an epistemic process. It is tempting to assume that psychedelics are fundamentally agents of misrepresentation (as suggested by the popular term 'hallucinogen'.) From this perspective, one might think that whatever psychedelic therapy is, it is not a process of knowledge acquisition. I suggest this is too quick, and there are good reasons to think that psychedelic experience can sometimes lead to knowledge gain.

Before I begin, some methodological preliminaries are in order. Here I take a naturalistic approach to the phenomenon of psychedelic transformation. Naturalism is a contested term in contemporary philosophy but minimally is taken to mean a rejection of supernatural entities or realities, and an acceptance that everything which exists forms part of a closed causal system populated by the kinds of entities and properties studied by the natural sciences (Papineau 2009). Many psychedelic experiences - especially profoundly transformative ones - involve the apparent apprehension of non-naturalistic realities, and many subjects believe these apprehensions are veridical (Masters and Houston 1966, Shanon 2002).

Given the preceding observation, why take a naturalistic approach? Firstly, naturalism may well be true. There are strong arguments in its favour which convince most current philosophers of mind that some version of it probably is true (Horst 2009). And if naturalism is true, the correct account of psychedelic transformation is naturalistic. Secondly, if naturalism is false, we can establish this only by reaching its limits. The principle of parsimony (Ockham's Razor) dictates that we posit only what is necessary to explain a phenomenon. This raises the question: what could establish the necessity of positing non-natural realities to explain psychedelic phenomena? The answer, I submit, is: trying, and failing, to explain such phenomena solely in terms of the independently well-motivated naturalistic posits already available to us. Cognitive neuroscience is an interdisciplinary attempt to explain mental phenomena - including consciousness - naturalistically. Its total success or partial failure remains to be seen. But the cognitive neuroscience of consciousness is in its infancy, and presuming its eventual partial failure at this point would be premature². So considerations of parsimony dictate that we assume, at least for the present, that psychedelic consciousness can be explained naturalistically.

There is another good reason to try to naturalize psychedelic transformation. The modern West is unusual among human societies in lacking socially approved ritual techniques for the deliberate induction of altered states of consciousness (Bourguignon 1973). Given the apparent safety and possible benefits of long-term ritual psychedelic use (e.g. Grob *et al.* 1996, Bouso *et al.* 2015) this may well be our loss. But in pluralistic, secular societies, this cultural situation is unlikely to be changed by an argument that psychedelics facilitate access to non-natural realities. Some envision the possibility of psychedelics finding a broader-than-merely-medical role in contemporary society (e.g. Tupper 2003). For this to be possible, we need a metaphysically innocent, naturalistically palatable way of understanding what exactly these drugs offer besides the mimicry or alleviation of psychopathology.

Psychedelic Transformation: An Overview

Psychedelic drugs form a chemically and neuropharmacologically diverse group of substances. What unites them is phenomenology: despite differences in molecular structures and receptor binding profiles, all drugs classified as 'psychedelic' induce a characteristic type of altered state of consciousness. Psychedelic experiences are notoriously variable, being strongly influenced by 'set and setting' - that is, an individual's state of mind prior to ingestion, as well as the interpersonal and aesthetic environment (Sessa 2012). However, virtually every psychedelic experience involves alterations to some non-negligible subset of the following phenomenal modalities: sensory experience, affective experience, spatial and temporal experience, somatic experience, thinking or reasoning, and the sense of self (Strassman 2005). These alterations can be very mild or very intense, and can take the form of distortion, intensification, or diminution of exogenous mental contents, as well as the generation of endogenous mental contents with no apparent basis in the physical environment (Masters and Houston 1966).

This class of drugs has been referred to by a bewildering variety of names, including, but not limited to: hallucinogens, psychotomimetics, mysticomimetics, entheogens, schizogens, psychotogens, and, of course, psychedelics (Osmond 1957, Szara 1967, Ruck *et al.* 1979). This abundance of terminology reflects an abundance of uses. Naturally occurring psychedelics such as mescaline, dimethyltryptamine, and psilocybin

have been used as religious sacraments in various cultures for centuries at least. Serious scientific interest in psychedelics was sparked by the accidental discovery in the 1940s of the extremely potent psychoactivity of lysergic acid diethylamide (LSD). In the 1950s and 60s, these drugs were studied and used in several different ways (Sessa 2012).

Psychologists and scholars of religion investigated the relationship between psychedelic intoxication and mystical experience, including, famously, experimental attempts to induce mystical states (Pahnke 1963). Various debates also occurred about whether psychedelic mysticism was equivalent to so-called 'genuine' mysticism or somehow inferior (Smith 1964). Meanwhile, psychiatrists used psychedelics in an attempt to understand mental illness - either by studying the psychedelic state as a 'model psychosis', or by experiencing it themselves in order to empathise with their patients. Somewhat surprisingly, psychiatrists also used psychedelics extensively as treatments. The 'psycholytic therapy' model involved classical psychoanalysis conducted under the influence of low doses of psychedelics, which were held to lower patients' defences and allow access to unconscious content. The 'psychedelic therapy' model, on the other hand, involved administration of high doses of psychedelics in order to induce an intense and overwhelming transformative experience (Sessa 2012).

Psychiatrists were very impressed by the efficacy of psychedelics, claiming extraordinary success rates with treatment-resistant alcoholism, as well as other conditions. However, much psychedelic science from the mid-20th century was methodologically sub-optimal by today's standards, often relying heavily on anecdotal evidence or lacking adequate controls, blinding, or follow-up (Doblin 1991, Strassman 1995a, Grob *et al.* 1998). In any case, after a brief heyday, psychedelics were soon made untenable both as treatments and as objects of research. Widespread uncontrolled use of these drugs in the context of the 1960s counterculture led to a moral panic which culminated in the banning of the substances and the subsequent cessation of human research for several decades (Sessa 2012).

Scientific study of the effects of psychedelic drugs on human subjects resumed in the late 1980s and early 1990s - most notably, with Rick Strassman's pioneering studies of dimethyltryptamine, or DMT (Strassman 1995b). These studies were investigations in basic psychopharmacology and did not involve any formal attempt to test for potential therapeutic or transformative effects (although the researchers were certainly aware of the possibility that such effects might occur; Strassman 2001.) Since then, however, many more studies have been conducted which suggest that the mid-20th century enthusiasm about psychedelics may not have been entirely unfounded. More research is still needed, but there is enough information to justify taking seriously the possibility of durable psychological change resulting from a single administration of a drug.

Psilocybin, the active ingredient³ in 'magic mushrooms', has been perhaps the most widely studied psychedelic since the resumption of human research. One study involved the administration of psilocybin to 9 patients with obsessive-compulsive disorder, all of whom showed significant reductions in symptoms not only during the drug experience but also at a follow-up 24 hours later (Moreno *et al.* 2006). Another study tested the effects of a psilocybin session on late-stage terminal cancer patients experiencing anxiety related to their illness, and these patients displayed significant reductions in anxiety during the experience as well as at two month follow-ups (Grob *et al.* 2011). A similar study has also been conducted with terminal cancer patients using LSD instead of psilocybin, and similar results were found (Gasser *et al.* 2014). Recent studies have found promising results treating tobacco and alcohol addiction with psilocybin (Garcia-Romeu *et al.* 2014, Bogenschutz *et al.* 2015). Meanwhile, the psychedelic dissociative anaesthetic ketamine has been shown to cause a rapid reduction in symptoms of treatment-resistant major depression, with this reduction lasting up to a fortnight. Ketamine used in conjunction with existentially-oriented psychotherapy has also been shown to lead to higher rates of abstinence in heroin addicts (Vollenweider and Kometer 2010).

Apart from these studies of therapeutic effects, perhaps the most famous study conducted to date in the 'psychedelic renaissance' involved the administration of psilocybin to mentally healthy subjects who had no prior experience with psychedelics and who reported regular participation in religious or spiritual activities. Of the 36 subjects in this study, 22 experienced a 'complete' mystical experience as determined by the Hood Mysticism Scale and the States of Consciousness Questionnaire. A mystical experience in this sense has seven essential phenomenological components: internal and external unity, transcendence of time and space, alleged ineffability and paradoxicality, sense of sacredness, noetic quality, and positive mood (Griffiths *et al.* 2006). What is even more interesting is that those subjects who had a complete mystical experience showed

significant increases in the core personality domain of Openness. The extent of these increases was predicted by the extent to which a given subject's mystical experience was 'complete', and the increases were diminished but still significant at 14-month follow-ups (MacLean *et al.* 2011). This is impressive because there are very few interventions which have been shown to cause durable and significant changes in any of the 'Big Five' personality domains in adult subjects.

Thus, there is good reason to think that a single administration of a psychedelic can, in conducive circumstances, lead to durable psychological benefit. I turn now to the nature of the transformative process whereby this occurs.

Evidence for the Causal Relevance of the ASC

As I mentioned earlier, bioethicists discussing psychopharmacological enhancement have expressed a concern that changing human personality by using drugs - especially for non-therapeutic or 'cosmetic' reasons - is inevitably dehumanising. One ground for this concern is expressed in the 2003 report of the President's Council on Bioethics as follows: '...biotechnical interventions act directly on the human body and mind to bring about their effects on a passive subject, who plays little or no role at all. He can at best *feel* their effects without *understanding their meaning in human terms*. Thus, a drug that brightened our mood would alter us without our understanding how and why it did so...' (President's Council on Bioethics 2003, 290; my italics).

Now, even if we disagree with the implicit normative verdict, we can recognise standard drug therapies such as antidepressants in this description. But psychedelic transformation seems different. My project begins by considering the intuition that the process of psychedelic transformation might be distinctive in some important or interesting way. This intuition is fuelled by the fact that most psychedelic subjects would *not* recognise psychedelic transformation in the above description. The phenomenology of psychedelic transformation is the antithesis of this kind of description - subjects feel themselves to be actively engaged in a process of transformation, one which crucially involves meaningful experiences and understanding (Masters and Houston 1966, Shanon 2002).

Of course, this feeling might be misleading. From the fact that it seems to subjects that they are engaged in a meaningful transformative process, it does not follow that they in fact are; they could be mistaken. The existence of this feeling does not preclude the possibility that the therapeutic benefits result from a direct pharmacological effect of the drug, unrelated to the ASC. If that were the case, the ASC - and associated feelings of meaningful transformation - would be a mere side-effect, therapeutically speaking. But note at this point that psychedelics certainly are distinctive psychopharmacological interventions in a couple of respects: they induce a vivid ASC which impresses subjects as meaningful, and they (sometimes) cause lasting benefits with a single dose. Standard drug therapies do neither of these things. This lends some initial plausibility to the idea that the subjects' impressions are accurate - that the durable changes to their cognitive and affective functioning are in fact brought about by a process which involves meanings and involves them as subjects. The idea that intense and meaningful experiences can have lasting effects on people is familiar and unmysterious enough, so we should consider the possibility that this is happening in psychedelic transformation.

Of course, to say that psychedelic transformation 'involves meanings' is imprecise and unsatisfactory. At this point I want to suggest an attractively clear and rigorous interpretation of this claim. I think it's quite natural to construe the claim as follows: the causal chain leading from (a) the ingestion of the psychedelic molecule to (b) the subject's enjoying a durable psychological benefit involves phenomenally conscious mental representations. Or, put differently, the psychological benefits are not caused solely by the drugs, but are caused at least partly by the altered state of consciousness, which itself is caused by the drugs.

Now, as I mentioned earlier, this is an empirical claim. We can render it yet more precise and tractable by understanding it as the hypothesis that the psychedelic experience is causally relevant to the psychological benefits. The philosopher Carl Craver, in his analysis of explanatory practice in neuroscience, develops an account of causal relevance as susceptibility to intervention and manipulation (Craver 2007). According to this analysis, the psychedelic experience is causally relevant to the long-term benefits if, and only if, we can reliably manipulate the latter by intervening to manipulate the former. I think there are at least four lines of evidence that this is the case, which I will discuss briefly in a moment.

Before I do discuss this evidence, though, I want to mention that the interpretation I've given of the claim of 'meaning involvement' is a fairly minimal one. I've construed this as the claim that the psychedelic ASC is causally relevant to the long term benefits. But of course that could be true without there being any interesting semantic relation between the acute drug experience and the durable benefits⁴. To see this, consider the possible world in which psychedelic drugs do nothing but induce in subjects vivid and colourful visions of Mickey Mouse. Suppose that in that world, as in the actual world, psychedelics show therapeutic potential for the treatment of conditions like OCD, addiction, and anxiety. And suppose, further, that the more vividly and realistically a given subject's visions portray Mickey Mouse, the greater the magnitude of the transformation enjoyed by that subject.

In the possible world I've just described, the psychedelic ASC is causally relevant to the long-term benefits. That is, we can manipulate the long-term benefits by intervening on the ASC. For instance, if there is a way to promote vividness and clarity of Mickey Mouse visions in certain subjects, then those subjects will enjoy greater benefits than other subjects who received identical doses but had less clear and vivid Mickey Mouse visions. So the minimal interpretation of the meaning-involvement claim is true in the possible world I've described. Nonetheless, it's pretty clear that the members of the President's Council on Bioethics would be unimpressed. This is because although meaning - by which I mean mental representation - is causally involved, that bare fact does not suffice to make the intervention one whose meaning can be "understood in human terms". What is needed for that is some more substantial semantic relation between the experience and the resultant benefits. A pretty clear example of the kind of relation at issue would be a case in which a person is forced under extreme circumstances to parachute out of a plane and consequently loses their lifelong fear of heights. This person has not only had their personality transformed by a conscious experience, but they can understand why that particular conscious experience should have had that particular effect on their personality.

That is the stronger sense of meaning-involvement: not just conscious mental representations being causally relevant to long-term benefits, but also bearing a semantic relation to those benefits such that the latter are rendered humanly comprehensible and able to be situated in a life narrative. Now, there is evidence that at least some cases of psychedelic therapy or transformation are meaning-involving in this stronger sense. Subjects do report having transformative experiences the content of which is transparently relevant to issues in their lives (Masters and Houston 1966; Shanon 2002). Suffice it to note, however, that the minimal interpretation of the meaning-involvement claim is necessary, if not sufficient, for the stronger one. So I shall here limit myself to arguing for the minimal claim that the ASC is causally relevant to the benefits. I turn now to the four lines of evidence for causal relevance.

The first is the phenomenology of psychedelic transformation mentioned earlier. Again, this is far from conclusive. But psychedelic subjects' belief that the benefits are caused by the intoxication merits attention. One possible naturalistic explanation of this belief is that it is true. Maybe subjects believe the intoxication transformed them because it did⁵. This point would be strengthened by some rigorous research into the relations between ASC contents and long-term benefits. Perhaps there are robust and identifiable correlations between certain specific experiential contents and certain specific benefits. Be that as it may, my claim here is just that the tendency of psychedelic subjects to report being transformed by a meaningful experience is a suggestive datum not to be ignored.

The second line of evidence concerns the existence of psychologically beneficial altered states of consciousness not induced by drugs. For example, there is evidence that psychological benefits can result from meditation practice (for reviews, see Davis and Hayes 2011, Goyal *et al.* 2014). Also, relationships have been found between non-drug-induced mystical experiences and measures of subjective well-being (Byrd *et al.* 2000) and durable, (largely) positive life changes are often reported following non-drug-induced near death experiences (Greyson 1997) - though there are obvious obstacles to studying such experiences in a controlled fashion. To the extent that these other altered states resemble psychedelic states, this suggests that the neuropharmacological action of the drugs is an inessential and in principle dispensable means to entering the altered state, which latter does the therapeutic work.

Of course, the question of similarities and differences between altered states is complex. But there is at least some evidence of commonalities between psychedelic and meditative states, as various psychedelic researchers have noted (Hoffmann *et al.* 2001, Palhano-Fontes *et al.* 2015, Stuckey *et al.* 2005). Deactivation

of the posterior cingulate cortex has been observed to correlate with subjective experiences of selflessness (or ego dissolution) occasioned by both methods (Brewer *et al.* 2013). Various kinds of unusual experiences popularly associated with psychedelics are also frequently occasioned by meditation practice. Indeed, the sample data in Jack Kornfield's (1979, 45-50) phenomenological study of Buddhist insight meditation read like a compendium of psychedelic phenomena. Such observations as these form part of the motivation for empirical studies, currently in progress, testing the efficacy of psilocybin as an adjunct to meditation training (Brown and Reitman 2010).

One might object that this line of evidence relies on a spurious and unsustainable ontological distinction between the psychedelic ASC and the neuropharmacological action of the drug. Surely the ASC and the neuropharmacological action are in fact one and the same phenomenon, viewed at different levels of description. The evidence suggests this is not the case, however. Certainly the ASC is entirely constituted by abnormal patterns of neural and synaptic activity. Not all of this activity directly involves the psychedelic molecule, however. Psilocybin, for example, causes its psychological effects mainly by stimulating 5-HT_{2A} receptors, which are found primarily on pyramidal neurons in cortical layer V (Carhart-Harris *et al.* 2014). The abnormal patterns of activity in these regions are at least partly constituted by the synaptic action of the drug itself. But abnormal activity in these regions leads in turn to abnormal patterns of activity in other (e.g. subcortical) regions to which these regions project, and those latter patterns are caused but not constituted by the drug action. We can express this by saying that the neuropharmacological action of the drug is a partially distinct sustaining cause of the ASC - only partially distinct because also partially constitutive, and sustaining because the ongoing drug action is necessary for the continuation of the psychedelic state. What I am suggesting, then, is that more variables than the direct drug action are involved in understanding the long-term effects. The downstream effects of that action in other brain regions are also important, and these effects are presumably a function not just of the direct drug action but of the prior state of the downstream systems (part of the 'set' in 'set and setting').

The third line of evidence that the ASC is causally relevant to the benefits relates to the fact that in some cases, variables quantifying the ASC have been found to predict variables quantifying the benefits. For instance, in the psilocybin mystical experience study mentioned earlier, the extent to which a subject's experience was mystical predicted the magnitude of increases in their personality domain of Openness (MacLean *et al.* 2011). This seems like clear evidence for causal relevance, in Craver's sense: evidence, that is, that if we could 'manipulate' psychedelic ASCs by creating conditions conducive to their being mystical, this would be a reliable means of manipulating the long-term benefits (in the direction of greater Openness increases.)

The question whether the psychedelic ASC is causally relevant to therapeutic benefits has been explicitly addressed by recent studies of ketamine. One small study of cocaine addicts involved three separate sessions: low dose ketamine, high dose ketamine, and lorazepam as an active placebo (Dakwar *et al.* 2014). These different infusions were given to subjects in a double-blind fashion on separate days. Psychological effects of each infusion were assessed using two different scales: one to measure dissociative type effects, and another to measure mystical type effects. It was found that higher mystical effect scores predicted increased motivation to quit cocaine, while higher dissociative effects did not. Further, variance in mystical effect scores predicted variance in increased motivation even across consistent dosages. That is to say that when you compare all the high-dose ketamine sessions, drug dosage is consistent between subjects. But there is variation in the kind of altered state subjects experienced, and this phenomenal variation predicts variation in therapeutic benefits 24 hours after the session. Once again, this is evidence that we can manipulate the lasting benefits by 'intervening on' the ASC, independently of drug dosage.

Such evidence for relationships between phenomenal variables and benefit variables is important. It would be considered a truism by those with firsthand experience of psychedelic research that the quality of the experience is relevant to the long-term outcome and can vary independently of dosage. Certainly it is possible for two people to consume an equivalent dose of a psychedelic but have vastly different experiences, one enjoying a mystic rapture and the other a nightmarish 'bad trip'. In such a case it would be very surprising if the long-term psychological consequences for the two people were the same⁶. But of course such clinical and anecdotal wisdom needs rigorous testing. It has been suggested that classic psychedelics and dissociative anaesthetics (such as ketamine) may share a common therapeutic mechanism in ultimately targeting glutamate-driven neuroplasticity (Vollenweider and Kometer 2010). One possibility is that this is sufficient for

relatively long-term (i.e. weeks) psychological change but further, more specific experiential factors are required to capitalise on this neuroplastic window of opportunity and yield truly durable change⁷. The current instruments used in psychedelic research are relatively coarse-grained; they measure such variables as 'visionary restructuralization' but not specific experiential contents. Perhaps the development of more fine-grained psychometric instruments might reveal further interesting correlations between acute and long-term drug effects.

So, to recap: I construed the conjecture that psychedelic transformation is meaning-involving as entailing (at least) the empirical claim that the altered state is causally relevant to the benefits. Thus far, I've mentioned three kinds of evidence for that causal relevance claim: first, phenomenology; second, beneficial non-drug altered states; and third, correlations between phenomenal variables and psychological benefit variables. I turn now to the fourth line of evidence. This is based on recent neuroimaging studies of the psychedelic state, and also leads to questions about the potentially epistemic nature of psychedelic transformation.

Recently a team led by Dr. Robin Carhart-Harris of Imperial College London gave intravenous injections of psilocybin to healthy volunteers lying in fMRI machines (Carhart-Harris *et al.* 2012). When given intravenously, psilocybin has a very rapid onset and a short duration, which not only makes it practicable for neuroimaging but also allows very precise imaging of the transition from ordinary consciousness to psychedelic consciousness.

The findings from this study were fascinating in a number of respects. Notably, the transition to the psychedelic state involved only decreases, and no increases, in brain activity. This contradicts previous assumptions that psychedelics work by increasing brain activity, and so requires some explaining in its own right. Even more interesting, however, was the localization of these decreases. They were mainly found in the much-discussed Default Mode Network, so-called because it has the interesting property of being most active when a subject is at rest and not engaged in any particular task. When a cognitive or behavioural task is begun, activity in the DMN decreases and activity in other networks correspondingly increases (Raichle *et al.* 2001).

The DMN is also interesting because it displays significantly higher metabolic activity than the average brain region and is extremely densely connected to many other regions. This suggestive set of observations has led to a lot of debate about what exactly the DMN does. Its role is still a matter of controversy, but it has been implicated in various self-referential and metacognitive functions, including daydreaming and 'mental time travel' - that is, the simulation of past and future events (Spreng and Grady 2010). Some theorists have also linked the DMN to the 'narrative self' (D'Argembeau *et al.* 2014).

In discussing their findings, Carhart-Harris and colleagues note that not only did activity in various DMN regions diminish under psilocybin, but many of the normal patterns of connectivity both internal and external to the DMN were disrupted. The result, from the standpoint of global neural dynamics, was a system in a much more disordered and unpredictable state than is ordinarily the case. The authors speculate that this is because the DMN ordinarily acts as a supervisory system which imposes inhibition and constraint on other cognitive systems, and hence disruption to the DMN in the psychedelic state results in a condition of unconstrained cognition (Carhart-Harris *et al.* 2014).

Most relevant to my concerns here are the speculations Carhart-Harris and colleagues make about the mechanisms of psychedelic therapy. They note that various conditions for which psychedelics show therapeutic promise - including OCD, depression, and addiction - can illuminatingly be characterised as 'over-rigid' conditions. A depressed or obsessive system is one trapped for whatever reason in a narrow region of state space. In light of this, they propose that a psychedelic experience might be therapeutic because it forcibly shakes the system out of its rut, freeing it from its rigid confinement and leaving a greater dynamical flexibility which outlasts the experience itself. The point of this, of course, is that this is a conjectural model of psychedelic therapy on which the experience is certainly causally relevant to the benefits⁸.

It is important to note that the psychedelic experience and subsequent benefits are here being described in purely dynamical terms. This raises interesting questions about different levels and kinds of explanation in the cognitive sciences. There are long-running debates about the explanatory credentials of dynamical models (e.g. van Gelder 1998, Kaplan and Bechtel 2011) and the case of psychedelic therapy readily reveals the

limitations of such models alone. Carhart-Harris *et al.* place great importance on entropy as a quantifiable explanatory construct in understanding the mechanisms of psychedelic therapy. However, transiently elevated entropy is insufficient for therapeutic benefit. This is because two different psychedelic experiences could be equally entropic even though one is a blissful mystical-type experience which leads to durably increased openness and the other is a hellish bad trip which leads to trauma and subsequent nightmares. (Such results do occur, although they are very rare in carefully conducted controlled research.) It may be that de-rigidifying the cognitive system by elevating its entropy is necessary but not sufficient for therapeutic benefit, and it is also necessary that the experience have contents of a certain kind. Acute psychotic experiences presumably involve elevated entropy, though they are distinct from controlled psychedelic states in at least three important ways: they are not voluntary, they are not of a fixed and known duration, and they are not accompanied by insight. All this notwithstanding, the neuroimaging studies combined with the rigidity characteristic of depression and other conditions do provide some evidence that increasing cognitive flexibility is an element of the therapeutic process. Therefore, the studies constitute a fourth and final line of evidence that the psychedelic experience is causally relevant to the long-term benefits. So there is good reason to accept the descriptive claim that psychedelic transformation is a distinctively meaning-involving psychopharmacological intervention.

I emphasise that this is a descriptive claim because I am refraining from entering into the normative bioethical debates about the relative merits of meaning-involving and non-meaning-involving transformative processes. I think the meaning-involving nature of psychedelic transformation is surely a fact highly relevant to policy debates, but my project here is purely descriptive. For me, the next interesting set of issues concerns exactly what kinds of meanings are respected or involved in psychedelic transformation, and how.

Epistemic Aspects of Psychedelic Transformation

Psychedelic subjects very often feel that they gain knowledge through their experiences. However, they also reasonably often come to entertain non-naturalistic metaphysical beliefs as a consequence of their experiences (Vaughan 1983, Strassman 2001, Shanon 2002). It is an interesting question what to say about this from a naturalistic perspective. Is the sense of epistemic benefit experienced by these subjects simply illusory? Or is there some kind of naturalistically palatable epistemic benefit which subjects might indeed be gaining, notwithstanding such metaphysical conversions? I think there are perfectly naturalistic epistemic benefits which may well result from psychedelic states.

The first one follows very naturally from the 'unconstrained cognition' theory of the psychedelic state. If this theory is correct, then cognitive systems, while psychedelically intoxicated, traverse wider regions of state space than they do at other times. Often they enter into completely novel and unfamiliar regions of state space, and often these are very distant from the more familiar regions of ordinary waking consciousness. In light of this, I propose that psychedelic experience can be a means of gaining knowledge by acquaintance of one's own vast psychological potential.

Consider, for example, Aldous Huxley. Prior to his famous mescaline experience in 1952, Huxley was a serious student of the mystical literature of the world's religions. He had knowledge by description - by testimony, in fact - of the existence of certain regions of the human phenomenal state space. He believed, truly and justifiably, that there existed certain possible ways for his mind to be - arguably intrinsically and instrumentally valuable ways featuring such things as intensified perceptions and emotions, greatly enhanced appreciation of the world, and a sense of kinship with all of existence. But he may well have doubted, as do many aspiring mystics, that these ways of being were genuine possibilities *for him*. After the 3rd of May, this doubt was gone. Huxley had acquired a new kind of knowledge about the potential of his mind - he had become directly acquainted with its ability to enter states of absorption, harmony, and unparalleled wonder and awe (Huxley 1994). In this context, it is worth mentioning that many senior Western teachers of meditation were inspired by experiences with psychedelic drugs in the 1960s (Badiner and Grey, eds. 2002, Dass 2005). It is plausible to think that they were drawn to the discipline of meditation because they gained knowledge about the immense potential of their own minds. They were then driven to investigate the possibility of realising the potential about which they had gained knowledge.

So psychedelic experiences may well be a means to acquire knowledge by acquaintance of one's own

psychological potential. This is, of course, perfectly naturalistic. A second naturalistically palatable kind of knowledge which such experiences might afford is slightly more controversial. This is knowledge by acquaintance⁹ with the metaphysical nature of the self.

As I mentioned earlier, some theorists have speculated that the Default Mode Network is the neurocognitive substrate of the narrative self - the sense of one's persistent identity as a distinct individual with a history, which sense is constituted by narratively structured representations of one's past and future. Some theorists inclined towards the research programme known as 'neuropsychanalysis' have even suggested that the DMN is the substrate of the Freudian ego (Carhart-Harris and Friston 2010). While remaining agnostic on the precise details, the fMRI studies of psilocybin mentioned earlier offer suggestive evidence about the neurocognitive substrate of the sense of self (Carhart-Harris *et al.* 2012). One of the key nodes of the DMN is the posterior cingulate cortex (PCC). In the studies in question, there was a strong correlation between decreases in PCC activity and ratings for one specific item on the psychological questionnaire used to quantify subjects' experiences. This item, scores for which tracked diminution in PCC activity, read as follows: "I experienced a dissolution of my self or ego" (Carhart-Harris *et al.* 2014).

Obviously a more careful treatment of this point is required. But I think the fact that the psilocybin-induced deactivation of the PCC leads subjects to report a dissolution of their ordinary sense of self suggests that, no matter how real and inviolable it feels in the normal course of things, the sense of self is in fact a model of some kind¹⁰ generated by specific cognitive systems in the brain (cf. Metzinger 2003, 2009). On the basis of ordinary waking life, it is tempting to assume that an experiencing subject is a transcendental precondition for the possibility of experience - that the idea of experience without a self is incoherent. The idea here is that this assumption is false - it is a case of incorrectly inferring necessity from constant conjunction - and psilocybin subjects become directly acquainted with this fact. That is, they gain experiential knowledge of the contingency of their own sense of self by experiencing its temporary subtraction from their phenomenal space. Even if this does not show the sense of self to be a *mere* model, the possibility of experience without the feeling of a subject - the possibility of "thoughts without a thinker" (Epstein 2004)—is a striking discovery to make.

It should be clear that there is no contradiction in the idea that psychedelic therapy is a person-involving process which happens to involve a dissolution of the sense of self. It is quite consistent and plausible to think that the sense of self is a mere part of persons considered as experiencing entities - this sense is just one type of experience which such entities typically, but not invariably, have. It is the person, in the sense of the conscious cognizing organism, which has, and unproblematically later remembers, the experience of ego dissolution.

The two suggestions I've made so far concern the possibility of direct epistemic benefit, of gaining knowledge during the psychedelic experience itself. I will finish by proposing one way in which a transformative or therapeutic psychedelic experience might lead to indirect epistemic benefit. This is by restoring the capacity to acquire modal knowledge¹¹ by restoring the subject's imaginative flexibility.

People suffering from depression, for instance, have difficulty imagining other ways that they could be or certain courses of action they could take. Part of the rigidity mentioned earlier is imaginative rigidity. Once again, the system is trapped in a narrow region of state space and tends not to envision creative solutions to problems or novel behavioural strategies. This seems straightforwardly to be a state of impoverished modal knowledge. There are possibilities available, but the suffering subject is unable to imagine these possibilities and thus unable to know of their availability. In this light, consider the conjectural dynamical model of psychedelic therapy: the system is temporarily unconstrained, conferring a degree of freedom and flexibility, some measure of which outlasts the acute experience. One way this greater flexibility could manifest is as an increased ability to imagine possibilities. And a greater ability to imagine possibilities is, at least, a higher level of access to putative modal truths about oneself and one's life.

Conclusion

To summarise: psychedelic drugs form a phenomenally defined class which includes serotonin agonists such as LSD, DMT, psilocybin and mescaline, as well as dissociative anaesthetics such as ketamine, and other drugs besides. Psychedelic drugs induce a distinctive and intense kind of altered state of consciousness which is different from the altered states induced by drugs of other classes. Psychedelics have been studied again in the last two decades as therapeutic and transformative agents, yielding evidence that they can cause lasting psychological benefits with a single dose or with very few doses.

Some bioethicists worry that cosmetic psychopharmacology is dehumanising because drugs transform passive subjects in a way that is subjectively opaque and not comprehensible in meaningful human terms. A passing acquaintance with the literature on psychedelic therapy leads to the intuition that psychedelic transformation is unlike this. In particular, it seems to be more transparent and meaning-involving. I precisified this intuition as the claim that the psychedelic ASC is causally relevant to the long term benefits, and reviewed four lines of evidence for this claim. These were: first, the phenomenology of psychedelic transformation; second, the existence of beneficial non-drug altered states; third, correlations between phenomenal variables and benefit variables in a dose-independent fashion; and fourth, the de-rigidifying model of psychedelic therapy based on neuroimaging results.

In closing, I suggested three naturalistically palatable kinds of epistemic benefit which might derive from psychedelic experiences: two direct, and one indirect. The first direct kind of benefit is acquiring knowledge by acquaintance of one's own psychological potential. The second direct kind of benefit is acquiring knowledge by acquaintance of the contingency of one's sense of self. And the indirect kind of benefit is a rejuvenation of cognitive capacities which are important for the acquisition of modal knowledge. Many questions remain, but the old idea of drug-induced epistemic benefit merits serious attention, even given naturalism¹².

The Epistemic Innocence of Psychedelic States

Statement of Authorship

Title of paper:	The Epistemic Innocence of Psychedelic States
Publication status:	Published
Publication details:	Letheby, C. (2016). The epistemic innocence of psychedelic states. <i>Consciousness and Cognition</i> 39, 28-37.
<i>Principal Author</i>	
Name of principal author (candidate):	Christopher Letheby
Contribution to the paper:	Devised the arguments and wrote, proofread, polished, formatted, and submitted the paper.
Overall percentage (%):	100%
Certification:	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.
Signature:	
Date:	

Introduction

The recent renaissance of scientific research into psychedelic drugs has provided evidence that these controversial substances can confer lasting psychological benefits in just one or two sessions. Small studies have found promising results using LSD and psilocybin to treat addiction (Johnson *et al.* 2014, Bogenschutz *et al.* 2015), obsessive-compulsive disorder (Moreno *et al.* 2006), and anxiety relating to terminal illness (Grob *et al.* 2011, Gasser *et al.* 2014). This method is unlike existing psychiatric treatments in that it relies crucially on the temporary induction of a dramatically altered state of consciousness (Letheby 2015), at least in some cases (Majiæ *et al.* 2015). While in this altered state, patients often have extremely convincing experiences of what seems to be a mystical, immaterial ultimate reality underlying the universe, and there is some reason to think that these mystical experiences can be important for the therapeutic effects (Bogenschutz and Johnson 2015, Majiæ *et al.* 2015). Indeed, one eminent scholar of religion has said that "the basic message of [these drugs is] that there is another Reality that puts this one in the shade" (Smith 2000). This raises difficult questions. Michael Pollan, writing in the *New Yorker* recently, put it well:

It's one thing to conclude that love is all that matters, but quite another to come away from a therapy convinced that "there is another reality" awaiting us after death ... or that there is more to the universe - and to consciousness - than a purely materialist world view would have us believe. Is psychedelic therapy simply foisting a comforting delusion on the sick and dying? (Pollan 2015).

In this paper I argue that even if psychedelic therapy is "foisting a comforting delusion on the sick and dying" it is not *merely* doing so. To make this argument, I rely on conceptual resources developed in epistemology, the philosophical study of knowledge. The philosopher Lisa Bortolotti and colleagues (Bortolotti 2015a, 2015b, Bortolotti and Miyazono 2016, Sullivan-Bissett 2015) have recently introduced the concept of 'epistemic¹³ innocence' to articulate the idea that certain epistemically sub-optimal states such as delusion and confabulation may have not just psychological but epistemic benefits. They urge us to resist a simplistic trade-off view according to which psychological benefits are purchased with epistemic costs. I apply their analysis to psychedelic therapy and argue that at least some psychedelic states are epistemically innocent imperfect cognitions.

I proceed as follows. In section 1, I give a brief overview of psychedelic therapy: the history and phenomenology of these drugs and recent evidence for their therapeutic and transformative efficacy. In section 2, I describe the concept of epistemic innocence and its application by Bortolotti to delusion. In section

3, I argue that psychedelic states can have significant epistemic benefits. This discussion touches on some interesting issues arising from the fact that psychedelic states, unlike many other imperfect cognitive conditions, are voluntarily entered. In section 4, I argue that often no alternative cognitions are available which will deliver the same epistemic benefits as psychedelics. In the course of making this argument, I highlight a potential ambiguity in the formal definition of epistemic innocence. Finally, in section 5, I reflect on the implications of my discussion. The epistemic innocence of psychedelic states has broad consequences. Mounting evidence for the benefits of psychedelics is leading to a re-evaluation of the potential of these substances as tools in psychiatry and beyond (e.g. Sessa 2005, 2008, Sessa and Johnson 2015). I argue that such policy discussions ought to be based on a comprehensive understanding of the merits and demerits, epistemic included, of the drugs.

Psychedelic Therapy: An Overview

There has recently been a renaissance of scientific research into the effects of psychedelic drugs on human subjects in controlled conditions. Although more evidence is required, including replications with larger sample sizes, the results to date give us good reason to take seriously the idea that these drugs can durably change personality and alleviate psychiatric distress with just one or two administrations. In this section, I give a brief history of psychedelics, followed by a review of the recent evidence for therapeutic and transformative efficacy, in order to set the scene for my discussion of epistemological issues.

Sporadic earlier research notwithstanding, the story of modern psychedelic science really begins in 1943, when the Swiss chemist Albert Hofmann accidentally ingested a minute quantity of lysergic acid diethylamide or LSD-25, a chemical he had been developing for medicinal purposes. He famously described the result as follows:

At home I lay down and sank into a not unpleasant intoxicated-like condition, characterised by an extremely stimulated imagination. In a dreamlike state, with eyes closed ... I perceived an uninterrupted stream of fantastic pictures, extraordinary shapes with intense, kaleidoscopic play of colours.
(Hofmann 1980).

Shortly afterwards, it was recognised that LSD's effects were similar to those of the naturally occurring substances mescaline, found in the peyote cactus, and psilocybin, found in 'magic' mushrooms. British psychiatrist Humphry Osmond (1957) coined the term 'psychedelic', meaning 'mind-manifesting', to describe this newly recognised class of drugs with its distinctive profile of effects. The three drugs I have mentioned not only produce similar effects but act by a common mechanism: stimulation of serotonin 2A receptors (Halberstadt 2015). Various other drugs which produce similar effects by different mechanisms of action have also been classified as 'psychedelic' (Sessa 2012). However, here I restrict myself to discussing the so-called 'classic' serotonin 2A agonist psychedelics (apart from a brief mention of the NMDA antagonist dissociative anaesthetic ketamine, also currently being studied for its therapeutic potential.)

Describing the effects of psychedelics in general is difficult for at least two reasons: first, because they are often held to defy description, and second, because they are extremely variable, being highly dependent on 'set and setting' - that is, the subject's prior state of mind and the external environment. One definition holds that a psychedelic is a drug

which, without causing physical addiction, craving, major physiological disturbances, delirium, disorientation, or amnesia, more or less reliably produces thought, mood, and perceptual changes otherwise rarely experienced except in dreams, contemplative and religious exaltation, flashes of vivid involuntary memory, and acute psychosis.
(Grinspoon and Bakalar 1979).

Such "thought, mood, and perceptual changes" might include seeing static objects as moving or warping, vast and colourful visions of ancient civilizations with the eyes closed, emotional extremes of euphoria or terror, or a sense that one's ordinary everyday identity has dissolved into a timeless ultimate reality - to name just a few (Masters and Houston 1966). Very often there is a sense of insightfulness, understanding, or new perspectives, either with respect to personal life issues or with respect to broader - often philosophical - themes (*ibid.*,

Shanon 2002).

Psychedelics have been used for religious and medicinal purposes by traditional cultures for centuries, if not millennia (Masters and Houston 1966). After coming to the attention of Western science in the 1940s, they were studied from several different perspectives. The 'psychotomimetic' paradigm emphasised the similarity between psychedelic states and naturally occurring psychotic states. Researchers in this paradigm studied psychedelics to try to gain clues about the biochemical basis of mental illness, while psychiatrists sampled psychedelics in order better to empathise with their psychotic patients (Osmond 1957).

Remarkably, as well as being used to model mental illness, psychedelics were used to treat it. Psychotherapy with psychedelics tended to follow one of two models: 'psychoalytic therapy' involved classical psychoanalysis conducted under the influence of low doses of psychedelics, on the grounds that the drugs lowered psychological defensiveness and provided access to unconscious material; 'psychedelic therapy', meanwhile, involved the administration of high doses of psychedelics in order to induce an overwhelming, transformative 'peak' or 'mystical' experience (Grinspoon and Bakalar 1986). Clinicians were very impressed with the therapeutic efficacy of psychedelics, but most of this early research was methodologically problematic by today's standards. Studies often lacked control groups, blinding, and follow-ups, and relied heavily on anecdotal evidence (Sessa 2005). However, one recent meta-analysis of early research did find evidence for the efficacy of high-dose LSD in the treatment of alcoholism (Krebs and Johansen 2012).

Ultimately scientific research and clinical practice with psychedelics was curtailed after the drugs became politically controversial in the 1960s. Widespread uncontrolled use by members of the hippy counterculture led to a highly emotive public debate which culminated in the prohibition of psychedelics. After this, virtually no human research was conducted for some decades (Sessa 2005).

Since the early 1990s, in a changed political climate, human research with psychedelics has slowly but steadily resumed, with intriguing results. Importantly, it has been demonstrated repeatedly and rigorously that it is possible safely to give moderate or high doses of psychedelics to carefully screened and prepared subjects in controlled clinical settings without serious adverse effects (Johnson *et al.* 2008). Various studies have also examined the drugs' therapeutic and transformative potential.

In one study, 9 patients with obsessive-compulsive disorder, all of whom had failed to respond to standard treatments, showed significant reductions in symptoms during the drug experience and at a follow-up 24 hours later (Moreno *et al.* 2006). Another study tested the effects of a psilocybin session on terminal cancer patients experiencing anxiety, and these patients showed significant reductions in anxiety and depression during the experience and at follow-ups two and six months later (Grob *et al.* 2011). A similar study using LSD instead of psilocybin found similar results (Gasser *et al.* 2014). Recent pilot studies have found promising results treating tobacco and alcohol addiction with psilocybin (Johnson *et al.* 2014, Bogenschutz *et al.* 2015). Promising results have also been found using ayahuasca, a South American shamanic beverage containing the classic psychedelic DMT, to treat addiction (Thomas *et al.* 2013) and depression (Osório *et al.* 2015).

Another study examined the effects of psilocybin on mentally healthy religiously-inclined volunteers. 22 of 36 subjects had a 'complete' mystical experience as measured by psychometric scales, involving feelings of unity, transcendence of time and space, and extreme positive affect (Griffiths *et al.* 2006). Those who had a mystical experience showed increases in the personality dimension of openness which were diminished but still significant 14 months later. Importantly, the intensity of their mystical experience predicted the extent of their increases in openness (MacLean *et al.* 2011).

This brief review of recent research reveals substantial evidence that psychedelic drugs can cause lasting psychological benefits with just one or two doses. The mechanism whereby they achieve this result remains uncertain but there is reason to think that the altered state of consciousness induced by the drugs is crucially involved. Specifically, mystical states of consciousness seem to have the greatest therapeutic benefit. In the psilocybin study just mentioned, the extent of mystical experience predicted the degree of personality change. And in the study of psilocybin for tobacco addiction, mystical experiences predicted therapeutic outcome (Garcia-Romeu *et al.* 2014). In the alcoholism study, mystical experience - as well as overall intensity of drug effects - predicted improvements to drinking behaviour (Bogenschutz *et al.* 2015).

From the perspective of philosophical naturalism - the view that the natural world is all that exists - such mystical experiences must be considered illusory, which leads to a troubling picture indeed. A treatment modality which has been hailed as the "next major breakthrough in mental health care" (Jacobson 2014) seems to work, at least in part, by inducing highly compelling metaphysical illusions. For help in thinking about this situation, I turn now to the recent philosophical work on epistemic innocence.

The Concept of Epistemic Innocence

The common saying 'ignorance is bliss' encapsulates the familiar idea that epistemically suboptimal cognitions can cause psychological benefits. Consider the case of motivated delusions. Delusions, in general, are very difficult to define without controversy, but one relatively neutral description holds that they are "strange beliefs¹⁴ which appear in the context of mental distress" (Bortolotti 2013). Delusions may be elaborate, complex narratives accounting for much of a person's experience, as in schizophrenia, or they may be more circumscribed. A certain subclass of circumscribed delusions has been described as 'motivated delusions' on the grounds that these delusions may serve a defensive function, protecting people from the negative psychological consequences of facing unpleasant or devastating truths. Examples of motivated delusions include Reverse Othello syndrome, the delusion that one's romantic partner is faithful; erotomania, the delusion that a stranger - often a famous person - is in love with one; and anosognosia, the denial of illness (Bortolotti 2015a).

Lisa Bortolotti (*ibid.*) notes that solely motivational explanations of the formation of many of these delusions are implausible. However, irrespective of what role, if any, motivation plays in their formation, it is plausible that such delusions can be psychologically beneficial in the kinds of circumstances in which they are typically held. Bortolotti discusses a few cases, including that of 'BX', a musician whose partner left him after he was rendered quadriplegic by a car accident. BX later developed Reverse Othello syndrome, becoming convinced that not only were he and his former partner still together, but that they had married. It is not difficult to see how this delusion could have psychologically benefited BX by saving him from confronting the totality of the awful truth of his situation. He was both quadriplegic *and* single; his life as he knew it had fallen apart. Accepting this could have been utterly debilitating.

Bortolotti discusses other cases for which similar points can be made. One is that of 'LT', a reserved, socially isolated and lonely person, who had recently been through an abrupt relationship break-up and had subsequently developed erotomania. Her erotomania took the form of a delusion that a former fellow student, to whom she had never spoken, was in love with her. It is easy to see how this delusion could have provided psychological consolation given her situation. And the same point can be made for anosognosia - the denial of illness - which often emerges in the context of a traumatic or sudden injury.

When we think about such cognitive conditions as these delusions, it is natural to assume a fairly straightforward 'trade-off' picture according to which, whatever psycho-emotional benefits they may have, their only epistemic effects are detrimental. They may increase a person's happiness or well-being, but they only ever decrease a person's knowledge. Bortolotti argues that often the picture is more complicated. She notes that in addition to the kinds of psychological benefits described above, delusions very often have detrimental psychological effects, putting strain and stress both on the agent's internal cognitive processes and on their relationships with others. The epistemic picture, she says, can be similarly mixed: despite their obvious epistemic costs, such cognitive conditions can deliver significant epistemic benefits which are otherwise unavailable to the agent at that time.

Bortolotti introduces the concept of 'epistemic innocence' to articulate this idea. She proposes two conditions which an imperfect, or epistemically sub-optimal, cognitive process must meet to nonetheless qualify as epistemically innocent. The first of these is the 'Epistemic Benefit' condition: "The [sub-optimal cognitive process] confers a significant epistemic benefit to the agent at the time of its adoption"; and the second is the 'No Alternatives' condition: "Other [cognitive processes] that would confer the same benefit are not available to that agent at that time" (Bortolotti 2015a, 496). Having defined the concept thus, she argues that motivated delusions sometimes satisfy these two conditions.

The concept of epistemic innocence is based on the legal notion of an 'innocence defence', in which an act

which would ordinarily be a criminal offence can be excused on the grounds that it was necessary to prevent catastrophe. As we will see shortly, Bortolotti thinks the epistemic status of motivated delusions is sometimes analogous to this: they prevent epistemic catastrophe by committing what would ordinarily be an epistemic crime. The epistemic offence is excused by virtue of its constituting an "acceptable response to an emergency" (Bortolotti 2015a, 495). It is worth noting that this is a richer notion than is captured by the minimal, formal definition of epistemic innocence cited above, which includes no explicit reference to the prevention of catastrophic outcomes. I think this is a strength, not a weakness, of the formal definition. The cases of psychedelic therapy which I will discuss below do not all fit the catastrophe-prevention or emergency-response model. But they do fit the more minimal conception of an epistemic cost being somewhat offset by a significant and otherwise unavailable epistemic benefit. This minimal conception has the advantage of illuminating such cases, not just those which more closely fit the original analogy. Like the cases of implicit bias-driven confabulation discussed by Sullivan-Bissett (2015), psychedelic therapy highlights the broader significance of the concept of epistemic innocence as formally defined. It is not merely a defence of epistemically costly emergency responses, but an invitation to "a more balanced view of the role of [imperfect cognitions] ... [and] a reflection on the role of contextual factors in epistemic evaluation" (Bortolotti 2015a, 490).

Let us return to Bortolotti's arguments for the epistemic innocence of motivated delusions. With respect to the Epistemic Benefit condition, the basic idea is that when these delusions cause psychological benefits, those in turn cause epistemic benefits. Insofar as the Reverse Othello syndrome developed by BX protected him from total psychological collapse, it enabled him to keep communicating and engaging with doctors, nurses and visitors. And communicating with others is one of the main methods - for BX, one of very few available methods - whereby we gain information about the world. BX's delusion may have compromised his access to knowledge about the condition of his romantic relationship, but it facilitated his access to knowledge about the course of his treatment, the weather outside, local politics, and any other topic about which he may have communicated with other people. This basic idea can be applied similarly to LT's erotomania and to cases of anosognosia. Insofar as delusions benefit a person psychologically by keeping them from despair, they thereby benefit a person epistemically by allowing them to keep engaging with other people and their environments, thereby acquiring knowledge (Bortolotti 2015a).

The No Alternatives condition is slightly less clear, since it is a substantive issue in itself what counts as the 'availability' of a cognitive process to an agent at a time. In discussing this, Bortolotti reviews empirical evidence suggesting that the belief revision capacities of patients with anosognosia are often compromised, limiting their ability to replace their delusion with the true belief that they are ill. She makes a more general point, however, which applies to all the motivated delusions discussed thus far. If it is true that these delusions provide psychological benefit by protecting against a potentially catastrophic or debilitating true belief, then it is unlikely that the relevant true beliefs would provide the same epistemic benefits conferred by the delusions. In the case of BX, for instance, a true belief about his relationship would provide one obvious epistemic benefit not provided by the delusion - that of knowing the truth about his relationship. But if this true belief were devastating, it would fail to yield all the epistemic benefits which arise from engagement with people and the environment. So it seems there is no alternative to the delusion available to BX *which would provide the same epistemic benefits gained from the delusion but without its epistemic costs*.

Epistemic Benefit and Epistemic Risk

If naturalism is true, and if mystical experience is the primary mechanism of psychedelic therapy, then it seems undeniable that psychedelic therapy is "foisting a comforting delusion on the sick and dying". However, equipped with the concept of epistemic innocence, we are now in a position to appreciate the respects in which it is not *merely* doing so. The epistemic appraisal of psychedelic therapy - like that of motivated delusion - is more complicated than it might at first appear.

When therapeutic psychedelic experiences alleviate debilitating psychiatric conditions such as alcoholism, they satisfy the Epistemic Benefit condition in the same way as motivated delusions. The recovered alcoholic is far more likely to engage with other people and the broader environment, and thus likely to acquire much more knowledge. Given the reversibility of some alcohol-induced cognitive impairment, the psychedelically recovered alcoholic is also likely to benefit from improved functioning of epistemically important cognitive capacities such as attention and working memory (Mann *et al.* 1999).

The same process of epistemic benefit will occur when psychedelic therapy is applied to other debilitating conditions such as anxiety associated with terminal illness. As a result of psychedelic therapy, patients with terminal illness experience significant and lasting reductions in anxiety and depression (Grob *et al.* 2011, Gasser *et al.* 2014). Lower levels of anxiety and depression will tend to lead to more engagement with the environment and therefore more knowledge acquisition. Successful psychedelic therapy has epistemic benefits.

There is another point worth making about therapeutic psychedelic trips, albeit one not directly relevant to their satisfaction of the Epistemic Benefit condition. It has often been noted that psychedelic subjects, unlike psychotic patients, typically retain insight; that is, they are aware that their unusual experiences have been occasioned by the ingestion of a drug (Hermle *et al.* 1992). This retention of insight does not uniformly beget epistemic conservatism: many psychedelic subjects are convinced that they have enjoyed drug-induced *but nonetheless veridical* metaphysical apprehensions. However, the important point is that non-naturalistic metaphysical conversions are not guaranteed to result even from spectacularly mystical psychedelic trips.

An instructive case is Benny Shanon (2002) the Israeli cognitive scientist, philosopher, and veteran of over 130 experiences with the South American psychedelic brew ayahuasca. It is fair to say that Shanon has experienced and been suitably impressed by a representative sample of psychedelic experience. But in his monograph on ayahuasca he treads a commendably cautious line, rejecting various paranormal claims associated with psychedelics, while flirting with a pragmatic or poetic approach to the idealistic pantheistic metaphysical vision often encountered on this drug. Shanon is living proof that it is possible to experience the full power of psychedelic mysticism in an epistemically responsible way - though of course he is unusually well equipped to do so by virtue of his professional training. It can bring further balance to our epistemic evaluation of psychedelic therapy to note that, while it may be an epistemically risky treatment modality from a naturalistic standpoint, the risk is not absolute.

This mention of epistemic risk highlights an interesting feature of the case of psychedelic therapy. The kinds of imperfect cognitions which have thus far received attention in the epistemic innocence literature - motivated delusions (Bortolotti 2015a), elaborated and systematized delusions (Bortolotti 2015b), implicit bias and confabulation (Sullivan-Bissett 2015) - are all involuntary. As such, the concept of epistemic innocence is applicable to the objective assessment of the epistemic status of these cognitions, as well as to practical questions about how others (e.g. clinicians) ought to interact with the subjects of such cognitions. The present case differs in that introducing the notion of epistemic innocence could perhaps help individuals deciding whether or not to *voluntarily* enter an epistemically detrimental or risky cognitive state.

Imagine a naturalistic, atheistic terminal patient who has been offered a place in a psychedelic therapy trial. A second- or third-hand acquaintance with the phenomenology of psychedelics might lead this patient, reasonably enough, to share Pollan's worry that they would be signing up to acquire a "comforting delusion". To appreciate the point, consider Pascal's Wager, which says we ought to try to bring about theistic belief in ourselves, because belief has higher expected utility than disbelief. That is, belief has vast advantages if God exists but only moderate disadvantages if God does not exist, while disbelief has vast disadvantages if God exists but only moderate advantages if God does not exist. Now, one possible response, in line with the 'evidentialist' tradition derived from the work of W. K. Clifford (1877), is that beliefs ought never be adopted for pragmatic reasons, but only ever on the basis of sufficient evidence. One need not hold such a strong view, however, to share the basic worry: that there is something ethically suspect (because intellectually dishonest) about adopting a belief solely for pragmatic reasons. The prospective psychedelic therapy subject might reasonably wonder if they are contemplating something a little too akin to following Pascal's advice.

Appreciating the epistemic innocence of psychedelic states could help such an epistemically responsible person to appreciate the disanalogies between psychedelic therapy and Pascal's prescription. The potential epistemic costs of psychedelics are offset not just by pragmatic benefits, but by significant and (I will later argue) otherwise unavailable epistemic benefits. It is not a simple purchase of pragmatic benefits with epistemic costs, and this makes a difference to the decision under consideration. It makes a further difference to note that, as illustrated by the case of Shanon, epistemic costs are not an inevitable result of psychedelic transformation but merely a risk.

How large a risk? It is hard to say. It would be interesting to know the statistical relationships between

therapeutic and transformative effects and the acquisition of new non-naturalistic metaphysical beliefs consequent on psychedelic ingestion, but such research has not been conducted to my knowledge. Prolonged psychosis is one epistemic (and, of course, psychological) risk of psychedelics. This risk is very low when the drugs are given in carefully controlled conditions, but is known to be higher when they are consumed in uncontrolled conditions (Strassman 1984, Johnson *et al.* 2008). Perhaps the risk of forming extravagant or unwarranted metaphysical beliefs is also higher when the drugs are consumed - especially repeatedly - in uncontrolled conditions. This possibility is suggested by some of the more colourful theories espoused by, for instance, the late Terrence McKenna (1991), countercultural icon and prolific high-dose psychedelic user. If (repeated) uncontrolled psychedelic use does carry elevated epistemic risk, then this might be another reason - in conjunction with the elevated pragmatic risk - militating against such use. At any rate, in the case of just one or very few clinically supervised psychedelic therapy sessions, significant epistemic and pragmatic benefits are highly likely, and epistemic costs not overwhelmingly likely, making the epistemic risk perhaps a reasonable one to take. Of course, the precise risk/benefit analysis will vary depending on the details of individual cases.

Returning to the epistemic benefits of psychedelic therapy, it is interesting to contrast the two recent studies of psilocybin for addiction - one on tobacco, the other on alcohol. I have argued that alcoholics who recover through a psychedelic trip will gain an epistemic benefit because their social and other functionality will be restored. The same argument cannot be made for a recovered smoker, since tobacco addiction, despite its adverse physiological consequences, does not tend to compromise socio-emotional and intellectual functioning in the same way as alcohol addiction. What are we to make of this difference?

First, I am not arguing that every therapeutic psychedelic trip is necessarily epistemically innocent - only that some are. Second, however, even the tobacco smoker who kicks the habit via psychedelic mysticism will gain a significant epistemic benefit. There are at least two ways this might occur. The first relates to the finding that psilocybin-induced mystical experiences had by healthy subjects can lead to durable increases in the personality trait of openness to experience. Even for an already well-functioning individual, it is likely that increased openness to experience will cause increases in epistemic functionality analogous to those gained by the recovered alcoholic. More openness means more engagement with the world, more curiosity, and more new experiences of more varied kinds - all potential sources of knowledge¹⁵.

The second way that the tobacco smoker might epistemically benefit relates to the psychedelic experience itself. I have argued elsewhere that psychedelic states can be a means of gaining direct knowledge of the potential of one's own mind and of the contingency of one's sense of self (Letheby 2015). On the first count, neuroscientific evidence suggests psychedelics temporarily decrease constraint or inhibition in the brain, making overall cognitive dynamics more fluid and flexible (Carhart-Harris *et al.* 2014, Tagliazucchi *et al.* 2014). This can amount to a radical and direct discovery that certain ways of thinking or feeling are genuinely possible for oneself. One mescaline subject thought the most important outcome of his first psychedelic experience was "the knowledge and certainty I now have that it is truly possible to attain to a sense of harmony with all creatures and things" (Masters and Houston 1966, 12).

On the second count, one striking feature of some psychedelic states, especially mystical ones, is the experience of 'ego dissolution', in which the ordinary sense of being a unitary, persisting self or 'I' is temporarily diminished or abolished (Pahnke 1969, Carhart-Harris *et al.* 2014). Various philosophers have defended the view that the self as we ordinarily experience it does not exist (Metzinger 2011). Moreover, attaining direct, experiential knowledge of its non-existence is a central goal of Buddhist meditation (Albahari 2014). Ordinarily, the idea of experience without a sense of self is difficult to understand and may even seem incoherent. However, some meditation practitioners and psychedelic subjects become directly acquainted with the contingency of their sense of self. The felt 'I' may be a ubiquitous feature of human consciousness, but the phenomenology of such altered states as these suggests it is not necessary - and this is a striking and significant discovery to make.

So even if alcoholics stand to benefit more, those who break a tobacco addiction via psychedelic mysticism may still derive a significant epistemic benefit, both in virtue of durably increased openness to experience, and in virtue of discoveries about their own mind made during the acute drug effects.

More generally, I conclude that psychedelic mystical experiences can cause significant epistemic benefits

which are independent of their non-naturalistic metaphysical contents. These benefits are of two kinds: direct benefits, in the form of knowledge gained during the psychedelic experience itself; and indirect benefits, in the form of improved epistemic functionality after the psychedelic experience. This establishes that these cognitive states meet the Epistemic Benefit condition. I turn now to the No Alternatives condition.

The No Alternatives Condition

Based on the evidence available to date - which is admittedly limited - the efficacy and rapidity with which controlled psychedelic use can change personality and alleviate psychiatric distress is, in many cases, unique (Young 2013). This means that when we consider the epistemic benefits gained by a psychedelically treated alcoholic as a consequence of their restored cognitive and socio-emotional functioning, there is no known alternative cognitive process which will deliver those benefits with such high probability or so rapidly. The antidepressant efficacy of the (non-classic) psychedelic anaesthetic ketamine provides a good illustration of this. Some researchers think ketamine could be important in cases of suicidal depression because it alleviates depressive symptoms within hours, rather than the timeframe of weeks or months required for standard pharmacological or cognitive antidepressant therapies (Vollenweider and Kometer 2010, 2). To put it starkly, if someone is about to take their own life, a therapeutic psychedelic experience may be the only cognitive process which will lead them to keep living and learning. One classic text on the phenomenology of psychedelics describes such a case: a subject who had decided to commit suicide before he underwent a death-and-rebirth experience in the LSD trip which was his last ditch effort (Masters and Houston 1966, 188-189).

It is true, on the other hand, that depression in particular is a condition for which rapid and effective non-psychedelic treatments - such as electroconvulsive therapy (ECT) and repetitive transcranial magnetic stimulation (rTMS) - exist (Ren *et al.* 2014). So it may be that some cases of antidepressant psychedelic therapy would not provide *uniquely* rapid access to indirect epistemic benefits consequent on restored socio-emotional functionality. These specific psychedelic states would then fail to qualify for epistemic innocence - at least, they would fail to do so in virtue of their indirect epistemic benefits. However, they might still qualify in virtue of their direct epistemic benefits. Further, depression-alleviating psychedelic states may provide uniquely rapid access to indirect epistemic benefits in some of the significant proportion of patients who do not respond to alternative rapid treatments (*ibid.*). Of course, my argument in this paper focuses on the epistemic innocence of trips occasioned by *classic* psychedelics, and it is as yet unclear whether drugs of this class have significant antidepressant potential (though see Baumeister *et al.* 2014, Osório *et al.* 2015).

When alternative rapid treatments for conditions such as depression are available, and a given patient has a good chance of responding to such treatments, does this mean that a psychedelic treatment taken by such a patient would not be epistemically innocent? Not necessarily. As mentioned in the preceding paragraph, even if alternative routes to psychedelics' indirect epistemic benefits are available, there may still be no alternative route available to their direct epistemic benefits. In this case, a psychedelic state still would yield significant and otherwise unavailable epistemic benefits; would still, that is, be epistemically innocent. This is not to say that the psychedelic treatment should therefore be preferred. Establishing epistemic innocence is the beginning, not the end, of such a practical deliberation, which must take into account all the kinds of facts mentioned in my earlier discussion of epistemic risk.

There is at least one kind of case in which a psychedelic treatment might fail to be epistemically innocent: that where the patient has already had multiple noteworthy psychedelic experiences. Such a patient might well already have the kinds of direct knowledge of the mind and the sense of self afforded by psychedelics, and in that case another psychedelic experience would be less likely to afford unique access either to indirect *or to direct* epistemic benefits. Above I speculated that repeated and uncontrolled psychedelic use might carry greater epistemic risk. Now I am suggesting that (excessively) repeated psychedelic use might promise lower epistemic benefit. Alan Watts famously said of the psychedelic experience: "If you get the message, hang up the phone" (Watts 1962). The implicit suggestion, which I am tentatively endorsing, is that there is a point at which the important knowledge to be had from psychedelics has already been gained and so the epistemic value of further use is considerably less.

I am endorsing this suggestion only tentatively because there is some evidence to the contrary. Veterans of dozens or even hundreds of ceremonial ayahuasca sessions claim still to learn something new and important

with each ingestion, or with most ingestions (Shanon 2002). What can be said conclusively is that these are the conditions that would have to be met for a particular instance of psychedelic therapy to *fail* the epistemic innocence test: not just availability of alternative routes to indirect epistemic benefit, but the lack (or low probability) of direct epistemic benefits uniquely available via the psychedelic route.

But one might wonder whether the direct benefits are in fact *uniquely* available via the psychedelic route. It could be objected that meditation practice can plausibly yield some of the direct epistemic benefits which I have argued may result from psychedelic ingestion, such as knowledge by acquaintance of the potential of the mind (Kornfield 1979, Delmonte and Kenny 1987) and of the contingency of the sense of self (Brewer *et al.* 2013). However, even if so, this method is slower and less reliable than psychedelics (Harris 2014). If an agent at a particular point in time is on the verge of undergoing a controlled psychedelic session, then they will begin to reap the epistemic benefits of doing so within a matter of hours. They would not reap the same epistemic benefits *at that time* if they decided to begin a meditation practice instead.

One worry which might arise at this point is that my discussion of the No Alternatives condition has focused not on whether the epistemic benefits of psychedelics are available via any alternative means whatsoever, but on whether they are available via any equally reliable and rapid means. This might seem to be a problematic misinterpretation of the condition as presented by Bortolotti *et al.*¹⁶ To this I respond, firstly, that the focus on the reliability or probability with which different methods deliver epistemic benefits is an unavoidable consequence of the voluntary nature of psychedelic therapy. When the concern is a practical one about whether individuals should choose to enter into epistemically costly or risky cognitive states, there is no choice but to consider the probabilities of the various costs and benefits which might be involved. The notion of an available alternative which *will* (without fail) yield the same benefits is shown to be an idealization by the forward-looking, practical nature of the present case. When reasoning about actions with uncertain future results, epistemic or otherwise, we cannot but trade in probabilities.

Secondly, the No Alternatives condition as originally formulated is somewhat ambiguous with respect to whether the alternative cognitive process *or its resultant benefits* must be unavailable at the time. If, as suggested above, meditation does provide some of the same direct epistemic benefits as psychedelics, then the objection is correct: strictly speaking, there are alternative cognitive processes *available to the agent at the time* which yield the same epistemic benefits. But such a strict interpretation of the condition may not capture what is most important in cases where one process takes hours to yield these benefits while another takes weeks, months, or years. The contrast between psychedelic and traditional antidepressant therapies shows that such a time difference can be practically as well as epistemically consequential. All this suggests either that the No Alternatives condition should be understood to refer to the availability of epistemic results, rather than of the processes which deliver them, or at least that the two possible interpretations of the condition should be borne in mind and disambiguated on a case-by-case basis.

1. Clearly the transformative process cannot be entirely transparent - subjects may be directly aware of the conscious mental states involved, but not of the neural and sub-neural mechanisms causing and/or constituting those states.
2. I take cognitive neuroscience to have explanatory, not merely descriptive, ambitions. As such, even if it were successful in thoroughly charting the neural correlates of consciousness (no small achievement) it would nonetheless have partially failed unless it could also provide a naturalistic explanatory bridge between those correlates and the qualitative character of conscious experience.
3. Strictly speaking, it is psilocin, the dephosphorylated metabolite of psilocybin, which is biologically active - but this loose way of speaking is common and harmless.
4. It could also be true even if the drug experience causes benefits only indirectly, by causing the adoption of new beliefs which themselves cause benefits. I suspect this is not the case, but I will not argue for that claim here.
5. It should be clear that the same point cannot be made with respect to some subjects' beliefs that they have glimpsed non-natural realms. Certainly, one possible explanation of the formation of those beliefs is that they are true - but that explanation is non-naturalistic and hence unavailable within the methodological parameters of my project.
6. Of course, on naturalism, such phenomenal differences as these amount to differences in neural activity. On naturalism, an altered state of consciousness is nothing other than a transient global alteration to neural information processing. My conjecture that the altered state is the therapeutic agent just amounts to saying that therapeutic results are caused by these dramatic, widespread information processing changes, rather than by some direct (e.g. intracellular) effect of the drug which (a) occurs invariably given a sufficient dose and (b) is insensitive to the details of *how* global information processing is differentially altered by particular ingestions.
7. I am grateful to Philip Gerrans for this suggestion
8. Recall: on naturalism, the experience - i.e. the altered state of consciousness - just is the alteration to global neural information processing; in this case, disinhibition of other regions consequent on DMN disintegration.
9. I do not intend to import all the details of Russell's (1910) original analysis of knowledge by acquaintance, but I take it the

- intuitive contrast with 'second-hand' knowledge by description or testimony is clear enough.
10. By 'model' I mean something like 'representation encoded by patterns of neural activation' - in this case, a representation whose contents contribute to phenomenal consciousness. I deliberately remain agnostic on details of representational formats, coding schemes, etc.
 11. 'Modal knowledge' is a philosophical term for knowledge about possibility, necessity, and so forth.
 12. For helpful feedback on earlier versions, I am grateful to Gerard O'Brien, Jon Opie, audiences at the University of Adelaide Philosophy Postgraduate Colloquium and the Australasian Association of Philosophy annual conference, and two anonymous referees for the *Journal of Consciousness Studies*.
 13. Where 'epistemic' simply means 'of or pertaining to knowledge'.
 14. Though it is highly controversial whether delusions are beliefs. Non-doxastic theorists (e.g. Gerrans 2014) argue that they are not, emphasising their non-belief-like features: they are often held with ambivalence, isolated from an agent's other beliefs, and relatively behaviourally inert. Doxastic theorists, on the other hand, emphasise the belief-like features of delusions, as well as suggesting that non-doxasticists deploy an overly strict or narrow conception of beliefs (e.g. Bortolotti and Gunn 2015).
 15. An intriguing question is whether personality traits such as openness to experience bear any interesting relation to the kinds of epistemic virtues discussed by philosophers. (I am indebted to James Morauta for raising this question.) Any attempt to draw such correspondences, however, faces difficult problems concerning how to understand the theoretical and ontological commitments of the Big Five model of personality (Miller 2014) - problems which do not arise for the mere claim that increased openness will typically be accompanied by epistemic benefit.
 16. I am indebted to an anonymous reviewer for raising this point.

Conclusion

Evidence for the transformative efficacy and possible therapeutic utility of psychedelics raises intriguing issues which deserve serious attention. Among these is the legitimate concern that we may be "foisting a comforting delusion on the sick and dying". I have applied the concept of epistemic innocence to argue that the situation is more complicated and that despite their potential (albeit not inevitable) epistemic faults, psychedelic states can confer significant and otherwise unavailable epistemic benefits, both direct and indirect. Possible direct benefits include knowledge by acquaintance of the potential of the mind and the contingency of the sense of self. Possible indirect benefits include improved epistemic functionality brought about either by restored cognitive, social, and emotional functionality, or by increased openness to experience. Thus, some psychedelic states are epistemically innocent imperfect cognitions.

This conclusion matters because the epistemic status of psychedelic therapy is highly relevant to policy debates. Some of the psychedelic scientists interviewed by Michael Pollan in the *New Yorker* seem quite unconcerned by the possibility that therapeutically efficacious psychedelic mystical experiences are non-veridical. I think they are right to be unconcerned, but not because the epistemic status of therapies is unimportant.

To appreciate the importance of epistemic issues, consider Nozick's (1974) famous 'Experience Machine' thought experiment. The hypothetical Machine provides users with a lifelong stream of virtual reality experiences as pleasurable, varied, and realistic as could be. It is a common intuition that plugging in to the Machine would be a mistake, and one explanation of this intuition is that we care not only about having pleasant experiences but about having experiences based in reality. Truth and knowledge matter. When some method of attaining well-being depends upon forming false or unsupported beliefs, or refraining from forming true or well-supported beliefs, we tend to view this as a flaw of that method.

Of course, such flaws are not inevitably fatal. This is one of the key points of the talk about epistemic innocence. The point is just that epistemic evaluation matters. So if we are to prescribe an epistemically costly method of attaining well-being, we want to do so because the epistemic costs are outweighed by the overall benefits - not because we have simply ignored the epistemic costs. The scientists interviewed by Pollan may well be correct that the potential non-veridicality of psychedelic mysticism is not a compelling objection to prescribing psychedelic therapy. But this is not because epistemic issues are unimportant. Rather, it is because psychedelic therapy is unlike the Experience Machine in the respects explored in this paper.

Finally, the case of psychedelic therapy helps advance the broader discussion of epistemic innocence. It shows that the concept illuminates a broader range of cases than just acceptable emergency responses, and also illuminates issues concerning voluntary imperfect cognitions. This in turn raises the issue of epistemic risk, and how such risk might be traded off against potential epistemic benefits, as well as pragmatic risks and benefits. Meanwhile, the contrast between psychedelics and prospective alternatives reveals the need for care in interpreting the No Alternatives condition. Most generally, this analysis of psychedelic therapy provides yet

more reason to think that the epistemic status of imperfect cognitions is more complex than it might at first appear.

Acknowledgements

I am grateful to Gerard O'Brien, audiences at the University of Adelaide Philosophy Club, the University of Adelaide Philosophy Postgraduate Colloquium, and the Australasian Association of Philosophy annual conference, and two anonymous referees for helpful feedback on earlier versions of this paper.

Naturalizing Psychedelic Spirituality

Statement of Authorship

Title of paper:	Naturalizing Psychedelic Spirituality
Publication status:	Submitted for publication
Publication details:	Under review for <i>Zygon: Journal of Religion and Science</i> . December 2016.
<i>Principal Author</i>	
Name of principal author (candidate):	Christopher Letheby
Contribution to the paper:	Devised the arguments and wrote, proofread, polished, formatted, and submitted the paper.
Overall percentage (%):	100%
Certification:	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.
Signature:	
Date:	

Introduction

One of the biggest philosophical problems of our age is how to respond to the existential anxiety caused by the 'disenchantment of the world' - the transition from a religious worldview to a naturalistic one which eschews transcendent foundations for meaning and value. Some philosophers have suggested that this problem is currently becoming more acute as the results from contemporary neuroscience filter into the zeitgeist, rendering a disenchanted naturalistic conception of human beings ever more unavoidable and undeniable, and leading to a distinctive 'neuroexistentialist' anxiety (Flanagan and Caruso forthcoming, Metzinger 2009).

Various theoretical solutions to this problem have been pursued, often in the form of efforts to 'naturalize' some central pillar of meaning or purpose such as free will or morality (Dennett 1984, Sturgeon 2006). But to (somewhat facetiously) paraphrase Douglas Adams (1979), this is odd because, on the whole, it isn't our theories that are unhappy. I want to explore the prospects for a different, more practical kind of solution, based on a remarkable empirical observation about the conditions under which human beings can regain a sense that life is worth living. I refer to the finding that classic psychedelic drugs such as LSD and psilocybin can, under conducive circumstances, occasion intense experiences leading to lasting psychological benefits - including the amelioration of existential distress resulting from terminal illness (Grob *et al.* 2011, Gasser *et al.* 2014, Griffiths *et al.* 2016, Ross *et al.* 2016). A renaissance of scientific research into an "existential medicine" (Grob 2007, 213) certainly warrants the attention of philosophers interested in the overcoming of disenchantment and nihilism.

There is a problem, however, in that some evidence suggests that the existential re-enchantment occasioned by psychedelics depends crucially on the induction of mystical experiences involving apparent encounters with transcendent non-natural levels of reality (Griffiths *et al.* 2006). If this is so, then psychedelics would seem less a means to making peace with a naturalistic worldview than a means to becoming persuaded of its falsity. In the spirit of the philosophical project of 'naturalizing spirituality', I will argue that such non-naturalistic metaphysical epiphanies are far from the whole story: that there are key elements of psychedelic spirituality which can be practiced by a naturalist in intellectual good faith, and which represent a promising path to re-enchanting the naturalistic world on its own terms.

I will begin, in section 1, by describing the neuroexistentialist predicament, and two possible kinds of solutions to it, in more detail. In section 2 I will describe the potential psychedelic solution and the apparent problem with that. In sections 3 and 4 I will draw on neuroscientific research into the mechanisms of psychedelic-induced transformative experience to argue that there is a naturalistically viable kind of spirituality to be found here. The basic idea is that psychedelics degrade the mechanisms of self-representation in the brain, in

some cases allowing subjects to see that the ordinary sense of self is a useful fiction, and in other cases freeing their attention from self-focused concerns leading to the availability of broader perspectives on life. The common thread is that a good deal of existential suffering stems from an excessively rigid focus on and reified belief in a solid and enduring self and its travails. Loosening the bonds of this egocentricity and ego identification, by psychedelics or other means, is one good idea about what naturalized spirituality might amount to, and a promising practical response to neuroexistential anxiety

The Neuroexistentialist Predicament

Naturalism, in its most basic form, is the denial of the existence of non-natural or supernatural entities such as gods, souls, spirits, and non-physical minds. It is a commonplace by now that the gradual transition from a religious to a naturalistic worldview over the past few centuries - due to the rise of modern science, and especially since Darwin - has created a crisis of meaning and value. The growing plausibility of naturalism, which leaves no room for any of the traditional foundations of value, evokes anxiety about what, if anything, could possibly replace those foundations, and how, if at all, we might come to terms with a world devoid of them. This naturalistic crisis is sometimes referred to as the *disenchantment of the world* (Dreyfus and Kelly 2011).

Owen Flanagan and Gregg Caruso (forthcoming) argue that we are currently experiencing a new and distinctive wave of existential anxiety as a result of impressive advances in neuroscience, and the popular dissemination of these findings. It is not that neuroscience has discovered anything new about how disenchanted the naturalistic world is; rather, it is painting a vivid and dramatic picture of what this disenchantment looks like, in the specific and cherished case of human beings. Part of what is hardest to swallow about a naturalistic view of the world, at least for some, is its implication that human beings are ultimately nothing special - a very smart animal, to be sure, but a mere accidentally evolved, cognitively limited, mortal animal none the less. We have no Cartesian ego, no contra-causal free will, and no eternal destiny. With every new stride in the neuroscientific mapping of the material bases of decision making, reward, emotion, and so forth, this reality becomes ever more palpable. And the popularization of neuroscience plausibly heralds the dissemination of this vision beyond intellectual and academic circles into the broader culture, threatening a wholesale and widespread crisis of meaning.

One useful way of conceiving of this situation is in terms of Wilfrid Sellars' (1963) manifest and scientific images of humankind. The manifest image is simply the potentially highly reflective and refined, but nonetheless pre-scientific, conception of what kinds of things human beings are, developed and expressed in the arts and humanities as well as in everyday talk and thought. The scientific image, meanwhile, is the theoretical conception of the nature of human beings delivered collectively by biology, psychology, anthropology and so on - including, most recently and notably, neuroscience. Sellars thought - and others since (e.g. Dennett 2013) have agreed - that dealing with the various apparent conflicts between the manifest and scientific images is one of the central tasks, if not the central task, of philosophy today. It is an important and difficult business trying to "know [our] way around" (Sellars 1963, 1) a world in which our most successful epistemic enterprise seems to keep telling us that we are something very different from what we automatically and perhaps unavoidably take ourselves to be.

Flanagan and Caruso (forthcoming) describe the manifest image as incorporating commitments to such things as transcendent foundations for meaning, contra-causal (or 'libertarian') free will, human exceptionalism, the immortality of the soul, and so forth. Doubtless there is considerable inter-individual and cross-cultural variation in the precise content of the manifest image. Nonetheless it seems plausible that most of us operate at least tacitly with an individual and collective self-conception that is significantly at odds with the disenchanted scientific image.

The philosophical attempt to grapple with this inter-image conflict is exemplified in the 'naturalizing' industry that has sprung up in the last several decades. The basic strategy is to take some component of the manifest image, such as free will, or objective moral requirements, that seems inadmissible to the scientific image, and give an analysis or account of the phenomenon in question which shows that it can be a real and unmysterious part of a purely naturalistic world. Such responses attempt to demonstrate theoretically that the problem is not as bad as it seems, because the naturalistic world is not as devoid of cherished foundations for meaning

and purpose as we thought it was.

There is nothing wrong with this kind of strategy, but an interestingly different sort of response is available. This is exemplified by the project that Flanagan (2007) has called *naturalistic eudaimonics* (from the Greek *eudaimonia*, usually translated as 'human flourishing') - an empirically based inquiry into the conditions under which human beings can flourish and live meaningful lives. This project arises from a recognition that the neuroexistential predicament, while it may have its genesis in our theoretical beliefs about how things are, is not merely a theoretical or intellectual problem - it is a practical, emotional problem, a problem of living. A reflective person who is struggling with the problem of nihilism and disenchantment in our neuroscientific, post-Darwinian age wants to know: how can I, or how can we, live well and find meaning and fulfillment, given the kinds of creatures it seems that we are and the kind of world it seems that we inhabit? Of course, one possible answer to that question is: by gaining a theoretical, philosophical understanding and appreciation of how phenomena such as morality and freedom fit into a naturalistic world. But the eudaimonic approach does not presuppose that the only or the best answers will take this form. Rather, it starts with the question: how do human beings achieve meaning, fulfillment and flourishing, when in fact they do?

This is similar in spirit to the methodology of the movement known as *neurophilosophy* (Churchland 1989). Practitioners of neurophilosophy - neurophilosophers - seek to bring the results of neuroscience to bear on traditional philosophical questions. For instance, a neurophilosopher interested in understanding the nature of morality might start her inquiry by looking at neuroscientific evidence concerning moral judgment: its alterations under various conditions, its relations to other cognitive capacities, its neural substrates, and so forth (Churchland 2011). She would then use these facts as a basis for philosophical conclusions about morality.

Here I want to suggest that a helpful contribution to Flanagan's naturalistic eudaimonics could be made by a neurophilosophical inquiry into certain kinds of experiences which are repeatedly referred to as 'spiritual', and which sometimes seem to help those who undergo them to overcome existential anxiety and find meaning in life. A golden opportunity for such an inquiry currently exists, in the form of a recent surge of cognitive neuroscience research into pharmacological agents which reliably evoke such experiences - namely, classic psychedelic drugs.

An Existential Medicine?

Patients with terminal illness undergo intense suffering of many kinds, not least of which is the sense that life is worthless or meaningless in light of their impending mortality. Powerful, effective, and reliable methods of ameliorating this existential distress have thus far proven elusive (Hench and Danielson 2008). Hence, the following quotations from terminal patients are noteworthy: "It was less about my illness. I was able to put it into perspective. ... Not to see oneself with one's sickness as center. There are more important things in life. ... The evolution of humankind for example. ... Your Inner Ego gets diminished, I believe, and you are looking at the whole ..." (Gasser *et al.* 2015, 62). And: "Dying is as usual or unusual as life itself. You cannot separate it. I simply have to familiarize myself with the idea and the process" (*ibid.*).

These patients are describing experiences and insights they had while intoxicated by the infamous drug lysergic acid diethylamide (LSD), as subjects in the first controlled study of LSD-assisted psychotherapy in over four decades. The 12 subjects in this study showed reductions in anxiety for 12 months after two LSD sessions (Gasser *et al.* 2014). Similar results were found in an earlier pilot study using psilocybin, another classic psychedelic, to treat anxiety relating to terminal illness (Grob *et al.* 2011). More recently, these results have been replicated in two larger, double-blind, placebo-controlled studies of psilocybin treatment for psychological distress in life-threatening cancer (Griffiths *et al.* 2016, Ross *et al.* 2016). As a result of these studies, the treatment of psychological distress in terminal illness is now the best-studied therapeutic application of classic psychedelics (Nutt 2016).

So, we have evidence for thinking that experiences occasioned by classic psychedelics can help people to deal with a sense of the meaninglessness of life in light of mortality. This suggests the possibility of an analogy with other kinds of existential anxiety: could psychedelics offer part of a solution to neuroexistential anxiety and the crisis of meaning resulting from the disenchantment of the world? This kind of solution might take two

forms: either the careful, controlled administration of psychedelics (in a vastly different legal, regulatory and cultural situation from that which exists today) to volunteers afflicted by neuroexistential anxiety, or the study of the mechanisms of psychedelic therapy in the hopes of finding non-pharmacological routes to their salutary psycho-existential effects.

Before I explore this possibility in more detail, a little background is in order. LSD and psilocybin are both 'classic' psychedelics, meaning that they alter consciousness mainly by agonism of the serotonin 2a receptor. Other famous exemplars of this class include mescaline (from the peyote cactus) and dimethyltryptamine (DMT, an ingredient in the South American beverage ayahuasca), both of which have long histories of traditional religious use (as does psilocybin.) I will confine my discussion here to these classic, serotonergic psychedelics, and reserve the term psychedelic for them alone, omitting other drugs which have been called 'psychedelic' on phenomenological grounds despite different mechanisms of action (e.g. ketamine, *Salvia divinorum*, and MDMA or 'Ecstasy'; Sessa 2012).

From the 1940s to the 1960s there was widespread scientific interest in psychedelics following Albert Hoffman's discovery in 1943 of the potent effects of LSD. The propensity of psychedelics to induce experiences which subjects described in spiritual or mystical terms was one factor responsible for this interest (Dyck 2008). It should be noted that another distinctive property of psychedelics is the extreme variability of their effects; as such, it is certainly not the case that every psychedelic experience is of a spiritual nature (Masters and Houston 1966). However, experiences of this broad kind seemed to keep happening unbidden, even when the drugs were being studied for quite different reasons, such as their putative capacity to mimic psychosis (Mangini 1998).

Since the early days of psychedelic science there has been a debate over the credentials of drug-induced mystical-type experiences: whether they are identical or different, equivalent or inferior, to 'genuine' mystical experiences which occur spontaneously or through religious practice (Zaehner 1958, Smith 1964, Walsh 2003). However, it is beyond doubt that many psychedelic experiences possess core phenomenological features characteristic of non-drug mystical states, such as the dissolution of the sense of self, a feeling of unity with the cosmos, ineffability, and profound joy and peace (Pahnke 1969, Griffiths *et al.* 2006). This similarity was emphasized by intellectuals such as Aldous Huxley (1954), Alan Watts (1962) and Timothy Leary (Leary *et al.* 1964), who explicitly connected the psychedelic state to the consciousness alterations aimed at by Eastern meditative practices, laying the conceptual groundwork for the psychopharmacological mysticism of the 1960s counterculture. Many psychedelic mystics of the era became dedicated practitioners, sometimes teachers, of meditation (Dass 1971, Osto 2016).

Ultimately, psychedelics became the focus of intense sociopolitical controversy due to widespread use by members of the counterculture, leading to their prohibition and the subsequent virtual cessation of human research for some decades (Dyck 2008). Since the 1990s, however, a slow but steady 'renaissance' of scientific interest in psychedelics has occurred (Sessa 2012). As well as establishing the feasibility and safety of careful and controlled psychedelic administration, this new wave of research has begun revisiting earlier claims for the therapeutic and transformative efficacy of psychedelics.

I already mentioned the four studies published to date of LSD and psilocybin in the treatment of anxiety resulting from terminal illness. Besides these, small studies have also investigated the potential of psilocybin, LSD, and ayahuasca to treat addiction, obsessive-compulsive disorder, and depression, with uniformly promising albeit tentative results (reviewed in dos Santos *et al.* 2016a). Meanwhile, studies in healthy non-patient populations have found that mystical-type experiences induced by psilocybin can lead to positive personality change lasting over a year (MacLean *et al.* 2011). One frequent result, both in these studies of healthy subjects and in studies of therapeutic applications, has been that the degree of mystical experience predicts the degree of long-term psychological benefits. This echoes results from earlier, often less methodologically rigorous mid-20th century studies (Pahnke 1969). So there is mounting evidence that classic psychedelics can durably improve quality of life by inducing mystical or spiritual experiences, bolstering the case for them as a potential remedy to neuroexistential anxiety and disenchantment.

There is a problem, however, in that such mystical experiences almost by definition involve the apparent apprehension of a transcendent, non-natural reality or 'Ground of Being' of precisely the kind that is denied by a naturalistic worldview. If inducing such experiences really is the means whereby psychedelics improve

individuals' psycho-existential situations, then it would seem that these drugs are suited not to reconcile subjects to naturalism but to persuade them of its falsity. This is a serious reason to doubt whether my proposed neurophilosophical response to the neuroexistentialist predicament can even get off the ground.

Naturalistic Entheogenics

To some - especially those who have themselves benefited from a psychedelic mystical experience - it may seem that there is little to worry about here. If psychedelic mysticism can solve our neuroexistentialist problem by showing us that there is no problem, because a transcendent foundation for meaning exists after all, then so much the better!

The problem is that those who are impressed by the theoretical arguments for naturalism will be unimpressed by this response. The arguments for naturalism strike many philosophers as very persuasive, and if these arguments are sound - if naturalism is true - then the apparently non-naturalistic dimensions of mystical experience must be illusory. I am not undertaking a defense of naturalism here, but rather asking: on the (plausible) supposition that naturalism is true, is there a way that we can existentially come to terms with that fact? It should be clear that taking non-naturalistic mystical realizations at face value cannot support an affirmative answer to this question. Becoming convinced that a fact does not obtain is not coming to terms with it.

My psychedelic response to the neuroexistentialist predicament is intended as a contribution to the project of eudaimonics, which Flanagan and Caruso explicitly define as an inquiry into the conditions of flourishing for beings "whose self understanding includes the idea that [a naturalistic world] is the only kind of world that there is" (forthcoming, 15). The key question then becomes: is it true that the basic way in which psychedelic experiences alleviate psychological and existential distress is by convincing subjects that some more comforting metaphysical worldview than naturalism is correct? I am going to argue that it is not: that there are core aspects of the psychedelic transformative process, even in intense peak experiences, which are independent of apparent non-naturalistic metaphysical apprehensions. Moreover, I am going to suggest that even some of the aspects of psychedelic experience which prompt the description *spiritual* fit this naturalistic bill.

At this point my project connects with recent philosophical endeavors to naturalize spirituality, as distinct from religious belief. The very idea of something worth calling 'spirituality' that is both separable from religion and consistent with a naturalistic metaphysics may sound strange at best, contradictory at worst. But there is a growing recognition that one common use of the word *spiritual* denotes something that is distinct from religion, that has more to do with a personal quest for meaning based on transformative practice and experience, specifically experience of a self-transcendent nature. And proponents of naturalistic spirituality have been arguing that at least some of these practices and experiences can be pursued and cultivated in a way that is compatible with both intellectual honesty (Metzinger 2014) and a commitment to naturalism (Stone 2012). My aim here is to extend these efforts into the psychedelic sphere.

The neologism *entheogen* (generating the divine within) has become popular as a way of referring to psychedelics when used specifically for spiritual purposes (Ruck *et al.* 1979, Smith 2000). It is not too much of a stretch to say that the history of psychedelic research has been witness to a tension between an 'entheogenic conception' of the drugs and a seemingly opposing 'psychotomimetic/hallucinogenic conception' favored by less mystically and more naturalistically inclined thinkers. This latter conception sees psychedelics as fundamentally agents of misrepresentation and cognitive distortion, lacking significant epistemic value; according to Nicolas Langlitz's (2012) anthropological studies of psychedelic science, such a conception is implicit in much of the language used to frame model psychosis research. But Langlitz notes a tendency to rapprochement on the part of certain psychedelic researchers who embrace what he calls a *mystic materialism*. In effect, I am going to be suggesting that such a view is legitimate, by trying to naturalize the entheogenic conception of psychedelics.

I will not be assuming or defending any specific philosophical account of naturalistic spirituality. Instead, in keeping with my neurophilosophical orientation mentioned earlier, I propose that we take a 'bottom-up' approach to understanding the phenomenon: look first at what is known neuroscientifically about the kinds of

experiences and practices in question, and try to identify the features and qualities which cause people to describe them as 'spiritual'. We can come to understand what spirituality is, and whether and how it is compatible with naturalism, by looking closely at the details of paradigm cases of putatively spiritual phenomena. And subjects' overwhelming tendency to describe them in such terms suggests that psychedelic peak experiences certainly fit the bill.

Even though I will not endorse any of the existing theories, I will note that one ground for optimism about naturalizing entheogenics is that many common features of psychedelic experiences are also common themes in theories of naturalistic spirituality. Among these are wonder and awe, especially at the natural world (Goodenough 1998), love of life (Solomon 2002), and dissolution or transcendence of the individual sense of self (Harris 2014, Simpson 2014).

Ursula Goodenough proposes that reflective contemplation of scientific descriptions of nature, along with philosophical mysteries such as why there is something rather than nothing, can evoke deeply thoughtful but also affective responses of wonder and reverence at the vastness, complexity, and improbability of the natural world and of our own existence. Part of this process involves broadening our perspective and learning to see the individual self in a much larger spatiotemporal context (Goodenough 1998, Goodenough 2001, Goodenough and Woodruff 2001). Robert Solomon's account of naturalistic spirituality echoes many of these themes, recommending a thoughtful (intellectually well-grounded) sense of love and reverence for life as it is, and a sense of existential gratitude grounded in an appreciation of the vulnerability and contingency of our own lives. Like Goodenough, Solomon (2002) suggests that a key part of the spiritual process is the broadening of perspective beyond the concerns of the individual self.

While these theories of spirituality recommend transcending the self by a broadening of perspective, Sam Harris (2014) holds that spirituality consists precisely in attaining and then deepening a direct, experiential insight into the nonexistence of the self. Of course, Harris does not deny that conscious biological organisms as constantly evolving processes exist. But the claim here, common to many mystical traditions and articulated most clearly in Buddhism (Albahari 2006), is that we tacitly and habitually take ourselves to be something more than a constantly changing process - a persisting and indivisible subject of experience sharply distinct from the rest of the world. Our acceptance of this 'self-illusion' is claimed to give rise to much psychological suffering, and as such, seeing through the illusion is prescribed as a remedy.

As Harris is at pains to emphasize, despite the stronger metaphysical commitments of some schools of Buddhism, there is nothing non-naturalistic about this core psychological claim (cf. Simpson 2014). Ultimately, I will suggest that psychedelics dissolve (or weaken) the self-illusion, and broaden our perspectives, by the same basic mechanism. Evidence to be discussed below suggests that the neurocognitive systems underpinning the sense of self are the same ones that constrain cognition in accordance with the goals and interests of that self, and disrupting these systems both diminishes the felt sense of self and liberates attention from bondage to self-centered concerns. Thus, despite their differences in emphasis, there is plausibly a mechanistic unity to the varieties of 'de-selfing' recommended in these various accounts of naturalistic spirituality.

A second ground for optimism is that some psychedelic researchers describe the entheogenic transformative process in terms which do not sound unambiguously non-naturalistic. For instance, Charles Grob, describing the use of psilocybin for anxiety relating to terminal illness, says this: "Under the influence of hallucinogens, individuals transcend their primary identification with their bodies and experience ego-free states before the time of their actual physical demise, and return with a new perspective and profound acceptance of the life constant, change ... This implicit acceptance of the inevitable cycles of life leads to a drastically altered approach to what time is remaining without the panic, fear, pain, and dependency that were previously so overwhelming" (Grob 2007, 213).

Of course, references to transcending a "primary identification" with the body, and to experiencing "ego-free states", may sound somewhat questionable to naturalistic ears. However, as I have indicated, I think that there is perfectly plausible naturalistic sense to be made of such talk. In order to establish this, I turn now to recent cognitive neuroscience research on the mechanisms of psychedelic transformation.

Mechanisms of Mysticism

Recent fMRI studies of the psychedelic state have provided a fascinating window onto the neurocognitive mechanisms whereby these substances exert their effects on consciousness. Robin Carhart-Harris and colleagues (2012) found, to their surprise, that intravenous psilocybin decreased brain activity, rather than increasing it as had previously been thought. Moreover, the decreases were concentrated mainly in the famous Default Mode Network (DMN), a network of densely connected brain regions, most active in resting task-free conditions, which has repeatedly been implicated in generating the sense of self, as well as in mind-wandering, mental time travel (the simulation of past and future events), and theory of mind (the attribution of mental states to self and others; Raichle *et al.* 2001, Spreng and Grady 2010).

The decrease in activity and connectivity within the DMN was accompanied by a global increase in the unpredictability of the patterns of functional connectivity throughout the brain, quantified by the information-theoretic construct of *entropy*. Carhart-Harris *et al.* (2014) used these findings to argue that the core mechanism of psychedelic consciousness-alteration is downregulation of the DMN, which they claim implements ego-functions and serves as a 'conductor' of the global neurocognitive orchestra, constraining the quality of cognition to maintain an efficient and adaptive functional profile. Disruption to the DMN disrupts its ability to constrain cognition and allocate resources in a goal-driven fashion, and this explains many key features of the psychedelic state such as the detaching of attention from personal concerns and the sense of 'mind expansion' resulting from a broadening and loosening of attention.

The dissolution of the ego or sense of self is a quintessential feature of psychedelic experiences, particularly mystical ones, which seems explicable on this model too. Remarkably, subjective ratings of ego dissolution in these studies correlated strongly with decreases in alpha oscillations in the posterior cingulate cortex (PCC), a key DMN hub whose activity is also decreased in states of 'effortless awareness' achieved by meditation (Brewer *et al.* 2013). A study of cortical thickness in long-term religious users of ayahuasca found thinning in the PCC relative to matched controls. This thinning correlated with the extent of prior ayahuasca use, as well as with psychometric scores for the personality trait of 'self-transcendence', and was not accompanied by any impairments to neuropsychological function (Bouso *et al.* 2015). Combined with the finding that acute ayahuasca intake increases capacities cultivated in mindfulness meditation, such as psychological 'decentering' and a nonjudgmental perspective on inner experience, this set of observations is intriguing to say the least (Soler *et al.* 2016).

Carhart-Harris *et al.* (2014) argue that ego dissolution and entropy elevation both result essentially from psychedelic-induced downregulation of the DMN and related networks. Discussing the therapeutic effects of psychedelics, they note that many of the conditions indicated - addiction, depression, anxiety, and obsessive-compulsive disorder - are characterized by cognitive rigidity. They propose that the psychedelic-induced entropy elevation shakes the cognitive system out of its rut, breaking down entrenched patterns (i.e. habits) of thought, feeling, and perception, creating the possibility of forming new, more adaptive patterns by potentiating novel and under-utilized pathways. This claim is further supported by the recent finding that the degree of entropy elevation induced by LSD predicted the magnitude of increases in the personality trait of Openness to Experience displayed by healthy subjects a fortnight later (Lebedev *et al.* 2016). This is the same personality trait that was increased following psilocybin-induced mystical experiences (MacLean *et al.* 2011).

Many questions remain unanswered in the cognitive neuroscience of psychedelics. Not all neuroimaging studies have produced consistent results (see dos Santos *et al.* 2016b for a review). However, the theoretical speculations I have just described represent a plausible synthesis of current mechanistic knowledge which is consistent with phenomenological observations and provides a useful basis for thinking about the spiritual dimensions of the experience.

Consider ego dissolution, a prominent element of mystical experience. As I mentioned earlier, it is a plausible naturalistic view that the self or 'I' as we ordinarily experience it does not exist - that the sense of self is a useful fiction created by the brain (Metzinger 2003). Buddhism contends that our belief in the reality of this fictitious entity, and our subsequent attachment to its fortunes, is the source of all psychological suffering (Albahari 2006). The DMN has repeatedly been implicated in both mind-wandering and the sense of self, and empirical studies have shown that time spent mind-wandering, and functional connectivity within the DMN,

both predict unhappiness. Matthew Killingsworth and Daniel Gilbert (2010) used an experience-sampling methodology implemented by a smartphone app to track participants' self-reported levels of mind-wandering vs. mindfulness, and happiness vs. unhappiness, throughout the day. They reported that mind-wandering was both extremely common and strongly correlated with state unhappiness. Meanwhile, Luo *et al.* (2015) found that high levels of functional connectivity within the DMN as revealed by resting-state fMRI were correlated with trait unhappiness as determined by psychometric questionnaires (cf. Machado and Cantilino 2016). There is no need to swallow Buddhist psychology whole to entertain the view that excessive rumination about the self and its travails causes much needless distress.

Granted this much, the ego dissolution experience may well involve a genuine insight, a veridical apprehension of the fact that our sense of being a discrete and persistent 'I' arises out of a habitual process of creating models and spinning narratives about an ultimately fictitious entity (Dennett 1991, Simpson 2014). And if breaking the cycle of self-focused rumination can help to alleviate anxiety and depression in general, then it theoretically ought to work for existential anxiety in particular. After all, existential anxiety, as much as any other kind, intrinsically involves a focus on the prospects of a subject, a self, in a threatening predicament.

There is an obvious worry about the very idea of an ego dissolution experience: that it is conceptually incoherent to claim that a subject has a memory of an episode during which their sense of self was absent (Metzinger 2005). It is possible to reply to this by suggesting that, despite appearances and intuitions, the cognitive processes supporting self-representation and autobiographical or episodic memory encoding are dissociable. But further empirical research would be needed to determine this. A simpler line of response is to concede that probably some minimal sense of self remains in most if not all psychedelic experiences (Pahnke 1969, Shanon 2002). Even if that is true, clearly the dramatic alteration or diminution of the ordinary sense of self is enough to have a profound impact on a subject.

Moreover, there is reason to think that a 'mere' loosening, rather than dramatic diminution, of the sense of self can have a profound effect. In the recent study of LSD in terminal illness, full-blown mystical experiences were not generally observed. Instead, reductions in anxiety were brought about by emotional, not metaphysical, peak experiences, characterized by loosening, not disintegration, of ego boundaries (Gasser *et al.* 2015). The researchers suggest that the reductions in existential angst were brought about by a liberation of attention from its bondage to personal concerns, enabling patients to access new, broader perspectives on their predicament: "Throughout an altered basic emotional experience, loosening of ego functions combined with pronounced self-referential processing of significant (intellectual/emotional) content, patients may gain a new perspective on themselves and a reduction of ruminations and ego-centeredness" (*ibid.*, 9).

One patient in this study commented that when the sense of ego is diminished, "you are looking at the whole ... you are indeed starting to build relations with plants or with the entire living world around. You think less about yourself, you are thinking - across borders" (*ibid.*, 6).

These observations suggest a mechanism, distinct from but continuous with the mechanisms of full-blown mystical or ego dissolution experiences, whereby psychedelics could induce the kinds of experiential qualities emphasized in Goodenough's and Solomon's theories of naturalistic spirituality: broader perspectives, wonder and awe, and appreciation of life. Recall, Carhart-Harris *et al.* propose that the DMN, as part and parcel of the 'ego functions' it implements, serves to constrain the quality of cognition, allocating attentional resources to stimuli in accordance with an organism's goals. We have seen that an excessive focus on the self and its travails, existential or otherwise, plausibly underpins much psychological distress, existential or otherwise. But this excessive focus also consumes attentional resources which otherwise could be devoted to non-self-related cognitive contents, such as reflection on the natural world and the human condition.

Indeed, psychedelic phenomenology suggests that when attention is freed from a bondage to self-related concerns by the downregulation of 'selfing' systems in the brain, these are exactly the kinds of places that it tends to go. It is very common, even in non-mystical states, for psychedelic subjects to report feeling a sense of kinship with the natural world, or wonder and awe at the miracle of existence in general and human existence in particular (Shanon 2002). In his literary analysis of narrative reports of psychedelic experience, R. A. Durr (1970) argues that changes to attention in the psychedelic state lead to the same kind of (essentially imaginative) appreciation of life that is expressed in verse by the Romantic poets. When the bonds of self-concern are weakened psychopharmacologically, "looking at the whole [and] starting to build relations with

plants or with the entire living world" becomes, it would seem, not only possible but probable.

The philosopher Jesse Prinz (2014) proposes an account of the cognitive dynamics of aesthetic appreciation which seems applicable to the broader kind of appreciation of life resulting from psychedelic states. On Prinz's view appreciation amounts to a positive feedback loop between attention and wonder. Attention alights on some element of a visual object which evokes wonder, and the sense of wonder leads to increased attention to the visual element, and so forth. Perhaps by disrupting the analogous but deleterious positive feedback loop between existential anxiety and excessive attention to the existential predicament of the self, psychedelics create space for other, more beneficial attentional cycles to arise (cf. Nichols *et al.* forthcoming).

Everything I have said so far points to a central common element in mystical and non-mystical psychedelic experiences: the weakening of the ego by disruption to the DMN, whether this results in an 'ego dissolution' experience or a broadening of attention and loosening of cognition. In either case, something occurs that could reasonably be called mind-expansion, a liberation of focus from a narrow, rigid and maladaptive preoccupation with a reified self and its predicament. There is evidence that this basic process is important to the therapeutic and transformative effects of psychedelics. And while clearly compatible with a naturalistic outlook, a transformative process which involves broadening perspective beyond the individual self, and apprehending its interconnectedness, transience, and ephemerality, surely deserves the name *spirituality* if anything does.

Conclusion

The disenchantment consequent upon naturalism, particularly in its current neuroexistentialist incarnation, is a pressing problem deserving of serious philosophical attention. A plethora of theoretical solutions have been proposed, but it is worth investigating practical, psychological solutions to this problem too, via a neurophilosophical inquiry into methods of re-enchantment. The renaissance of scientific research into psychedelics, the 'existential medicine', constitutes an extremely promising avenue for this inquiry, but the apparent centrality of non-naturalistic mysticism to psychedelic transformation poses a *prima facie* obstacle. I have argued that despite appearances, entheogenic spirituality can be naturalized, by understanding that much of its transformative efficacy is a matter of breaking the spell of narrowly self-focused rumination, temporarily disrupting the mechanisms of self-representation and liberating attention from bondage to personal concerns. This disruption and liberation afford an opportunity to access new perspectives, see oneself and one's concerns in a broader context, and develop a sense of wonder and appreciation for life.

Anatomy of an Avatar: Ego Dissolution in Psychedelic Experience

Statement of Authorship

Title of paper:	Anatomy of an Avatar: Ego Dissolution in Psychedelic Experience
Publication status:	Submitted for Publication
Publication details:	Under review for Neuroscience of Consciousness, December 2016.
<i>Principal Author</i>	
Name of principal author (candidate):	Christopher Letheby
Contribution to the paper:	Contributed equally to devising the arguments and writing the paper; additionally proofread, polished, formatted, referenced, and submitted the paper.
Overall percentage (%):	55%
Certification:	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.
Signature:	
Date:	

Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

- i) the candidate's stated contribution to the publication is accurate (as detailed above);
- ii) permission is granted for the candidate to include the publication in the thesis;
- and
- iii) the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Name of co-author:	Philip Gerrans
Contribution to the paper:	Contributed equally to devising the arguments and writing the paper.
Signature:	
Date:	

Introduction

In this paper we argue that 'ego dissolution' in psychedelic experience is a window on the nature of self-awareness. Ego dissolution provides evidence that the mind uses a predictive coding strategy to represent the self as a unified entity. Under this strategy higher-level models predict that features represented at lower levels are attributes of objects. Thus, for example, in perceptual and sensorimotor feature binding the higher-level modelling of objects is used to integrate representations of features into coherent wholes. We see and feel objects, not concatenations of features, as a result of this top down process. Similarly, we argue, the self is postulated by higher-level processes as an entity to facilitate the binding or integration of information. We argue that this explains a Cartesian intuition that the self is a simple indivisible entity. In psychedelic experience these integrative processes are disrupted leading to the phenomenology of ego dissolution.

We support this idea with evidence from two related sources. The first is research by Sui and Humphreys (2015) on cognitive binding and the role of the self-model in enhancing it. The postulation of a self to which information is relevant makes a difference to the integrity of representations. However the mechanisms involved model that self as a heuristic, a way of making information 'sticky', rather than as a way of tracking the fluctuating cognitive fortunes of an actual entity. Self-awareness is the experience of cognitive processes in which these binding processes are intact.

The second is evidence provided by the phenomenon of 'ego dissolution' reported in psychedelic experience. We argue that that evidence suggests that psychedelics target mechanisms on which self-binding depends. Psychedelic experience degrades these binding processes, enabling us to experience cognition not bound by self-models. We emphasise that the 'self' which dissolves in psychedelic experience is not an actual entity or an object of perception, interoception or introspection but an entity inferred by the mind to predict the flow of

experience in and across cognitive modalities. As Thomas Metzinger (2016) put it: "There is just no entity there, no individual substance, and scientifically we can predict and explain everything we want to predict and explain in a much more parsimonious way."

We can contrast Metzinger's view with Hohwy and Michael's (forthcoming) "causal realist interpretation of the reference of intentional terms and of the self" which argues that the causal power exerted by the object(s) of self representation implies a causally efficacious entity underlying the hierarchy of integrative processes. Sui and Humphreys similarly argue that the binding power of self-representation cannot be explained by a self that is essentially a theoretical or narrative abstraction. Ultimately we will argue, on the basis of the psychedelic evidence, that Sui and Humphreys, and Hohwy and Michael, are right to reject a purely narrative view of the underpinnings of self-awareness: the self-model is deeper, more robust, and more cognitively ubiquitous than that. But we will side with Metzinger and against Hohwy and Michael in denying that anything, including the self-model, possesses the right attributes to qualify as a self. The issues are subtle since all sides agree on the integrative role of the self-model and the predictive hierarchical nature of self-modeling. We return to this issue in the final section. At this point we note only that the phenomenology of ego dissolution seems easier to explain if the subject is experiencing the disintegration of a system whose integration she normally experiences in terms of an indivisible mental (but not immaterial) substance.

We offer a diagnosis of this phenomenon in terms of self-binding. We discuss the nature of binding and self-binding, arguing that self-binding does not entail the existence of an object to which attributes are bound (though it does require the *representation* of an object, to which representations of attributes are bound.) We then turn to the main sections of the paper that explain the phenomenon of ego dissolution in psychedelic experience. We describe the phenomenology and neural correlates and show that ego dissolution is best understood as a phenomenon of unbinding. In a final section we situate our account in the conceptual framework of contemporary debates and address some puzzles and remaining issues. In brief our conclusion is that even in florid psychedelic experience the self-model is never entirely destroyed. Rather as the coherence it normally imposes degrades, we become aware that our normal experience of unity depends on a modeling process. Just as disorders of feature binding help disclose the nature of object representation, ego dissolution discloses the nature of self-awareness.

Furthermore our normal experience of unity compels the inference not just that we are a self, but that that self is a Cartesian substance. The self-model is a hypothesis of a unitary and persisting entity - a bare particular - which owns and inhabits the body, which thinks the thoughts, which feels the feelings, and which was present in past experience and will be present in future experience. It turns out that controversy over the neural basis of self awareness turns on whether the mechanisms which produce the experience of self awareness perform their integrative tasks by modeling the self as a simple indivisible substance.

Binding

The notion of binding solves a problem first discussed for visual representation. It arises from evidence that different visual features of objects, such as colour, shape, and motion, are processed in separate areas of visual cortex. These representational elements are combined into coherent, unified percepts by feature binding processes. Evidence for the existence of feature binding mechanisms comes from conditions in which they fail, such as 'illusory conjunctions' of features induced experimentally or resulting from pathology (Burwick 2014). An example might be a subject presented with a green circle and a blue square who misperceives a blue circle and a green square.

Binding can be described and explained at the phenomenal (unity of experience), cognitive (coherence and integration) and neural (mechanism) levels (Revonsuo 1999). There is no single theory of the nature of binding. Some theorists treat it as a global phenomenon involving the integration of information across large-scale networks. Others focus on local or modular integration of percepts or even perceptual elements, proposing that such modular integration is necessary before candidate representations can be globally integrated. However in the absence of a final theory it remains true that cognition and experience require the coherent functioning of spatially and temporally distributed neuronal populations to produce integrated representations.

The resurgence of interest in the phenomenon by neuroscientists in the 1990s produced different theories of binding. It has been explained in terms (i) of phase synchrony (peaks and valleys of waveforms occurring simultaneously) across distributed circuitry (Singer 1999); (ii) convergence zones: specialized circuitry that integrates signals fed forward from lower level systems (Damasio 1989). Such convergence zones could be hierarchically organized, elements of a percept being individually unified then fed forward to higher levels to be integrated into the overall representation of a perceived scene and (iii) integrative properties of layer V pyramidal cells (Bachmann 2015). Refined versions of this idea suggest that information thus bound passes through a processing bottleneck enabling it to become the object of metacognitive processes. (i)-(iii) fit well with attentional and global broadcast theories of consciousness, which make availability to executive processing the essence of consciousness. On these views unbound information cannot be the object of conscious awareness. (iv) is a Hebbian theory that proposes that binding consists in the frontally regulated construction and maintenance of transient activity patterns in frontal-posterior (especially parietal) circuitry, stabilizing and integrating perceptual representations, which then become available to executive functioning (Ballard *et al.* 1983, van Essen *et al.* 1992).

These theories (i)-(iv) can be made consistent. Conscious representation might depend on synchronous firing within convergence zones whose properties are modulated from above according to context.

Here we want to note several aspects of binding theory. It is proposed to explain the unity and stability of conscious experience and it implies a neural mechanism. It explains the unity of conscious experience in terms of functional coupling between cognitive subsystems and there is a strong implication that that such coupling requires frontal involvement. Relatedly, the disruption, alteration or disappearance of consciousness is produced by the disruption of frontal-posterior coupling (Mashour 2013). This last point reminds us that consciousness comes in different forms and that a binding theory of consciousness should help explain (ultimately) sleep, dreaming, coma, alert waking, psychosis, as well as the state we are interested in here, psychedelic experience.

We think that the phenomenon of ego dissolution is a case where understanding the nature of phenomenal unbinding can illuminate the nature of binding: in this case the binding which underlies the phenomenon of self awareness.

Binding and Predictive Processing

When the binding problem and tentative solutions were first proposed the predictive processing theory had not undergone its recent revival but we note that it provides an elegant framework for understanding the binding of information. Predictive processing theories of cognition treat the brain as a neurally instantiated processing hierarchy, with generative models at higher levels predicting bottom-up inputs from lower levels. The discrepancy between predicted and actual input takes the form of an error signal, with the magnitude of error known as 'surprisal'. Error signals are cancelled either by updating the model or by taking action (known as 'active inference') to alter the discrepant input (Hohwy 2013).

This framework provides an elegant solution to the binding problem. Regularities in sensory input over time constrain the combinations of features into objects that will be predictively successful. So the binding of features together into coherent percepts is determined by the probabilistic model which best matches present and past inputs, and the choice of probabilistic model is determined by regularities in the inputs. Binding, on this view, is essentially a kind of top-down abductive inference. If I am fielding in a game of cricket, the patterns in the inputs to my visual system are efficiently and reliably predicted by a model that posits various entities - a batter, a bat, a ball - each possessing various attributes such as colour, shape, location, speed and direction of motion. So object perception is a matter of hypothesizing an ontology that makes the flow of sensory input intelligible and predictable.

Our central idea here is that self-representation can be understood using the predictive framework in the same way as object representation. The self model is an object posited to make sense of, to unify, and to predict ongoing patterns of egocentric, salient, autobiographical experience.

The notion of egocentric experience reflects the fact that experience is spatially and temporally organized from a perspective or location such as point of view or origin of movement. For example sensorimotor control

depends on integration of sensory input at a point or points internal to the organism.

Salience refers to that fact that the organism is bombarded by information only a fraction of which is relevant to its goals and interests. Each of those goals, which range from phylogenetically ancient (maintaining homeostasis) to recent (maintaining social reputation) creates a dimension of salience which requires the mind to allocate processing resources adaptively. The ability to represent these multiple dimensions of salience depends on hierarchically organized neural circuitry, which focuses perception and cognition, biasing cognitive systems to the processing of relevant information. This salience network coordinates the interaction of neural systems that detect and evaluate relevance, bias processing accordingly and, importantly, allow us to feel the consequences as affective states.

Thus we are organisms which experience and act in the world from a unique physical (spatiotemporal) and affective orientation. In the process the mind makes models to explain the way the world makes us feel as we move through it. When we deliberate on and communicate the resultant experience we represent the integrated functioning of the egocentric and salience systems as unified through time in a single entity.

In effect the salience system links the egocentric self-models that manage the physical/bodily interface with autobiographical models which situate the organism historically. The salience system tells the organism which information matters and how it matters, moment to moment and through time. In effect the salience system is constantly creating the illusion of substantial selfhood by binding information into a representation not of the world in itself, but the world as it matters to the organism. As Seth (2013) put it:

emotion and embodied selfhood are grounded in active inference of those signals most likely to be 'me' across interoceptive and exteroceptive domains. In humans, self-related predictive coding simultaneously engages multiple levels of self-representation, including physiological homeostasis, physical bodily integrity, morphology and position, and - more speculatively - the metacognitive and narrative 'I'.

The narrative 'I' is Dennett's (1991) narrative self or Damasio's (2010) autobiographical self, which allows us to experience high level (attempted) control of the integrated functioning of the egocentric and salience systems.

We think that Dennett is correct that the narrative self is a model which abstracts from integrated functioning of a hierarchy of egocentric and salience systems to posit a simple unified entity of which those integrated features are attributes. This view is consistent with the idea that as that functioning disintegrates we will no longer feel like unified entities.

The Architecture of Subjectivity

That the salience network is activated in such a variety of conditions, ranging from basic perceptual processing to high-level metacognition, is a consequence of the way minds evolved. The mind did not evolve to range impartially over all the information available in the world and represent its causal and conceptual structure. Rather the mind naturally allocates resources to information relevant to organismic goals and interests, represented at different time scales and levels of explicitness. Simplifying, we can say that emotional processes represent relevance to goals and interests, the salience/reward system focuses attention cognition and behaviour adaptively and affective systems allow us to experience the consequences as felt body states.

A consistent finding in the case of self-referential processing is activity in a network of midline structures whose hubs are involved in emotion, motivation, salience, and the switching of attention between neutral/external and self-relevant information (Qin and Northoff 2011). These structures are not active when attention is focused 'outward' so to speak on the perceptual environment or on problem solving with an impersonal aspect. It seems clear that smooth context-sensitive activity in this midline-centred network is an essential aspect of cognitive function. Organisms need to be able to allocate cognitive resources appropriately according to whether the context requires accurate representation of the world (concrete or abstract) or the significance of the world to their goals and interests.

The cortical midline structures repeatedly implicated in self reference divide into two distinct, large-scale

functional networks. The salience network (SLN), which we have been discussing, is centered on key hub regions such as the anterior cingulate cortex (ACC), implicated in error detection and task switching, and the anterior insular cortex (AIC), implicated in interoceptive processing and conscious emotional feelings (Seth *op. cit.*). Meanwhile, the much-discussed default mode network (DMN) is centered on key hub regions such as the medial prefrontal cortex (MPFC), posterior cingulate cortex (PCC), and inferior parietal lobule (IPL; Davey *et al.* 2016).

The DMN was originally identified in neuroimaging studies of the resting brain (Raichle *et al.* 2001). It has since been found to be activated by many self-referential tasks, including 'prospection' or 'mental time travel' - the autobiographical simulation of past and future experience - and theory of mind (Spreng and Grady 2010). Of course, mind-wandering in task free conditions often concerns self-referential themes. The pioneering studies of Damasio, and much subsequent research, support the view that the MPFC is centrally involved in the attribution of personal relevance to (that is to say, the egocentric evaluation of) actual or simulated autobiographical episodes (D'Argembeau 2013). In relation to the different levels of self-modelling, the DMN is implicated in higher-level, narrative self-representation, as opposed to the more minimal embodied form of self-awareness supported by the SLN.

The PCC is a functionally enigmatic region with considerably higher metabolic activity and structural connectivity than most brain areas, leading some to describe it as the "core node of the DMN" (Davey *et al.* 2016, 390). In their analysis aimed at delineating the substrates of self-reference within the DMN, Davey *et al.* found that the optimal model was one in which "self-related processes were driven by PCC activity and moderated by the regulatory influences of MPFC." Activity in the left IPL, which is known to play a role in the retrieval of semantic information, was also part of the core self-reference network identified by this analysis.

Many intriguing findings support a key role for the PCC in self-reference. Neurofeedback studies of meditators suggest that PCC activity co-varies not with mind wandering as such, but with the experience of becoming 'caught up' in a train of thought. Brewer proposes that the PCC is involved in getting 'caught up in experience', whether it be a particularly compelling rumination or a drug craving - and numerous results find meditation downregulating PCC activity (Brewer *et al.* 2013, Brewer and Garrison 2014). Experimental results implicate the PCC in processes such as integrating information about spatial self-location and body ownership (Guterstam *et al.* 2015). According to one recent model, the PCC is a key hub involved in 'tuning' the connectivity of many other brain networks, regulating the balance between internal and external attention, as well as the breadth of attention (Leech and Sharp 2014). These observations about PCC function cohere with findings from psychedelic neuroscience which we will discuss below. First, however, we turn to the generic function of 'self-binding' which seems to be implemented by these self-processing networks.

Self Binding

The concept of self-binding comes from the work of Sui and Humphreys (2015) who proposed it to explain why cognition across domains and levels is enhanced for self-relevant information. Subjects asked to classify words on the basis of self-relatedness or meaning remember more self-related words, and moreover remember more episodic details surrounding the learning of those words, suggesting enhanced mnemonic binding. Meanwhile, in face processing studies, subjects are quicker to recognise their own face than friends' or strangers' faces, whether upright or inverted. Friends' faces also have an advantage over strangers', but only in the upright condition, suggesting that self-reference confers a unique Gestalt or integrative advantage. (Other results militate against an explanation of this effect in terms of mere familiarity).

Sui and Humphreys also cite evidence that self-reference enhances coupling between different temporal stages of information processing, possibly related to temporal binding, and enhances functional coupling between specific brain regions such as the aforementioned DMN, implicated as a neural substrate for autobiographical thought.

Sui and Humphreys also show that there is an attentional and decisionmaking enhancement effect for self-related information which speeds up sequential processing by enhancing "binding between different states of processing" (*ibid.*, 724). They noted that these effects could best be explained in terms of an integrative role for self-related processing and hence they employed the concept of self-binding. Given the necessity for

binding in establishing coherence they then examined a role for functional coupling of distributed circuitry consistently implicated in self-referential effects. Dynamic causal modelling of fMRI data showed "self-reference enhances the neural coupling between regions concerned with a core self-representation (vmPFC) and with distinct domain-specific regions associated with different components of the self, including self-related attention (LpSTS)" (*ibid.*, 724). Meta-analyses of fMRI data reinforced the finding that "there is enhanced neural coupling for self-processing (vs other-related processing) between the vmPFC/pregenual anterior cingulate (pACC) and several other regions including the bilateral anterior insula, left striatum, right thalamus, and amygdala" (*ibid.*, 724).

Further evidence that self-relevance enhances processing comes from the role of affect in perceptual binding. It is well-established that perceived objects evoke immediate affective responses. Lebrecht *et al.* (2012) found that even paradigmatically neutral objects such as coffee cups and clocks evoke measurable 'micro-valences', concluding that "valence should be construed as a property of object representations" - that is, the affective response is bound to the percept, forming a unified representation, rather than being merely associated with it. The view that affect is bound to perception is supported by studies showing the role of the amygdala in driving eye saccades in face perception, directing the gaze to emotionally salient stimuli.

At higher levels affect is bound to memories and imaginative states invoked in reverie or deliberation, a process which seems to depend particularly on the vmPFC. This suggests that the vmPFC is recapitulating at a higher level the role played by the amygdala in perceptual processing. Namely co-ordinating the construction of self-relevant representations. One nice demonstration of this phenomenon arises in a distinction drawn between episodic (the episodic representation of previously acquired information) and autobiographical memory (in which the subject feels herself present as part of the recollected experience) by Asaf Gilboa:

the Ventromedial PFC ... does so by establishing a *self-related* retrieval template that sets up the parameters against which retrieved memories are evaluated based on (or which give rise to) a general intuitive 'feeling of rightness'. Thus, the monitoring of the veracity and perhaps cohesiveness of autobiographical memories is primarily mediated by the ventromedial PFC. (2004; our italics).

For Sui and Humphreys the self plays far too important a role to be dismissed as a narrative fiction. They frame their argument in terms of the *function* of self-reference and argue that the integrative function they identify could not be performed by a fiction. They describe different elements of a self-referential processing system as constituting an "integrative hub for information processing" (Sui and Humphreys *op. cit.*, p. 719). The key nodes of the network they identify are those implicated in emotional binding and affective experience, the enhancement of attention for emotionally salient and motivating information and the switching of attention according to context.

These systems together perform the function identified by Sui and Humphreys. At perceptual and sensory levels they ensure that salient information is preferentially attended to and integrated into coherent percepts. At higher, conceptual or metacognitive levels they ensure that the same is true of the representations that figure in memories, stories, and causal explanations. In a review of the neural correlates of self-binding, which are the substrates of the network we outline above Lou *et al.* concluded "self-awareness is an integral function of all conscious experiences, binding conscious experiences together into a single construct with a sense of unity of consciousness" (in press).

This is a very strong claim. It implies that conscious experience would not arise in the absence of self-binding. The work of Sui and Humphreys suggests that this is not the case. Their subjects did not lose consciousness when perceiving self-irrelevant information. The same is true for patients with lesions to components of this network. However such patients do seem to have compromised abilities in tasks (such as autobiographical simulation, personal decision making and social interaction) that we pretheoretically describe as requiring self-awareness. As we will show, the neuroscience and phenomenology of psychedelic ego dissolution put further pressure on the view that self-binding is necessary for consciousness. Indeed, psychedelic consciousness, in which the mechanisms of self-binding are compromised, is usually described as a state of increased or intensified consciousness, leading Carhart-Harris *et al.* (2014) to suggest that self-referential processes constrain consciousness rather than enabling it.

Recognizing the role of self-binding and self-models leaves questions about the mechanisms of binding unanswered. As we noted earlier many candidates have been proposed, including various forms of temporal correlation (synchronisation of cells or functional coupling of brain regions; Singer 1999), convergence of modular input streams in high-level integrative brain regions (Damasio 1989), and even intracellular processes implemented in cortical pyramidal cells (Bachmann 2015).

At this point evidence about the actions of psychedelic drugs and the experience of ego dissolution becomes relevant. Under the influence of psychedelics we continue to represent the world and our bodies but attention is captured and allocated in different ways, experience seems both more intense and less personal, and salience, affective feeling and motivation become detached from personal goals and history.

Psychedelic Ego Dissolution

'Classic' (serotonin 2a receptor agonist) psychedelics such as mescaline, psilocybin, and dimethyltryptamine (DMT, a key constituent in the South American beverage ayahuasca) have a long history of religious, spiritual, and medicinal use in various cultures (Sessa 2012). Serious scientific interest in these substances was sparked by the accidental discovery in the 1940s of the extremely potent psychoactive effects of lysergic acid diethylamide (LSD). During the 1950s and 60s, psychedelics were studied intensively as models of psychotic and mystical states, as psychotherapeutic agents, and as instruments for mapping the varieties of human consciousness (Osmond 1957). However, uncontrolled use of the drugs in the context of the hippie counterculture led to socio-political controversy, culminating in the virtual cessation of human psychedelic research for some decades.

Since the early 1990s this research has been slowly but steadily resuming, with greater methodological rigour and aided by the many new technical and theoretical innovations in the mind and brain sciences. Several small studies have revealed preliminary evidence for the safety and therapeutic efficacy of carefully controlled psychedelic administration (dos Santos *et al.* 2016a, Griffiths *et al.* 2016, Ross *et al.* 2016). Meanwhile, neuroimaging investigations have begun to identify neural correlates of the changes in cognition and consciousness caused by psychedelics (dos Santos *et al.* 2016b).

The remarkable variability and alleged ineffability of psychedelic experiences pose barriers to describing these states. But some progress has been made on charting the phenomenology of the 'antipodes of the mind' (Huxley 1954, Masters and Houston 1966, Shanon 2002). Psychedelics alter many aspects of experience: sense perception, emotion, cognition, and the apprehension of time and space. But of all their effects, perhaps none is more provocative than the profound alteration to the ordinary sense of self or ego known as 'ego dissolution'. The countercultural icon Alan Watts, responsible for popularising both Eastern religion and psychedelics in the 1960s, wrote: "there are certain types of change which are usual enough to be considered characteristic of psychedelics: the sense of slowed or arrested time, and the alteration of "ego boundary" - that is, of the sensation of one's own identity" (Watts 1964). Psychedelic subjects often report that their sense of being a self, or 'I', distinct from the rest of the world 'out there', is weakened, altered, or abolished during the intoxication.

Ego dissolution experiences often occur in the context of mystical states in which the ordinary sense of self is replaced by a sense of union with an ultimate reality underlying all of manifest existence - the famous 'cosmic consciousness' experience. Shanon describes the content of this apparent metaphysical realisation as "idealistic monism with pantheistic overtones" (2002, 163). The propensity of psychedelics to occasion such experiences goes some way to explaining their history of religious use. Indeed, intellectuals such as Watts and Aldous Huxley were initially drawn to psychedelics due to their prior interest in mysticism. With respect to ego dissolution, it is worth noting that apprehending the non-existence of the individual self is a central goal of Buddhist meditation (Albahari 2014). There is evidence that mystical states are important for the therapeutic effects of psychedelics (Garcia-Romeu *et al.* 2016) so explaining the ego dissolution experience is a crucial step in theorising the mechanisms of psychedelic treatment.

Early theoretical speculations about psychedelic ego dissolution strikingly anticipated the key ideas we will develop in the rest of the paper. In accordance with the dominant orientation of mid-20th century psychiatry, Freudian constructs were used liberally - some of which have recently been revived and linked to the

predictive coding framework in the context of the 'neuropsychanalysis' movement (Carhart-Harris and Friston 2010). However, Freudian commitments are inessential to the central points here. Savage (1955) and Klee (1963) both suggest explanations of ego dissolution in terms of failures of mechanisms of cognitive integration, and view these compromised perceptual (particularly interoceptive) integrative mechanisms as the basis of the self-model: "Lacking a reliable inflow (or integration) of stimuli, particularly from his body, the subject has lost much of the basis for his self-percept" (Klee 1963, 465). Both authors also liberally refer to the breakdown of cognitive processes of expectation (cf. prediction) in explaining psychedelic ego dissolution. These early speculations sit very comfortably with currently influential theoretical frameworks in cognitive neuroscience, and suggest a bidirectional, mutually reinforcing relationship between self-representation and cognitive integration.

Another important point is that ego dissolution is not an all-or-nothing affair. Different aspects of self-awareness may be more or less disrupted in different ways on psychedelics. Various authors have suggested that some vestiges of self-awareness are preserved in most, if not all, psychedelic experiences (Pahnke 1969, Shanon 2002), which helps explain the puzzling fact that autobiographical memories can apparently be formed of these putatively selfless episodes (Metzinger 2005). Moreover, there are aspects of psychedelic phenomenology which may not be explicitly described as ego dissolution but are still readily explained by our account. Subjects often find their attention drawn to stimuli which they normally would not notice; as Watts puts it, psychedelics "make the spotlight of consciousness a floodlight which... brings to light unsuspected details - details normally ignored because of their lack of significance" (Watts 1964). Attention is no longer guided exclusively by adaptive and egocentric goals and agendas; salience attribution is no longer bound to personal concern.

Subjects often become less defensive and better able to view their own thoughts and feelings dispassionately - one main rationale for early therapeutic use of these drugs (Eisner and Cohen 1958). A very common experience is to see one's own dysfunctional emotional or behavioural patterns, and the possibility of alternatives, with striking clarity (Shanon 2002, 162-163). There is a tendency toward decentring and the objectification of self-related phenomena which ordinarily are taken very personally and evoke strong emotional reactivity (Soler *et al.* 2016). All of this suggests that, even when ego is not felt to dissolve altogether, the contents of consciousness are less filtered through considerations of self-relevance than is usual. Various cognitive processes which are usually tightly coupled to self-representation become decoupled from it.

Ego Dissolution as Unbinding

When Carhart-Harris *et al.* (2012) published the first human fMRI studies of serotonergic psychedelics their results caused quite a stir. Not only did they find that intravenous psilocybin caused an overall reduction in neural activity (measured by BOLD signal), apparently contradicting earlier PET studies showing metabolic hyperactivity on psilocybin (Vollenweider *et al.* 1997). The decreases were concentrated mainly in the DMN. Psilocybin decreased activity in DMN regions but also the integrity of the network as measured by functional connectivity analyses. The normal functional coupling between key DMN hubs diminished - in particular, between the MPFC and the PCC. Despite apparent inconsistency with PET results, the fMRI results cohered with MEG investigations in the same study. These electrophysiological measures revealed another striking result: decreased alpha power (which has been linked to inhibitory processes) in the PCC correlated with psychometric ratings of ego dissolution.

Not all of these findings have been replicated in all subsequent psychedelic fMRI studies. But various observations support a role for downregulation of the DMN, and the PCC in particular, in psychedelic ego dissolution. Palhano-Fontes *et al.* reported decreased BOLD signal in DMN nodes including the MPFC and PCC during ayahuasca intoxication (Palhano-Fontes *et al.* 2015). Meanwhile, thinning of the PCC has been observed in long-term ayahuasca users, with the degree of thinning correlated with both psychometric scores for 'self-transcendence' and the extent of prior ayahuasca use (Bouso *et al.* 2015). Interestingly, these ayahuasca users equaled or outperformed matched controls on tests of neuropsychological function, suggesting that this PCC thinning is not linked to cognitive impairment. The connection with findings about PCC deactivation in 'effortless awareness' meditation is obvious, and bolstered by the finding that acute ayahuasca intoxication increases mindfulness-related capacities (Soler *et al.* 2016).

Matters have been complicated by recent apparently conflicting findings concerning the neural correlates of ego dissolution. Lebedev *et al.* (2015) reported that disintegration not of the DMN but of the SLN correlated with psilocybin-induced ego dissolution. Referring to the apparent conflict with prior results, Lebedev *et al.* note that while the DMN is frequently linked to the narrative self, the salience network has been linked with a more minimal or embodied sense of self. The distinction between narrative and minimal selfhood is commonplace in recent literature and will form part of our explanation of this pattern of results too.

Lebedev *et al.* also found ego dissolution significantly correlated with reductions in interhemispheric connectivity and in functional coupling between the medial temporal lobes (MTL) and cortical regions. This connects with earlier findings of MTL-DMN decoupling under psilocybin (Carhart-Harris *et al.* 2014) - though in the earlier studies, this decoupling did not significantly predict ego dissolution, while in the Lebedev *et al.* study, the MTL decoupled from broader cortical regions than just the DMN.

Most recently, Tagliazucchi *et al.* (2016) found that ego dissolution induced by LSD correlated with increased global functional connectivity, with these increases resulting from a breakdown in the integrity of resting-state networks. Here the DMN and SLN were both affected. The strongest specific correlation with ego dissolution was decoupling of the parahippocampal cortex, in the MTL, from the precuneus - a key DMN node anatomically adjacent to the PCC. In this study disintegration of the DMN also tracked reduced mental time travel to the past, consistent with this network's hypothesized role in this process (Speth *et al.* 2016).

What are we to make of these varied results? The major networks whose breakdown is linked to psychedelic ego dissolution have all independently been linked to one aspect or another of self-representation. Lebedev *et al.* suggest a need for fine-grained psychometric instruments capable of distinguishing between alterations to different aspects of self-awareness. Unfortunately, the ego dissolution inventory recently introduced and validated by Nour *et al.* (2016) does not make such discriminations, so for now we are stuck with theoretical speculation. However, our account will predict that Lebedev *et al.* are correct: different aspects or levels of the self-model are disrupted by psychedelics in different conditions, compromising the integrity of distinct self-binding processes.

We submit that an elegant resolution of these issues is as follows. Sui and Humphreys are correct that the self is more than a mere narrative posit. There is a robust self-representation implemented in cortical midline nodes of the DMN and SLN which plays a causal role in processing at all levels, both enabling and emerging from cognitive binding processes. This model includes both the "pre-reflective experience of "I" as a continuously existing being" (White 2015) as well as the various perceptual, embodied, and narrative representations of attributes, goals, and events which are bound to that being. The plastering of these representations with "sticky 'I, me, mine' labels that... prove painful to remove" (Kingsland 2016, 285) explains the anxiety and dysphoria which sometimes attend ego dissolution experiences (albeit relatively rarely in controlled settings.) The multi-layered nature of the self-model explains the apparently variable neural correlates of the same coarse-grained subjective report of "ego dissolution". Savage suggests that on LSD "[changes] in body ego feeling usually precede changes in mental ego feeling and sometimes are the only changes" (1955, 11), evincing an early clinical recognition that different aspects of the self-model might be affected to varying extents by psychedelics.

However, it does not follow from the existence of a robust, causally efficacious, and multi-layered self-model that the entity represented by this model exists. We agree with Metzinger that it does not, and reject Dennett's identification of the self with narrative self-models as well as Hohwy and Michael's identification of the self with hierarchical predictive self-models. This is because the self-model does not posit the existence of a narrative or a model; it posits the existence of a substance, a prime mover, a persistent entity that underlies, owns, and initiates thoughts, experiences, and actions. In accordance with predictive coding principles of cognitive efficiency, the postulation of a bare particular to which events and attributes are bound is an economical means to prediction and explanation of the flow of information.

Hohwy and Michael (forthcoming) argue that the self-model qualifies as a self because it performs many of the functions attributed to the self, such as being a "hidden cause" of thought, emotion, and behaviour. It is true that the self-model does these things. But on our view it is more like an existential placebo than a successfully self-referential model. The expectation that a drug treatment will work is a mental representation which is capable of performing some of the functional role which it attributes to its hypothesized object - viz., the

pharmacological capacities of a chemical. But this does not mean such an expectation is a pharmacological capacity of a chemical. It does some of the right stuff, but it is the wrong kind of thing.

Analogously, consider a theist who mentally represents God as the source of meaning and purpose in her life. Plausibly her mental model of God as an objectively existing and meaning-providing entity will cause her to have experiences of meaning and purpose in life - that is, it will perform some of the functional role that it attributes to its hypothesized entity. But this does not mean the theist's God-model qualifies as God. Again, it does some of the right stuff, but it is the wrong kind of thing. It is not a transcendent deity; it is a system of mental representations in someone's head. One way of putting the point, then, is that the concept SELF is not purely functionally defined but has crucial *substantial* metaphysical commitments; playing part of the functional role of a self is an insufficient qualification for selfhood.

Our argument, then, turns crucially on the claim that the content of the self-model is that of an enduring substance which is distinct from all the mental, emotional, and physical activity of which it is supposed to be the owner. As Metzinger (2003) points out, our 'phenomenal avatar' (conscious self-model) is *transparent*: one does not feel like an avatar encoded in a "biological data format", one feels like a unitary, persisting substance or entity. The prevalence of this Cartesian intuition in philosophical discussions of selfhood (as exhibited by standard responses to Parfit's famous teleportation thought experiments) supports this psychosemantic claim. Some may find this claim introspectively dubitable. But it is precisely the ubiquity of this sense of 'I' that makes it difficult to isolate phenomenologically - and this is the reason why phenomena such as ego dissolution are theoretically valuable. As Savage says in his discussion of LSD phenomenology, "By and large the individual is not aware of the ego boundaries of his mind and body and becomes aware of them only when a change has occurred in them" (1955, 6). Psychedelics, by deconstructing the avatar, render it opaque and acquaint subjects directly with its representational nature. The transformative existential shock which often attends this discovery testifies to the fact that a mere avatar is not what we, in the ordinary and sober course of things, deeply feel and take ourselves to be.

Conclusion

Dennett (1992) famously argued that the self is a 'centre of narrative gravity' akin to the centre of gravity of an object: an abstraction, but no less real or predictively useful for that. On the predictive coding construal of the integrative self, the self-model functions as a centre of *representational* gravity - not just in narrative processes but also in lower-level affective, bodily, and spatial processes. It is a representation of the 'point' around which everything else revolves. But crucially, it is not a representation as of a mere point, but as of a particular. This is why we reject Dennett's metaphysics of selfhood, as well as Hohwy and Michael's claim that the self-model has enough of the right attributes to qualify as a self.

If our conjectures are right, the content of the self-model is Cartesian: it is of a substance or an entity which *has* the properties and experiences. And here we agree with Metzinger (2016) that there "is just no entity there, no individual substance, and scientifically we can predict and explain everything we want to predict and explain in a much more parsimonious way." Thus, our view is structurally analogous to Mackie's (1977) error theory of morality, conjoining the psychosemantic claim that the self-representation is as of a Cartesian substance with the ontological claim that no such substance exists. It is also substantially similar to Metzinger's views about the self-model, developed in the context of a specific case and a mechanistic proposal (self-binding on predictive processing principles) about how the model performs its adaptive functions for the organism.

All of this fits smoothly with the idea that the DMN and SLN implement narrative and embodied aspects of self-representation respectively. The self-representation as we conceive it serves as an organising principle at different levels and in different domains of processing. Perceptual representations are organised in an egocentric space; interoceptive representations are interpreted as signals of adaptively relevant events or 'core relational themes' (Prinz 2004); and narratives are structured around the fortunes and prospects of a protagonist. At all levels, salience is attributed, attention directed, and information integrated in accordance with the relevance of information to the organism's goals. As James Kingsland (2016, 209) puts it, "we have evolved into an ape that takes things personally".

Conclusion

In this final section, I summarise the arguments given in the four papers, and suggest directions for future research. The arguments given have made significant inroads into the philosophical analysis of psychedelic phenomena, but many unanswered questions remain, most of which will benefit from further results of the ongoing scientific study of psychedelics and their transformative effects.

Summary of Arguments

Overall, I have argued that the 'entheogenic conception' of psychedelics as effective agents of knowledge gain and spiritual growth is both consistent with naturalism and plausible in light of current scientific knowledge.

In the first paper, after arguing for the importance of a naturalistic approach and reviewing the history and transformative effects of psychedelics, I used Craver's account of causal relevance to argue that four lines of evidence implicate the psychedelic ASC in the transformative process: the apparently epistemic phenomenology of psychedelic transformation, the analogy with non-drug-induced transformative experiences, the correlations between phenomenal variables and psychological benefit variables, and the 'de-rigidifying' theory of psychedelic therapy. From the causal relevance of the ASC I concluded that psychedelic transformation is a distinctively 'humanistic', that is meaning-involving, psychopharmacological intervention. I then suggested three kinds of epistemic benefits that plausibly result from psychedelic transformation. These were: knowledge by acquaintance of the mind's potential, suggested by the de-rigidifying theory of psychedelic therapy; knowledge by acquaintance with the metaphysical nature of the (sense of) self, suggested by the phenomenon of ego dissolution; and revitalized capacities for the acquisition of modal knowledge, suggested by claims of increased cognitive flexibility following psychedelic ingestion (and again, bolstered by the de-rigidifying theory.)

In the second paper, after once again reviewing the history and science, I used Bortolotti's notion of epistemic innocence to argue that psychedelic states sometimes have the potential to cause significant and otherwise unavailable epistemic benefits to subjects. These epistemic benefits include the three kinds argued for in the first paper, plus indirect epistemic benefits consequent on psychological benefits. By freeing subjects from debilitating forms of psychological distress, and increasing their Openness to Experience, psychedelics plausibly result in a more pro-epistemic behavioural profile characterised by increased curiosity, positive mood, and engagement with the world. I suggested that the epistemic innocence of psychedelic states is highly relevant to policy debates because, even if they are not the only thing that matters, truth and knowledge do matter; the concern that psychedelic therapy may be "foisting a comforting delusion on the sick and dying" is real and legitimate. I also argued that the analysis of psychedelic states reveals an ambiguity in the formal definition of epistemic innocence; namely, it is not clear whether what is required for epistemic innocence is the unavailability at-a-time of alternative cognitive processes leading to equivalent epistemic benefits, or of the beneficial epistemic *results* of those processes. I suggested that the distinction between these two possibilities ought to be borne in mind in future research into the complex epistemic profiles of intuitively 'imperfect' cognitions.

In the third paper, I explored the potential of psychedelic experiences to alleviate the 'neuroexistentialist' anxiety consequent on the disenchantment of the world, arguing that their capacity to alleviate other kinds of existential distress provides a prima facie case for investigating this potential. I noted the problem arising from the apparent centrality of non-naturalistic metaphysical apprehensions to psychedelic transformation. I then used current theoretical knowledge about the mechanisms of psychedelic transformation, and phenomenological reports of psychedelic subjects, to argue that such metaphysical apprehensions are far from the whole story. I proposed that there are central means whereby psychedelics can combat existential distress that are naturalistically acceptable; namely, disruption to processes of self-representation, leading to a decoupling of attention from personal concerns and the consequent availability of broader perspectives, leading to wonder and appreciation for life and a diminution of excessive and distressing focus on the existential plight of the vulnerable individual self. I also suggested that such neurophilosophical investigations of psychedelic-induced re-enchantment can make a valuable contribution to the related philosophical projects of naturalistic eudaimonics and naturalizing spirituality, which collectively constitute a promising and distinctive kind of response to neuroexistential anxiety. I pointed out that the central means of psychedelic re-enchantment

which I identified are common themes in philosophical accounts of naturalistic spirituality

In the fourth and final paper, Philip Gerrans and I argued that psychedelic ego dissolution can best be explained as the result of a neuropharmacological disruption to 'self-binding' processes implemented in accordance with predictive coding principles. We suggested that hierarchical predictive models unify and integrate information in various domains and at various levels of processing by positing a discrete and unitary entity - a 'bare particular' or Cartesian substance - to which representations of attributes are bound. The psychosemantic claim that we represent ourselves as such a particular explains both prevalent Cartesian intuitions in philosophical discussions of self-awareness as well as the transformative existential shock that attends the psychedelic-induced realization that no such particular in fact exists. On the basis of this theoretical account we sided with Metzinger and against other theorists such as Hohwy and Michael in arguing that the neurocognitive nature of self-awareness and alterations thereto support an error theory of self-representation; even if predictive models themselves perform some of the functional role attributed to the entity which they posit, they do not have the right attributes to qualify as a self, because the self-model posits a unitary, persistent and indivisible entity. Self-models play some of the roles that a self would play, but they are the wrong kind of thing. This theory of ego dissolution contributes to naturalizing the entheogenic conception by showing that psychedelic experiences can afford subjects veridical and transformative insights into the very nature of their own existence.

Directions for Future Research

Several questions in the philosophy of psychedelic transformation remain to be addressed, either at all or in more detail. For instance, in the first paper, I gave reasons for thinking that epistemic gains such as knowledge by acquaintance with the potential of the mind and with the constructed nature of the sense of self are causally involved in the psychedelic transformative process. But more work needs to be done to show precisely how such epistemological concepts map onto specific components of the transformative mechanism.

This is an area where the ideas from philosophy of science about causal relevance and multi-level explanation, briefly touched on in the first paper, could be extremely helpful. Theoretical investigation into psychedelic transformation has continued to emphasise a common mechanism operating across different conditions: downregulating connector hubs, thereby rendering the brain more flexible and allowing the potentiation of more adaptive pathways. As I argued in the first paper, certain considerations from the philosophy of science suggest that a purely dynamical interpretation of this theory needs to be supplemented with mechanistic details to be genuinely explanatory. So far, so good. But it is still an open question whether those mechanistic details need refer to specific (kinds of) representational contents.

Nichols *et al.* (forthcoming) emphasise changes to whole brain dynamics, as mechanistically realized in changing connectivity between hub regions and other networks, as a common mechanism in different cases of psychedelic transformation. They suggest that the mystical experience might simply be an epiphenomenon of this transformative mechanism, rather than causally involved itself. Perhaps knowledge by acquaintance with the mind's potential is still plausibly part of the mechanism even on a reasonably content-independent story such as this. But clearly these questions need further analysis and (like most of the unanswered questions I will mention) will become even more tractable in light of further developments in psychedelic science.

Meanwhile, the notion of epistemic risk discussed in the second paper could be developed further, both in relation to psychedelics and other cognitive conditions with mixed epistemic profiles. The idea of assessing a pharmacological or other intervention in terms of its physiological, psychological, or social risk/benefit ratio is familiar enough, but if my arguments in this paper are sound, then epistemic risk/benefit ratios may constitute an important and relatively understudied dimension of an overall risk/benefit analysis of psychiatric treatment modalities. Biddle (2016) suggests that the notion of epistemic risk can help illuminate issues in philosophy of science relating to theory confirmation. It would be an interesting project to see if the idea as he develops it could help in further analysis of the risks involved in voluntarily entering epistemically complex cognitive conditions such as the psychedelic state.

It should be noted that empirical data are clearly relevant to epistemic risk assessment too. I have cited the case of Benny Shanon as showing that the epistemic risk of psychedelics from a naturalistic perspective is not

absolute - it is possible to undergo psychedelic mystical experiences while subsequently retaining a critical attitude toward the apparent metaphysical apprehensions occurring therein. But it would be good to have some rigorous data on the likelihood of doing this. If it is genuinely a concern that psychedelic therapy might be "foisting a comforting delusion on the sick and dying" - and I have argued that it is - then it would seem important to conduct quantitative research into the frequency with which psychedelic treatment results in patients acquiring new beliefs which must be regarded as delusive from a naturalistic perspective, and also into the relevance or otherwise of such doxastic changes to clinical outcomes.

Such research could also help make more progress on issues concerning naturalizing psychedelic spirituality, discussed in the third paper. In that paper I argued that mystical metaphysical apprehensions are not the whole story, and that key elements of the psychedelic transformative process are both independent of such apprehensions and consistent with a commitment to naturalism. I think the evidence which I synthesized provides a strong theoretical case for this claim, but it ought to be subjected to direct empirical scrutiny. Quantitative studies into the relationships, if any, between psychedelic-induced metaphysical conversions and psychological benefits could give us much greater confidence regarding the respects in which and extents to which psychedelic transformation is a matter of metaphysical consolation or, alternatively, of naturalistic spiritual transformation. Such research might also be able to support or undermine the suggestion made by Nichols *et al.* (forthcoming) that the psychedelic mystical experience is essentially an epiphenomenon of the dismantling of rigid connectivity patterns by downregulation of connector hubs.

My proposal in the third paper that psychedelics can help overcome existential anxiety in non-patient populations merits empirical scrutiny too. There is certainly extensive anecdotal evidence that healthy subjects can experience a renewed sense of optimism and meaning in life as a result of psychedelic use (Masters and Houston 1966, Shanon 2002). But it would be an interesting project to rigorously and quantitatively examine psychedelics' capacity to combat disenchantment and nihilism in ostensibly mentally healthy people who nonetheless report significant levels of existential distress. This could be combined with the research questions suggested in the preceding paragraphs, to yield a project examining the ability of psychedelics to alleviate existential distress in non-patient populations and the role of metaphysical consolation in their ability (if such they have) to do so. Such research would amount to an interesting kind of experimental philosophy focused on what Langlitz (2016) calls "live philosophical questions".

Further research is also needed on the predictive coding-inspired unbinding theory of ego dissolution developed in the fourth paper. Gerrans and I argue that disruption to different large-scale brain networks in different conditions would result in phenomenologically distinct ego dissolution experiences by disrupting different aspects of the self-model. As we remark in the paper, more fine-grained psychometric instruments capable of discriminating between different varieties or aspects of ego dissolution would help in testing this prediction (although some progress could be made by examining the written or verbal descriptions of subjects in different neuroimaging studies identifying different neural correlates of ego dissolution.) The development of such instruments could be an interesting interdisciplinary project because it ideally ought to draw on descriptions of ego dissolution experiences as well as on recent research, both scientific and philosophical, on the different varieties and aspects of self-awareness.

There are also outstanding questions concerning the potential epistemic benefits of psychedelics which I have not addressed in any of these papers. As I mentioned in the introduction, both Metzinger (2003) and Shanon (2010) endorse the venerable claim that psychedelics can allow subjects to gain therapeutically valuable insights into their own psychology. The account of ego dissolution as unbinding goes some way to establishing a mechanistically plausible explanation of how this could work: by disrupting self-binding processes, psychedelics decouple salience and attention from self-representation, allowing the sort of decentering and objectivity which permits uncomfortable mental contents to be viewed with detachment, honesty and clarity, rather than denial. But much more work needs to be done to develop this proposal, including answering Roche's (2010) arguments against the possibility of any kind of drug-induced propositional knowledge.

Further questions in psychedelic epistemology have more strongly conceptual dimensions. Shanon (2010) lists a variety of other putatively epistemic benefits of psychedelics, including the apprehension or appreciation of truths about nature in new and deeper ways; increases to well-being and stamina; more harmonious social interaction; novel solutions to problems in one's field of specialisation; and enhanced abilities of artistic performance and creativity. Determining whether psychedelics do in fact cause each of these results is an

empirical matter, and determining the extent to which each of them amounts to a genuinely epistemic benefit demands rigorous conceptual analysis. Although he does not explicitly use an epistemic vocabulary, similar points apply to Tupper's (2002, 2003) proposals that psychedelics can provide access to mythic and somatic forms of understanding and can enhance existential intelligence.

Another epistemological concept which may be useful in analysing psychedelic transformation is that of 'knowledge how' (Fantl 2016). I have not explicitly mentioned this possibility so far, but some of Shanon's and Tupper's proposals, plus the connections between psychedelics and meditation, raise the question whether certain aspects of psychedelic-assisted learning might best be conceived of as the acquisition of cognitive, emotional, or attentional skills or abilities.

In sum, many more questions remain than I have even attempted to answer here, and most of them require further empirical, theoretical, and conceptual work. With any luck, the psychedelic renaissance, and the philosophical exploration thereof, is only just beginning.

References

1. Adams, D. (1979). *The hitchhiker's guide to the galaxy*. Pan Books.
2. Albahari, M. (2006). *Analytical Buddhism: the two-tiered illusion of self*. Palgrave Macmillan.
3. Albahari, M. (2014). Insight knowledge of no self in Buddhism: an epistemic analysis. *Philosopher's Imprint* 14, 1-30.
4. Aragona, M. (2013). Philosophy of clinical psychopharmacology. *Psychiatria Danubina* 25, 4-10.
5. Bachmann, T. (2015). How a (sub) cellular coincidence detection mechanism featuring layer-5 pyramidal cells may help produce various visual phenomena. *Frontiers in Psychology* 6. doi: 10.3389/fpsyg.2015.01947
6. Badiner, A. H. and A. Grey (eds.) (2002). *Zig zag zen: Buddhism and psychedelics*. Chronicle Books.
7. Ballard, D. H., G. E. Hinton, and T. J. Sejnowski (1983). Parallel visual computation. *Nature* 306, 21-26.
8. Biddle, J. B. (2016). Inductive risk, epistemic risk, and overdiagnosis of disease. *Perspectives on Science* 24, 192-205.
9. Bogenschutz, M. P., A. A. Forcehimes, J. A. Pommy, C. E. Wilcox, P. C. R. Barbosa, and R. J. Strassman (2015). Psilocybin-assisted treatment for alcohol dependence: a proof-of-concept study. *Journal of Psychopharmacology* 29, 289-299.
10. Bogenschutz, M. P., & Johnson, M. W. (2015). Classic hallucinogens in the treatment of addiction. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*. DOI: <http://dx.doi.org/10.1016/j.pnpbp.2015.03.002>.
11. Bortolotti, L. (2013). Delusions. In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Summer 2015 Edition)*. <http://plato.stanford.edu/archives/sum2015/entries/delusion>
12. Bortolotti, L. (2015a). The epistemic innocence of motivated delusions. *Consciousness and Cognition*, 33, 490-499.
13. Bortolotti, L. (2015b). Epistemic benefits of elaborated and systematized delusions in schizophrenia. *British Journal for the Philosophy of Science*. DOI: <http://dx.doi.org/10.1093/bjps/axv024>.
14. Bortolotti, L. & Gunn, R. (2015). Philip Gerrans: The measure of madness: philosophy, cognitive neuroscience, and delusional thought. *British Journal for the Philosophy of Science*. DOI: <http://dx.doi.org/10.1093/bjps/axv032>.
15. Bortolotti, L. & Miyazono, K. (2016). The ethics of delusional belief. *Erkenntnis* 81, 275-296.
16. Bourguignon, E. (1973). *Religion, altered states of consciousness, and social change*. Ohio State University Press.
17. Bouso, J. C., F. Palhano-Fontes, A. Rodriguez-Fornells, S. Ribeiro, R. Sanches, J. A. S. Crippa, J. E. C. Hallak, D. B. de Araujo, and J. Riba (2015). Long-term use of psychedelic drugs is associated with differences in brain structure and personality in humans. *European Neuropsychopharmacology* 25, 483-492.
18. Brewer, J. A. and K. A. Garrison (2014). The posterior cingulate cortex as a plausible mechanistic target of meditation. *Annals of the New York Academy of Sciences* 1307, 19-27.
19. Brewer, J. A., K. A. Garrison, and S. Whitfield-Gabrieli (2013). What about the "self" is processed in the posterior cingulate cortex? *Frontiers in Human Neuroscience* 7, 1-7.
20. Brogaard, B. and D. E. Gatzia (2016). What can neuroscience tell us about the hard problem of consciousness? *Frontiers in Neuroscience* 10. doi: 10.3389/fnins.2016.00395
21. Brown, D. J. and L. Reitman (2010). Psilocybin studies and the religious experience: an interview with Roland Griffiths, Ph.D. *MAPS Bulletin* 20, 22-25. <http://www.maps.org/news-letters/v20n1/v20n1-22to25.pdf>
22. Burwick, T. (2014). The binding problem. *Wiley Interdisciplinary Reviews: Cognitive Science* 5, 305-315.
23. Byrd, K. R., D. Lear, and S. Schwenka (2000). Mysticism as a predictor of subjective well-being. *The International Journal for the Psychology of Religion* 10, 259- 269.
24. Carhart-Harris, R. L. and K. Friston (2010). The default mode, ego-functions, and free-energy: a neurobiological account of Freudian ideas. *Brain* 133, 1265- 1283.
25. Carhart-Harris, R. L., D. Erritzoe, T. Williams *et al.* (2012). Neural correlates of the psychedelic state as determined by fMRI studies with psilocybin. *Proceedings of the National Academy of Sciences* 109, 2138-2143.
26. Carhart-Harris, R. L., R. Leech, P. J. Hellyer, *et al.* (2014). The entropic brain: a theory of conscious states informed by neuroimaging research with psychedelic drugs. *Frontiers in Human Neuroscience* 8, 1-22.
27. Cavanna, A. E. (2015). On the philosophy of psychopharmacology. *Cognitive Neuropsychiatry* 20, 551-554.
28. Churchland, P. S. (1986). *Neurophilosophy: toward a unified science of the mind-brain*. MIT Press.
29. Churchland, P. S. (2011). *Braintrust: what neuroscience tells us about morality*. Princeton University Press.
30. Clifford, W. K. (1877). The ethics of belief. *Contemporary Review*.
31. Craver, C. F. (2007). *Explaining the brain: mechanisms and the mosaic unity of neuroscience*. Oxford University Press.

32. Dakwar, E., C. Anerella, C. L. Hart, F. R. Levin, S. J. Mathew, and E. V. Nunes (2014). Therapeutic infusions of ketamine: do the psychoactive effects matter? *Drug and Alcohol Dependence* 136, 153-157.
33. Damasio, A. R. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation* 1, 123-132.
34. Damasio, A. R. (2010). *Self comes to mind: constructing the conscious brain*. Pantheon.
35. D'Argembeau, A. (2013). On the role of the ventromedial prefrontal cortex in self-processing: the valuation hypothesis. *Frontiers in Human Neuroscience* 7. doi: 10.3389/fnhum.2013.00372
36. D'Argembeau, A., H. Cassol, C. Phillips, E. Baiteau, E. Salmon, and M. van der Linden (2014). Brains creating stories of selves: the neural basis of autobiographical reasoning. *Social Cognitive and Affective Neuroscience* 9, 646-652.
37. Dass, R. (1971). *Be here now*. Modern Press.
38. Dass, R. (2005). Walking the path: psychedelics and beyond. In R. Walsh and C. S. Grob (eds.) *Higher wisdom: eminent elders explore the continuing impact of psychedelics*, 207-221. SUNY Press.
39. Davey, C. J., J. Pujol, and B. J. Harrison (2016). Mapping the self in the brain's default mode network. *Neuroimage* 132, 390-397.
40. Davis, D. M., and J. A. Hayes (2011). What are the benefits of mindfulness? A practice review of psychotherapy-related research. *Psychotherapy* 48, 198-208.
41. Delmonte, M. M. & Kenny, V. (1987). Conceptual models and functions of meditation in psychotherapy. *Journal of Contemporary Psychotherapy* 17, 38- 59.
42. Dennett, D. C. (1984). *Elbow room: the varieties of free will worth wanting*. MIT Press.
43. Dennett, D. (1991). *Consciousness explained*. Little Brown & Co.
44. Dennett, D. (1992). The self as a center of narrative gravity. In F. Kessel, P. Cole, and D. Johnson (eds). *Self and consciousness: multiple perspectives*. Erlbaum.
45. Dennett, D. (2013). Kinds of things: towards a bestiary of the manifest image. In D. Ross, J. Ladyman, and H. Kincaid (eds.) *Scientific metaphysics*, 96-107. Oxford University Press.
46. Doblin, R. (1991). Pahnke's Good Friday experiment: a long-term follow-up and methodological critique. *Journal of Transpersonal Psychology* 23, 1-28.
47. dos Santos, R. G., F. L. Osório, J. A. S. Crippa *et al.* (2016a). Antidepressive, anxiolytic, and antiaddictive effects of ayahuasca, psilocybin and lysergic acid diethylamide (LSD): a systematic review of clinical trials published in the last 25 Years. *Therapeutic Advances in Psychopharmacology* 6, 193-213.
48. dos Santos, R. G, F. L. Osório, J. A. S. Crippa *et al.* (2016b). Classical hallucinogens and neuroimaging: a systematic review of human studies: hallucinogens and neuroimaging. *Neuroscience & Biobehavioral Reviews* 71, 715-728.
49. Durr, R. A. (1970). *Poetic vision and the psychedelic experience*. Syracuse University Press.
50. Dyck, E. (2008). *Psychedelic psychiatry: LSD from clinic to campus*. Johns Hopkins University Press.
51. Eisner, B. G. and S. Cohen (1958). Psychotherapy with lysergic acid diethylamide. *Journal of Nervous and Mental Disease* 127, 528-539.
52. Epstein, M. (2004). *Thoughts without a thinker: psychotherapy from a Buddhist perspective*. Basic Books.
53. Fantl, J. (2016). Knowledge how. In E. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2016 edition). <https://plato.stanford.edu/archives/spr2016/entries/knowledge-how/>
54. Fischman, L. G. (1983). Dreams, hallucinogenic drug states, and schizophrenia: a psychological and biological comparison. *Schizophrenia Bulletin* 9, 73-94.
55. Flanagan, O. (2007). *The really hard problem: meaning in a material world*. MIT Press.
56. Flanagan, O. and G. D. Caruso (forthcoming). Neuroexistentialism: Third-wave Existentialism. In G. D. Caruso and O. Flanagan (eds.) *Neuroexistentialism: meaning, morals and purpose in the age of neuroscience*. Oxford University Press.
57. Garcia-Romeu, A., R. R. Griffiths, and M. W. Johnson (2014). Psilocybin-occasioned mystical experiences in the treatment of tobacco addiction. *Current Drug Abuse Reviews* 7, 157-164.
58. Garcia-Romeu, A., B. Kersgaard, and P. H. Addy (2016). Clinical applications of hallucinogens: a review. *Experimental and Clinical Psychopharmacology* 24, 229-268.
59. Gasser, P., D. Holstein, Y. Michel *et al.* (2014). Safety and efficacy of lysergic acid diethylamide-assisted psychotherapy for anxiety associated with life-threatening diseases. *Journal of Nervous and Mental Disease* 202, 513-520.
60. Gasser, P., K. Kirchner, and T. Passie. (2015). LSD-assisted psychotherapy for anxiety associated with a life-threatening disease: a qualitative study of acute and sustained subjective effects. *Journal of Psychopharmacology* 29, 57-68.
61. Gerrans, P. (2014). *The measure of madness: philosophy, cognitive neuroscience, and delusional thought*. MIT Press.
62. Gilboa, A. (2004). Autobiographical and episodic memory - one and the same? Evidence from prefrontal activation in neuroimaging studies. *Neuropsychologia* 42, 1336-1349.
63. Glannon, W. (2008). Psychopharmacological enhancement. *Neuroethics* 1, 45-54.
64. Goodenough, U. (1998). *The sacred depths of nature*. Oxford University Press.
65. Goodenough, U. (2001). Vertical and horizontal transcendence. *Zygon* 36, 21-31.
66. Goodenough, U. and P. Woodruff (2001). Mindful virtue, mindful reverence. *Zygon* 36, 585-595.
67. Goyal, M., S. Singh, E. M. S. Sibinga *et al.* (2014). Meditation programs for psychological stress and well-being: a systematic review and meta-analysis. *Journal of the American Medical Association Internal Medicine* 174, 357-368.
68. Graham, G. and G. L. Stephens (1994). *Philosophical psychopathology*. MIT Press.
69. Greyson, B. (1997). The near-death experience as a focus of clinical attention. *Journal of Nervous and Mental Disease* 185, 327-334.
70. Griffiths R. R., W. A. Richards, U. McCann, and R. Jesse (2006). Psilocybin can occasion mystical-type experiences having substantial and sustained personal meaning and spiritual significance. *Psychopharmacology* 187, 268-283.
71. Griffiths, R. R., M. W. Johnson, M. A. Carducci *et al.* (2016). Psilocybin produces substantial and sustained decreases in depression and anxiety in patients with life-threatening cancer: a randomized double-blind trial. *Journal of Psychopharmacology* 30, 1181-1197.

72. Grinspoon, L. & Bakalar, J. B. (1979). *Psychedelic drugs reconsidered*. Basic Books.
73. Grinspoon, L. & Bakalar, J. B. (1986). Can drugs be used to enhance the psychotherapeutic process? *American Journal of Psychotherapy* 40, 393-404.
74. Grob, C. S. (2007). The use of psilocybin in patients with advanced cancer and existential anxiety. In M. Winkelman and T. B. Roberts (eds.) *Psychedelic medicine: new evidence for hallucinogenic substances as treatments*, 205-216. Praeger.
75. Grob, C. S., D. J. McKenna, J. C. Callaway, G. S. Brito, E. S. Neves, G. Oberlaender, O. L. Saide, E. Labigalini, C. Tacla, C. T. Miranda, R. J. Strassman, and K. B. Boone (1996). Human psychopharmacology of Hoasca, a plant hallucinogen used in ritual context in Brazil. *Journal of Nervous and Mental Disease* 184, 86-94.
76. Grob, C. S., G. R. Greer, and M. Mangini (1998). Editors' introduction: hallucinogens at the turn of the century. *Journal of Psychoactive Drugs* 30, 315-319.
77. Grob, C. S., A. L. Danforth, G. S. Chopra *et al.* (2011). Pilot study of psilocybin treatment for anxiety in patients with advanced-stage cancer. *Archives of General Psychiatry* 68, 71-78.
78. Grof, S. (1975). *Realms of the human unconscious: observations from LSD research*. Viking.
79. Guterstam, A., M. Björnsdotter, G. Gentile *et al.* (2015). Posterior cingulate cortex integrates the senses of self-location and body ownership. *Current Biology* 25, 1416-1425.
80. Halberstadt, A. (2015). Recent advances in the neuropsychopharmacology of serotonergic hallucinogens. *Behavioral Brain Research* 277, 91-120.
81. Harris, S. (2014). *Waking up: a guide to spirituality without religion*. Simon and Schuster.
82. Hermle, L., Fünfgeld, M., Oepen, G., Botsch, H., Borchardt, D., Gouzoulis, E. ... & Spitzer, M. (1992). Mescaline induced psychopathological, neuropsychological and neurometabolic effects in normal subjects: experimental psychosis as a tool for psychiatric research. *Biological Psychiatry* 32, 976-991.
83. Hofmann, A. (1980). *LSD: my problem child*. McGraw-Hill.
84. Hoffmann, E., J. M. K. Hesselink, and Y.-W. M. da Silveira Barbosa (2001). Effects of a psychedelic, tropical tea, Ayahuasca, on the electroencephalographic (EEG) activity of the human brain during a shamanistic ritual. *MAPS Bulletin* 11, 25-30
85. Hohwy, J. (2013). *The predictive mind*. Oxford University Press.
86. Hohwy, J. and J. Michael (forthcoming). Why should any body have a self? In F. de Vignemont and A. Alsmith (eds.) *The body and the self, revisited*. MIT Press.
87. Hood, R. W. (2006). The common-core thesis in the study of mysticism. In P. McNamara (ed.) *Where God and science meet: how brain and evolutionary studies alter our understanding of religion*, vol. 3, 119-138. Praeger.
88. Horst, S. (2009). Naturalisms in philosophy of mind. *Philosophy Compass* 4, 219- 254.
89. Huxley, A. (1932). *Brave new world*. Vintage.
90. Huxley, A. (1945). *The perennial philosophy*. Harper and Brothers.
91. Huxley, A. (1954). *The doors of perception*. Chatto and Windus.
92. Huxley, A. (1962). *Island*. Harper Collins.
93. Huxley, A. (1994). *The doors of perception and heaven and hell*. Flamingo
94. Jacobson, R. (2014, August 14). Turn on, tune in, get better: psychedelic drugs hold medical promise. *Scientific American*. <http://www.scientificamerican.com/article/turn-on-tune-in-get-better-psychedelic-drugs-hold-medical-promise/>
95. Johnson, M. W., Richards, W. A., & Griffiths, R. R. (2008). Human hallucinogen research: guidelines for safety. *Journal of Psychopharmacology* 22, 603-620.
96. Johnson, M. W., Garcia-Romeu, A., Cosimano, M. P., & Griffiths, R. R. (2014). Pilot study of the 5-HT_{2A}R agonist psilocybin in the treatment of tobacco addiction. *Journal of Psychopharmacology* 28, 983-992.
97. Kaelen, M., F. S. Barrett, L. Roseman *et al.* (2015). LSD enhances the emotional response to music. *Psychopharmacology* 232, 3607-3614.
98. Kaplan, D. M. and W. Bechtel (2011). Dynamical models: an alternative or complement to mechanistic explanation? *Topics in Cognitive Science* 3, 438- 444.
99. Katz, M. M., I. E. Waskow, and J. Olsson (1968). Characterizing the psychological state produced by LSD. *Journal of Abnormal Psychology* 73, 1-14.
100. Killingsworth, M. A. and D. T. Gilbert (2010). A wandering mind is an unhappy mind. *Science* 330, 932-932.
101. Kingsland, J. (2016). *Siddhartha's brain: the science of meditation, mindfulness and enlightenment*. Robinson.
102. Klee, G. D. (1963). Lysergic acid diethylamide (LSD-25) and ego functions. *Archives of General Psychiatry* 8, 461-474
103. Kornfield, J. (1979). Intensive insight meditation: a phenomenological study. *Journal of Transpersonal Psychology* 11, 41-58.
104. Krebs, T. S., & Johansen, P.-Ø. (2012). Lysergic acid diethylamide (LSD) for alcoholism: meta-analysis of randomized controlled trials. *Journal of Psychopharmacology* 26, 994-1002.
105. Langlitz, N. (2012). *Neuropsychedelica: the revival of hallucinogen research since the decade of the brain*. University of California Press.
106. Langlitz, N. (2016). Is there a place for psychedelics in philosophy? Fieldwork in neuro- and perennial philosophy. *Common Knowledge* 22, 373-384.
107. Leary, T., R. Metzner, and R. Alpert (1964). *The psychedelic experience*. University Books.
108. Lebedev, A. V., M. Lövdén, G. Rosenthal *et al.* (2015). Finding the self by losing the self: neural correlates of ego-dissolution under psilocybin. *Human Brain Mapping* 36, 3137-3153.
109. Lebedev, A. V., M. Kaelen, M. Lövdén *et al.* (2016). LSD-induced entropic brain activity predicts subsequent personality change. *Human Brain Mapping* 37, 3203-3213.
110. Lebrecht, S., M. Bar, L. F. Barrett *et al.* (2012). Micro-valences: perceiving affective valence in everyday objects. *Frontiers in Psychology* 3. doi: 10.3389/fpsyg.2012.00107
111. Leech, R. and D. J. Sharp (2014). The role of the posterior cingulate cortex in cognition and disease. *Brain* 137, 12-32.
112. Letheby, C. (2015). The philosophy of psychedelic transformation. *Journal of Consciousness Studies*, 22, 170-193.
113. Lou, H. C., J. P. Changeux, and A. Rosenstand (in press). Towards a cognitive neuroscience of self-awareness. *Neuroscience and Biobehavioral Reviews*. doi: 10.1016/j.neubiorev.2016.04.004

114. Luke, D. P. and D. B. Terhune (2013). The induction of synaesthesia with chemical agents: a systematic review. *Frontiers in Psychology* 4. doi: 10.3389/fpsyg.2013.00753
115. Luo, Y., K. Feng, S. Qi *et al.* (2015). Resting-state functional connectivity of the default mode network associated with happiness. *Social Cognitive and Affective Neuroscience*. doi: 10.1093/scan/nsv132
116. Machado, L., and A. Cantilino (2016). A systematic review of the neural correlates of positive emotions. *Revista Brasileira de Psiquiatria*. doi: 10.1590/1516-4446-2016-1988
117. Mackie, J. L. (1977). *Ethics: inventing right and wrong*. Penguin.
118. MacLean, K. A., M. W. Johnson, and R. R. Griffiths (2011). Mystical experiences occasioned by the hallucinogen psilocybin lead to increases in the personality domain of openness. *Journal of Psychopharmacology* 25, 1453-1461.
119. Majiæ, T., T. T. Schmidt, and J. Gallinat, (2015). Peak experiences and the afterglow phenomenon: when and how do therapeutic effects of hallucinogens depend on psychedelic experiences? *Journal of Psychopharmacology* 29, 241-253.
120. Malitz, S., H. Esecover, B. Wilkens *et al.* (1960). Some observations on psilocybin, a new hallucinogen, in volunteer subjects. *Comprehensive Psychiatry* 1, 8-17.
121. Mangini, M. (1998). Treatment of alcoholism using psychedelic drugs: a review of the program of research. *Journal of Psychoactive Drugs* 30, 381-418.
122. Mann, K., Günther, A., Stetter, F., & Ackermann, K. (1999). Rapid recovery from cognitive deficits in abstinent alcoholics: a controlled test-retest study. *Alcohol and Alcoholism* 34, 567-574.
123. Mashour, G. A. (2013). Cognitive unbinding: a neuroscientific paradigm of general anesthesia and related states of unconsciousness. *Neuroscience and Biobehavioral Reviews* 37, 2751-2759.
124. Masters, R. E. L. and J. Houston (1966). *The varieties of psychedelic experience*. Holt, Rinehart and Winston.
125. McKenna, T. (1991). *The archaic revival*. Harper Collins.
126. Metzinger, T. (2003). *Being no-one: the self-model theory of subjectivity*. MIT Press.
127. Metzinger, T. (2005). Précis: being no one. *Psyche* 11, 1-35.
128. Metzinger, T. (2009). *The ego tunnel: the science of the mind and the myth of the self*. Basic Books.
129. Metzinger, T. (2011). The no-self alternative. In S. Gallagher (ed.), *The Oxford handbook of the self*, 279-296. Oxford University Press.
130. Metzinger, T. (2013). Why are dreams interesting for philosophers? The example of minimal phenomenal selfhood, plus an agenda for future research. *Frontiers in Psychology* 4. doi: 10.3389/fpsyg.2013.00746
131. Metzinger, T. (2014). Spirituality and intellectual honesty: an essay. http://www.blogs.unimainz.de/fb05philosophieengl/files/2013/07/TheorPhil_Metzinger_SIR_2014_English.pdf
132. Metzinger, T. (2016). All about the ego tunnel: interview by Richard Marshall. *3AM Magazine*. <http://www.3ammagazine.com/3am/all-about-the-ego-tunnel/>
133. Miller, C. (2014). Moral virtues, epistemic virtues, and the Big Five. In A. Fairweather & O. Flanagan (eds.) *Naturalizing epistemic virtue*, 92-117. Cambridge University Press.
134. Moreno, F. A., C. B. Wiegand, E. K. Taitano, and P. L. Delgado (2006). Safety, tolerability and efficacy of psilocybin in 9 patients with obsessive-compulsive disorder. *Journal of Clinical Psychiatry* 67, 1735-1740.
135. Nichols, D. E. (2004). Hallucinogens. *Pharmacology and Therapeutics* 101, 131-181.
136. Nichols, D. E. (2016). Psychedelics. *Pharmacological Reviews* 68, 264-355.
137. Nichols, D. E., M. W. Johnson, and C. D. Nichols (forthcoming). Psychedelics as medicines: an emerging new paradigm. *Clinical Pharmacology & Therapeutics*. doi: 10.1002/cpt.557
138. Nour, M. M., L. Evans, D. Nutt *et al.* (2016). Ego-dissolution and psychedelics: validation of the ego-dissolution inventory (EDI). *Frontiers in Human Neuroscience* 10. doi: 10.3389/fnhum.2016.00269.
139. Nozick, R. (1974). *Anarchy, state and utopia*. Basic Books.
140. Nutt, D. (2016). Psilocybin for anxiety and depression in cancer care? Lessons from the past and prospects for the future. *Journal of Psychopharmacology* 30, 1163-1164.
141. Osmond, H. (1957). A review of the clinical effects of psychotomimetic agents. *Annals of the New York Academy of Sciences* 66, 418-434.
142. Osório, F. L., Sanches, R. F., Macedo, L. R., dos Santos, R. G., Maia-de-Oliveira, J. P., Wichert-Ana, L. ... & Hallak, J. E. (2015). Antidepressant effects of ayahuasca in patients with recurrent depression: a preliminary report. *Revista Brasileira de Psiquiatria* 37, 13-20.
143. Osto, D. (2016). *Altered states: Buddhism and psychedelic spirituality in America*. Columbia University Press.
144. Pahnke, W. N. (1963). Drugs and mysticism: an analysis of the relationship between psychedelic drugs and the mystical consciousness: a thesis. Doctoral dissertation, Harvard University.
145. Pahnke, W. (1969). The psychedelic mystical experience in the human encounter with death. *The Harvard Theological Review* 62, 1-21.
146. Pahnke, W. N. and W. A. Richards (1966). Implications of LSD and experimental mysticism. *Journal of Religion and Health* 5, 175-208.
147. Palhano-Fontes, F., K. C. Andrade, L. F. Tofoli, A. C. Santos, J. A. S. Crippa, J. E. C. Hallak, S. Ribeiro, and D. B. de Araujo (2015). The psychedelic state induced by ayahuasca modulates the activity and connectivity of the default mode network. *PLoS ONE* 10, e0118143.
148. Papineau, D. (2009). Naturalism. In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Spring 2009 Edition)*. <http://plato.stanford.edu/archives/spr2009/entries/naturalism/>
149. Parens, E. (2005). Authenticity and ambivalence: toward understanding the enhancement debate. *Hastings Center Report* 35, 34-41.
150. Pollan, M. (2015, February 9). The trip treatment. *The New Yorker*. <http://www.newyorker.com/magazine/2015/02/09/trip-treatment>. 28 Aug 2015.
151. Prinz, J. J. (2004). *Gut reactions: a perceptual theory of emotion*. Oxford University Press.
152. Prinz, J. (2014). Seeing with feeling. In Gregory Currie, Matthew Kieran, Aaron Meskin *et al.* (eds.) *Aesthetics and the*

- sciences of mind, 143-158. Oxford University Press.
153. Qin, P. and G. Northoff (2011). How is our self related to midline regions and the default mode network? *Neuroimage* 57, 1221-1233.
 154. Raichle, M. E., A. M. MacLeod, A. Z. Snyder *et al.* (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences* 98, 676-682.
 155. Ren, J., Li, H., Palaniyappan, L., Liu, H., Wang, J., Li, C., & Rossini, P. M. (2014). Repetitive transcranial magnetic stimulation versus electroconvulsive therapy for major depression: a systematic review and meta-analysis. *Progress in Neuro-Psychopharmacology and Biological Psychiatry* 51, 181-189.
 156. Revonsuo, A. (1999). Binding and the phenomenal unity of consciousness. *Consciousness and Cognition* 8, 173-185.
 157. Roberts, T. B. (2013). *The psychedelic future of the mind: how entheogens are enhancing cognition, boosting intelligence, and raising values*. Park Street.
 158. Roche, G. T. (2010). Seeing snakes: on delusion, knowledge and the drug experience. In D. Jacquetted (ed.) *Cannabis - philosophy for everyone: what were we just talking about?*. John Wiley and Sons.
 159. Ross, D. (2014). *Philosophy of economics*. Palgrave Macmillan.
 160. Ross, S., A. Bossis, J. Guss *et al.* (2016). Rapid and sustained symptom reduction following psilocybin treatment for anxiety and depression in patients with lifethreatening cancer: a randomized controlled trial. *Journal of Psychopharmacology* 30, 1165-1180.
 161. Ruck, C. A. P., J. Bigwood, D. Staples, J. Ott and G. Wasson (1979). Entheogens. *Journal of Psychedelic Drugs* 11, 145-146.
 162. Russell, B. (1910). Knowledge by acquaintance and knowledge by description. *Proceedings of the Aristotelian Society* 11, 108-128.
 163. Savage, C. (1955). Variations in ego feeling induced by d-lysergic acid diethylamide (LSD-25). *The Psychoanalytic Review* 42, 1-16.
 164. Schermer, M. H. N. (2007). Brave new world versus island - utopian and dystopian views on psychopharmacology. *Medicine, Health Care and Philosophy* 10, 119-128.
 165. Schmid, Y., F. Enzler, P. Gasser *et al.* (2015). Acute effects of lysergic acid diethylamide in healthy subjects. *Biological Psychiatry* 78, 544-553.
 166. Sellars, W. (1963). *Science, perception and reality*. The Humanities Press.
 167. Sessa, B. (2005). Can psychedelics have a role in psychiatry once again? *British Journal of Psychiatry* 186, 457-458.
 168. Sessa, B. (2008). Is it time to revisit the role of psychedelic drugs in enhancing human creativity? *Journal of Psychopharmacology* 22, 821-827.
 169. Sessa, B. (2012). *The psychedelic renaissance: reassessing the role of psychedelic drugs in 21st century psychiatry and society*. Muswell Hill Press.
 170. Sessa, B. & Johnson, M. W. (2015). Can psychedelic compounds play a part in drug dependence therapy? *British Journal of Psychiatry* 206, 1-3.
 171. Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences* 17, 565-573.
 172. Shanon, B. (2002). *The antipodes of the mind: charting the phenomenology of the ayahuasca experience*. Oxford University Press.
 173. Shanon, B. (2010). The epistemics of ayahuasca visions. *Phenomenology and the Cognitive Sciences* 9, 263-280.
 174. Simpson, W. (2014). The mystical stance: the experience of self-loss and Daniel Dennett's "center of narrative gravity". *Zygon* 49, 458-475.
 175. Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron* 24, 49-65.
 176. Smith, H. (1964). Do drugs have religious import? *Journal of Philosophy* 61, 517- 530.
 177. Smith, H. (2000). *Cleansing the doors of perception: the religious significance of entheogenic plants and chemicals*. Tarcher/Putnam.
 178. Sober, E. (2000). *Philosophy of biology*, 2nd ed. Westview Press.
 179. Soler, J., M. Elices, A. Franquesa *et al.* (2016). Exploring the therapeutic potential of ayahuasca: acute intake increases mindfulness-related capacities. *Psychopharmacology* 233, 823-829.
 180. Solomon, R. C. (2002). *Spirituality for the skeptic: the thoughtful love of life*. Oxford University Press.
 181. Speth, J., C. Speth, M. Kaelen *et al.* (2016). Decreased mental time travel to the past correlated with default-mode network disintegration under lysergic acid diethylamide. *Journal of Psychopharmacology* 30, 344-353.
 182. Spreng, R. N. and C. L. Grady (2010). Patterns of brain activity supporting autobiographical memory, prospection, and theory of mind, and their relationship to the default mode network. *Journal of Cognitive Neuroscience* 22, 1112-1123.
 183. Stace, W. T. (1960). *Mysticism and philosophy*. Lippincott.
 184. Stein, D. J. (2012). *Philosophy of psychopharmacology*. Cambridge University Press.
 185. Stein, D. J. (2012). Psychopharmacological enhancement: a conceptual framework. *Philosophy, Ethics and Humanities in Medicine* 7, 5.
 186. Stephens, G. L. and G. Graham (2007). Philosophical psychopathology and self-consciousness. In M. Velmans and S. Schneider (eds). *The Blackwell companion to consciousness*, 194-208. Blackwell.
 187. Stone, J. A. (2012). Spirituality for naturalists. *Zygon* 47, 481-500.
 188. Strassman, R. J. (1984). Adverse reactions to psychedelic drugs: a review of the literature. *Journal of Nervous and Mental Disease* 172, 577-595.
 189. Strassman, R. J. (1995a). Hallucinogenic drugs in psychiatric research and treatment. *Journal of Nervous and Mental Disease* 183, 127-138.
 190. Strassman, R. J. (1995b). Human psychopharmacology of N,N-dimethyltryptamine. *Behavioral Brain Research* 73, 121-124.
 191. Strassman, R. J. (2001). *DMT: the spirit molecule*. Park Street Press.
 192. Strassman, R. (2005). Hallucinogens. In M. Earlywine (ed.) *Mind-altering drugs: the science of subjective experience*, 49-85. Oxford University Press.
 193. Stuckey, D. E., R. Lawson, and L. E. Luna (2005). EEG gamma coherence and other correlates of subjective reports during

- Ayahuasca experiences. *Journal of Psychoactive Drugs* 37, 163-178.
194. Studerus, E., M. Komter, F. Hasler *et al.* (2011). Acute, subacute and long-term subjective effects of psilocybin in healthy humans: a pooled analysis of experimental studies. *Journal of Psychopharmacology* 25, 1434-1452.
195. Sturgeon, N. (2006). Ethical naturalism. In D. Copp (ed.) *The Oxford handbook of ethical theory*, 91-121. Oxford University Press.
196. Sui, J. and G. W. Humphreys (2015). The integrative self: how self-reference integrates perception and memory. *Trends in Cognitive Sciences* 19, 719-728.
197. Sullivan-Bissett, E. (2015). Implicit bias, confabulation and epistemic innocence. *Consciousness and Cognition* 33, 548-560.
198. Szara, S. (1967). The hallucinogenic drugs - curse or blessing? *American Journal of Psychiatry* 123, 1513-1518.
199. The President's Council on Bioethics (2003). *Beyond therapy: biotechnology and the pursuit of happiness*. <https://bioethicsarchive.georgetown.edu/pcbe/reports/beyondtherapy/index.html>. 10 Jun 2014.
200. Tagliazucchi, E., Carhart-Harris, R., Leech, R., Nutt, D., & Chialvo, D. R. (2014). Enhanced repertoire of brain dynamical states during the psychedelic experience. *Human Brain Mapping* 35, 5442-5456.
201. Tagliazucchi, E., L. Roseman, M. Kaelen *et al.* (2016). Increased global functional connectivity correlates with LSD-induced ego dissolution. *Current Biology* 26, 1043-1050.
202. Thomas, G., Lucas, P., Capler, N. R., Tupper, K. W., & Martin, G. (2013). Ayahuasca-assisted therapy for addiction: results from a preliminary observational study in Canada. *Current Drug Abuse Reviews* 6, 30-42.
203. Tupper, K. W. (2002). Entheogens and existential intelligence: the use of plant teachers as cognitive tools. *Canadian Journal of Education/Revue canadienne de l'education* 27, 499-516.
204. Tupper, K. (2003). Entheogens & education: exploring the potential of psychoactives as educational tools. *Journal of Drug Education and Awareness* 1, 145-161.
205. Turton, S., D. J. Nutt, and R. L. Carhart-Harris (2014). A qualitative report on the subjective experience of intravenous psilocybin administered in an fMRI environment. *Current Drug Abuse Reviews* 7, 117-127.
206. Vaughan, F. (1983). Perception and knowledge: reflections on psychological and spiritual learning in the psychedelic experience. In L. Grinspoon and J. B. Bakalar (eds.) *Psychedelic reflections*. Human Sciences Press.
207. van Essen, D. C., C. H. Anderson, and D. J. Felleman (1992). Information processing in the primate visual system: An integrated systems perspective. *Science* 255, 419-423.
208. Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* 21, 615-628.
209. Vollenweider, F. X., K. L. Leenders, C. Scharfetter *et al.* (1997). Positron emission tomography and fluorodeoxyglucose studies of metabolic hyperfrontality and psychopathology in the psilocybin model of psychosis. *Neuropsychopharmacology* 16, 357-372.
210. Vollenweider, F. X. and Komter, M. (2010). The neurobiology of psychedelic drugs: implications for the treatment of mood disorders. *Nature Reviews Neuroscience* 11, 642-651.
211. Walsh, R. (2003). Entheogens: true or false. *International Journal of Transpersonal Studies* 22, 1-6.
212. Wasson, R. G. (1957). Seeking the magic mushroom. *Life* 42, 100.
213. Watts, A. (1962). *The joyous cosmology: adventures in the chemistry of consciousness*. Pantheon.
214. Watts, A. (1964). A psychedelic experience: fact or fantasy? In D. Solomon (ed.) *LSD: the consciousness-expanding drug*. Putnam.
215. White, P. A. (2015). The pre-reflective experience of "I" as a continuously existing being: the role of temporal functional binding. *Consciousness and Cognition* 31, 98-114.
216. Windt, J. M. (2011). Altered consciousness in philosophy. In E. Cardeña and M. Winkelmann (eds.) *Altering consciousness: multidisciplinary perspectives*, 229-254. ABC-CLIO.
217. Yaden, D. B., K. D. Le Nguyen, M. L. Kern *et al.* (2016). Of roots and fruits: a comparison of psychedelic and nonpsychedelic mystical experiences. *Journal of Humanistic Psychology*. doi: 10.1177/0022167816674625
218. Young, S. N. (2013). Single treatments which have lasting effects: some thoughts on the antidepressant effects of ketamine and botulinum toxin and the anxiolytic effect of psilocybin. *Journal of Psychiatry and Neuroscience* 38, 78-83.
219. Zaehner, R. C. (1958). *Mysticism: sacred and profane*. Clarendon Press.