The Philosophy of Psychedelic Transformation

Chris Letheby

(Forthcoming in Journal of Consciousness Studies. Please quote from published version.)

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 somewhat[1] 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 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

[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.

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 premature[2]. 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.

[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 an naturalistic explanatory bridge between those correlates and the qualitative character of conscious experience.

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 ingredient[3] 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

[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.

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, p. 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 benefits[4]. 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

[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.

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 did[5]. 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,

[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.

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, pp. 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-HT2A 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 same[6]. 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 change[7]. 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

[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.

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 benefits[8].

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,

[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.

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 acquaintance[9] 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 'neuropsychoanalysis' 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

[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.

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 kind[10] 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

[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.

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 knowledge[11] 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

[11] 'Modal knowledge' is a philosophical term for knowledge about possibility, necessity, and so forth.

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 naturalism[12].

[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.

References

Badiner, A. H. and A. Grey (eds.) (2002). Zig zag zen: Buddhism and psychedelics. San Francisco: Chronicle Books.

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.

Bourguignon, E. (1973). Religion, altered states of consciousness, and social change. Columbus: Ohio State University Press.

Bouso, J. C., F. Palhano-Fontes, A. Rodríguez-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.

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(647), 1-7.

Brown, D. J. and L. Reitman (2010). Psilocybin studies and the religious experience: an interview with Roland Griffiths, Ph.D. MAPS Bulletin 20(1), 22-25. <http://www.maps.org/news-letters/v20n1/v20n1-22to25.pdf>. 17 October 2014.

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.

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.

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.

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(20), 1-22.

Craver, C. F. (2007). Explaining the brain: mechanisms and the mosaic unity of neuroscience. New York: Oxford University Press.

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.

D'Argembeau, A., H. Cassol, C. Phillips, E. Balteau, 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(5), 646-652.

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. New York: SUNY Press.

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.

Doblin, R. (1991). Pahnke's Good Friday experiment: a long-term follow-up and methodological critique. Journal of Transpersonal Psychology 23, 1-28.

Epstein, M. (2004). Thoughts without a thinker: psychotherapy from a Buddhist perspective. New York: Basic Books

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.

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(7), 513-520.

Glannon, W. (2008). Psychopharmacological enhancement. Neuroethics 1, 45-54.

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.

Greyson, B. (1997). The near-death experience as a focus of clinical attention. Journal of Nervous and Mental Disease 185, 327-334.

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.

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.

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.

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.

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

Horst, S. (2009). Naturalisms in philosophy of mind. Philosophy Compass 4, 219-254.

Huxley, A. (1932). Brave new world. London: Vintage.

Huxley, A. (1962). Island. New York: Harper Collins.

Huxley, A. (1994). The doors of perception and heaven and hell. London: Flamingo.

Kaplan, D. M. and W. Bechtel (2011). Dynamical models: an alternative or complement to mechanistic explanation? Topics in Cognitive Science 3(2), 438-444.

Kornfield, J. (1979). Intensive insight meditation: a phenomenological study. Journal of Transpersonal Psychology 11, 41-58.

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.

Masters, R. E. L. and J. Houston (1966). The varieties of psychedelic experience. New York: Holt, Rinehart and Winston.

Metzinger, T. (2003). Being no-one: the self-model theory of subjectivity. Cambridge, Mass.: MIT Press.

Metzinger, T. (2009). The ego tunnel: the science of the mind and the myth of the self. New York: Basic Books.

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.

Osmond, H. (1957). A review of the clinical effects of psychotomimetic agents. Annals of the New York Academy of Sciences 66, 418-434.

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.

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(2), e0118143.

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/>.

Parens, E. (2005). Authenticity and ambivalence: toward understanding the enhancement debate. Hastings Center Report 35, 34-41.

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(2), 676-682.

Ruck, C. A. P., J. Bigwood, D. Staples, J. Ott and G. Wasson (1979). Entheogens. Journal of Psychedelic Drugs 11, 145-146.

Russell, B. (1910). Knowledge by acquaintance and knowledge by description. Proceedings of the Aristotelian Society 11, 108-128.

Schermer, M. H. N. (2007). Brave new world versus island – utopian and dystopian views on psychopharmacology. Medicine, Health Care and Philosophy 10, 119-128.

Sessa, B. (2012). The psychedelic renaissance: reassessing the role of psychedelic drugs in 21st century psychiatry and society. London: Muswell Hill Press.

Shanon, B. (2002). The antipodes of the mind: charting the phenomenology of the ayahuasca experience. Oxford: Oxford University Press.

Smith, H. (1964). Do drugs have religious import? Journal of Philosophy 61, 517-530.

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(6), 1112-1123.

Stein, D. J. (2012). Psychopharmacological enhancement: a conceptual framework.

Philosophy, Ethics and Humanities in Medicine 7, 5.

Strassman, R. J. (1995a). Hallucinogenic drugs in psychiatric research and treatment. Journal of Nervous and Mental Disease 183, 127-138.

Strassman, R. J. (1995b). Human psychopharmacology of N,N-dimethyltryptamine. Behavioral Brain Research 73, 121-124.

Strassman, R. J. (2001). DMT: the spirit molecule. Rochester: Park Street Press.

Strassman, R. (2005). Hallucinogens. In M. Earlywine (ed.) Mind-altering drugs: the science of subjective experience, 49-85. New York: Oxford University Press.

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.

Szara, S. (1967). The hallucinogenic drugs—curse or blessing? American Journal of Psychiatry 123, 1513-1518.

The President's Council on Bioethics (2003). Beyond therapy: biotechnology and the pursuit of happiness. <https://bioethicsarchive.georgetown.edu/pcbe/reports/beyondtherapy/index.html> . June 10 2014.

Tupper, K. (2003). Entheogens & education: exploring the potential of psychoactives as educational tools. Journal of Drug Education and Awareness 1, 145-161.

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. New York: Human Sciences Press.

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. Behavioral and Brain Sciences 21(05), 615-628.

Vollenweider, F. X. and Kometer, M. (2010). The neurobiology of psychedelic drugs: implications for the treatment of mood disorders. Nature Reviews Neuroscience 11, 642-651.