Review of Philosophy and Psychology
https://doi.org/10.1007/s13164-023-00709-6

From Altered States to Metaphysics: The Epistemic Status of Psychedelic-induced Metaphysical Beliefs

Paweł Gładziejewski[1]

Accepted: 26 September 2023

Abstract
Psychedelic substances elicit powerful, uncanny conscious experiences that are thought to possess therapeutic value. In those who undergo them, these altered states of consciousness often induce shifts in metaphysical beliefs about the fundamental structure of reality. The contents of those beliefs range from contentious to bizarre, especially when considered from the point of view of naturalism. Can chemically induced, radically altered states of consciousness provide reasons for or play some positive epistemic role with respect to metaphysical beliefs? In this paper, I discuss a view that has been underexplored in recent literature. I argue that psychedelic states can be rationally integrated into one's epistemic life. Consequently, updating one's metaphysical beliefs based on altered states of consciousness does not have to constitute an instance of epistemic irrationality.

Keywords DMT · LSD · Metaphysical beliefs · Mystical experience · Predictive processing · Psychedelics · Psilocybin · REBUS

1 Introduction

Growing evidence lends support to the idea that serotonergic psychedelic drugs—LSD, psilocybin, mescaline, and DMT—can contribute to alleviating different forms of mental suffering, like depression, obsessive-compulsive disorder, addiction, or end-of-life anxiety (for a philosophically-oriented discussion of this vast literature, see Letheby 2021). It is commonly assumed that at least one crucial factor underlying the therapeutic efficacy of psychedelics is the experience that these compounds produce. From this perspective, psychedelics work by eliciting subjectively mean-

✉ Paweł Gładziejewski
pawel.gladz@gmail.com

[1] Department of Cognitive Science, Nicolaus Copernicus University, Ul. Gagarina 39, Toruń 87-100, Poland

Published online: 10 October 2023

ingful and emotionally powerful conscious states that enable people to re-frame their outlook on themselves and the world in psychologically beneficial ways.

But sometimes, and especially at higher doses, psychedelics elicit experiences that altogether break the frame of ordinary consciousness. The phenomenology of those experiences resembles the states described by mystics across different cultures and historical periods (for accounts of mystical experience, see Stace 1960; James 2008/1902; Jones and Gellman 2022; for the connections between mystical and psychedelic experience, see Griffiths et al. 2006; Richards 2018). So, in a deep psychedelic experience, the very subject-object structure is sometimes lost, and one's sense of being a bounded, separate self becomes dissolved into what feels like an all-encompassing unity. These experiences also involve the sense of transcending time and space. For example, one's ordinary subjective temporal flow may be replaced by a sense of apprehending the world from an eternal point of view. Other relevant forms of psychedelic phenomenology (which however do not meet the traditional criteria of mystical experience) involve experiences of being transported to another "dimension" populated by autonomous, conscious entities capable of non-verbal telepathic communication (Davis et al. 2020).

Uncanny as they are, these conscious states possess what William James (2008/1902) called a "noetic" quality. They are associated with the sense of veridicality or reality, often reported as surpassing the sense of reality that characterizes normal sensory perception. Those who undergo acute psychedelic states sometimes describe them as awakenings to a "true" reality hidden underneath the veil of the commonsense world.

Both anecdotal evidence and recent studies (see Nayak et al. 2023; TimMerman et al. 2021) support the idea that intense psychedelic experiences reliably cause the acquisition or revision of beliefs regarding the fundamental structure of reality. They tend to nudge non-believers toward some form of theism (Griffiths et al. 2019). They inspire panpsychist or idealist beliefs regarding the nature of reality (Nayak et al. 2023; TimMerman et al. 2021). In particular, these experiences lead some to embrace a belief in a "universal" consciousness that both transcends and grounds individual minds (see Richards 2018, Ch. 4–5). Relatedly, in light of those experiences, the very existence of selves—construed as enduring, substantial entities—is sometimes questioned (see Letheby 2021, Ch. 7). To take yet another striking example, some who experience encountering telepathic extra-dimensional entities endorse the belief that said entities and the realm(s) they inhabit really do exist (Davis et al. 2020).

Many of the metaphysical conclusions that people draw from their psychedelic experiences seem to reach beyond, or be at odds with, what our common experience or our best science allows us to reasonably believe. How could a chemically induced altered state of consciousness possibly reveal metaphysical truths? Could a person be ever rationally entitled to her newfound metaphysical beliefs in light of her psychedelic experience? This seems unlikely on the face of it, especially from the point of view of philosophical naturalism. The dubious epistemic standing of psychedelic-induced beliefs casts a shadow on the value of psychedelic therapy. According to what has come to be known as the "Comforting Delusion Objection", psychedelics produce therapeutic outcomes through a mechanism that is deeply epistemically

deficient (see Letheby 2021). This, in turn, renders psychedelic therapy ethically problematic.

My aim here is to argue that updating metaphysical beliefs in light of deep psychedelic experiences does *not* have to constitute an instance of epistemic irrationality. I want to paint a picture of the role that psychedelic experiences can realistically play within a person's wider epistemic life. This picture renders psychedelic states capable of enriching one's epistemic life in a way that rationally contributes to one's attempts to make comprehensive sense of reality. That is, we can make sense of the idea that, at least sometimes, psychedelic states are capable of playing an evidential or reason-providing role (in a restricted and qualified sense to be fleshed out in Sect. 3 of this paper) with respect to metaphysical beliefs about fundamental furnishing of reality.

The idea that psychedelic states could serve as a legitimate source of metaphysical insights is not new. Traditionally, it has been championed by proponents of approaches rooted in religion or spirituality (see e.g. Huxley 2009/1954; Richards 2018). Here, I want to advance this sort of position without relying on any religious or spiritual ideas. In fact, this paper is meant to offer an alternative to two prevalent ways of approaching the epistemology of mystical states. One approach, associated with religiously-minded authors, rests on simply following the noetic force of mystical experiences to claim that when undergoing those states, one apprehends ultimate reality, thus obtaining a kind of direct, non-inferential justification for metaphysical beliefs (for roughly such a view expressed in analytic philosophy, see Alston 1991; see also Broad 1939; Plantinga 2000; Tucker 2011). The second, directly opposite approach rests on noting the inconsistencies and conflicts between the mystical-experience-acquired metaphysical beliefs and the metaphysics that underlies naturalism or physicalism, and on this basis, recommends the rejection of the former (this approach is exemplified in Letheby 2021). Here, I will avoid *both* (1) treating psychedelic states as sources of direct justification and (2) appraising the epistemic status of psychedelically induced metaphysical beliefs based on whether their contents are consistent with any particular metaphysical framework, including physicalism. Instead, the focus here is on the *cognitive processes* that underlie psychedelic-based belief acquisition and revision. My approach is to evaluate those beliefs based on whether their etiology is rational or epistemically appropriate.

The plan is as follows. To start, I will apply Laurie Paul's notion of transformative experience to interpret the core structure of psychedelic-based induction of metaphysical beliefs (Sect. 2). I will argue that psychedelic experiences are epistemically transformative and draw some general conclusions from this claim.

Next, I will focus on the epistemic status of the cognitive processes that underlie the psychedelic epistemic transformation (Sect. 3). Drawing on a predictive processing-based model of psychedelic states, I will argue that psychedelic experiences are instances of a radical epistemic exploration of one's representational state space. As such, psychedelic experiences allow people to construct their overall picture of the world by drawing from a wider repertoire of cognitions than what is usually available for a neurotypical human. This is epistemically good, I will argue, because it allows epistemic subjects to peer beyond the Sellarsian manifest image.

I will also revisit the problem at hand as it relates to the aforementioned Comfortable Delusion Objection (Sect. 4). Chris Letheby recently advocated for a view that

psychedelic-induced metaphysical beliefs are largely therapeutically irrelevant. I will argue that this view is not entirely satisfying. There is no escaping metaphysics when addressing the epistemic risks of psychedelic therapy.

In closing, I will briefly delineate some possible limitations to the optimistic view of psychedelic epistemology on offer here.

2 Psychedelic Experience as Epistemically Transformative

Let me anchor the discussion by building on a case heuristically introduced by Letheby (2021, p. 162). Imagine Mary, a person committed to a thoroughly naturalistic and secular worldview. Mary suffers from treatment-resistant depression. Driven by curiosity, but also out of growing desperation, she partakes in a scientific trial of psilocybin.[1] During a high-dose psilocybin session, her first ever psychedelic experience, Mary undergoes a mystical(-type) state. Subsequently, she is changed. For the first time in years, she feels free of her depression, serene and hopeful.

But Mary also claims to have learned something during her session. Some of the lessons she claims to have drawn from her experience pertain to personal relations with significant others. However, Mary also treats her experience as world-revealing in a way that requires adjusting her view of the structure of reality. Now, let me stipulate that this is not simply an exercise in motivated cognition. Let us assume that the opposite is the case. Mary feels internally conflicted about whether she should treat the experience as veridical. She eventually decides to shift her view of reality out of the strive for intellectual honesty. She wants to give proper justice to (what she interprets as) the deep meaning of her experience.

For now, I will not make specific assumptions about the content of the metaphysical beliefs that Mary forms. Let us only assume that the doxastic shift exhibits the worrisome directionality of straying away from a naturalistic/secular set of beliefs that acted as Mary's starting point.

Mary's case is an instance of a *transformative experience*, understood technically in light of Laurie Paul's work on this notion (Paul 2014). We may assume that before her psilocybin session, Mary had a purely intellectual interest in psychedelics and read a lot about the phenomenology of mystical states. But this left her unmoved in her metaphysical convictions. Only after her session does Mary acquire first-hand knowledge of what it is like to undergo a mystical(-type) experience. Her transformation is crucially dependent on *having* this experience.

There are two aspects of Mary's transformation. Her *personal* transformation involves Mary becoming a psychologically different version of herself, with new emotional dispositions and preferences (Paul 2014). However, Mary's *epistemic* transformation is more relevant for the present purposes. Mary's experience is epistemically transformative because it (purportedly) teaches her something she could not have learned without having that kind of experience (Paul 2014). Although she may have previously recognized purely intellectual reasons in favor of certain meta-

[1] Psilocybin is the main psychoactive compound found in "magic" mushrooms.

physical propositions, it had to take a profoundly non-ordinary subjective experience to sway Mary to endorse them.

My claim is that the notion of an epistemically transformative experience descriptively captures the basic structure of psychedelic-induced alternations of metaphysical beliefs. But more importantly for present purposes, this perspective is also helpful in highlighting some epistemologically (hence, normatively) relevant aspects of the process of the sort that Mary undergoes.

When queried about what epistemically grounds or justifies her new/altered metaphysical beliefs, the best Mary can do is to point to the phenomenology of her mystical experience. But she cannot convey this phenomenology in words. One of the core properties of mystical states is ineffability (James 2008/1902; Stace 1960).[2] These experiences are also uncommon, so for most people, it is impossible to imaginatively project onto Mary's epistemic perspective. And even those that have undergone mystical states may never be fully confident about the degree to which their own experiences resemble Mary's. Thus, most outside observers have no access to what Mary points to as a reason for her belief. This naturally raises a worry about what to make of a situation like Mary's, where no intersubjective check on purported evidence is available. But note also that it would be presumptuous for those lacking access to Mary's transformative experience to dismiss it as devoid of evidential value. Notably, this point pertains to the relation between pre-session and post-session Mary herself. In fact, following her psylocybin session, Mary may claim to be in a better epistemic position than before it. This is because after her psychedelic experience, Mary (1) gains access to a new experience that was previously inaccessible to her, while (2) she can still understand and rationally respond to all the evidence she possessed prior to her psilocybin session. Note, however, that this case also raises complications, since pre-session Mary might consider her post-session self to be epistemically compromised (Paul 2021).

It is also essential to notice that Mary's psilocybin session does not impede her autonomy or capacity as an epistemic subject. To understand this point, compare psychedelic states with monothematic delusions. The latter are often thought to constitutively involve the inability of a person to be rationally responsive to evidence against her delusion (see Bortolotti 2005; Coltheart et al. 2011). Now, notice that even if the psychedelic state itself may resemble a psychotic episode (see Carhart-Harris et al. 2016), it is transient. Following the experience, Mary regains the core judgmental or cognitive skills that characterize her in a sober state.[3] In particular, she

[2] An important question naturally arises regarding the relation between the content of the psychedelic experience and the content of beliefs formed on its basis. That psychedelic states are ineffable suggests that they possess non-conceptual content of some kind. It is thus unlikely that a psychedelic-induced belief directly derives its content from a corresponding psychedelic experience. Perhaps the content of the belief stems from an attempt to explain the phenomenology of the psychedelic state. Or maybe the contents of psychedelic-induced beliefs are heavily compressed versions of the more fine-grained experiential contents? Here, I will have to set this problem aside.

[3] Two caveats should be added to this point. First, in a population of people at risk of schizophrenia, a psychedelic experience may trigger a full psychotic episode. However, this is very unlikely in neurotypical individuals, and I assume Mary is one of them. Second, as one reviewer points out, it is an open empirical possibility that some psychedelic experiences possess such an overwhelming noetic force (feeling of truthiness or deep insight) that they effectively render a person unable to respond to counterarguments.

is able to understand and respond to any evidence against her newfound beliefs. For example, she can understand and be responsive to a debunking argument on which her beliefs may be merely powerful illusions that feel veridical but are not. Thus, if Mary decides to change her metaphysical beliefs, it is despite her *recognizing* that different parts of her evidence point in different directions. In this sense, she is not delusional. To generalize this, much of the epistemologically relevant part of psychedelic transformation takes place in a sober state, where a person faces the challenge of rationally integrating the (purported) lessons drawn from the experience with her previous worldview.

Suppose, for the sake of argument, that what Mary purportedly learns from her experience is quite bizarre. For example, she might think to have been telepathically contacted by trans-dimensional beings. Suppose also that Mary knows that there is a well-supported neuroscientific explanation that unequivocally establishes that her experience is a hyper-realistic but non-veridical fiction created by her mind. Mary's belief in telepathic entities clearly seems ill-grounded. Notice, however, that for Mary to form and maintain this belief, she would presumably have to be unable to fully comprehend or appropriately weigh the counterevidence. Given the considerations above, this would *not* be due to the psychedelic experience itself, but due to a failed sober-state integration. Mary's uncritically taking her experience at face value would have to stem from preexisting shortcomings of epistemic capacity. In a sense, the psychedelic experience would act as a trigger of an irrational belief but, so to speak, would not itself constitute the source of the irrationality.

However, there is no need to assume that Mary's doxastic shift is of this radical kind. And it might be a caricature to think that psychedelic-induced shifts in metaphysical beliefs are usually like this. Psychedelic epistemic transformations, like transformative experiences in general (Paul 2014, p. 104), come in degrees. Of course, on one side of the spectrum, they may involve dogmatic endorsements of preposterous beliefs or resemble complete religious conversions. But perhaps more often than not, they involve changes that are subtler and less epistemically risky. To illustrate this with a real-world example, consider how Rachael Petersen, a writer who participated in a clinical trial of psilocybin, recounts her reaction to a researcher asking her if, after her psychedelic experience, she still identified as an atheist:

> Suddenly, the label felt like a shirt that had shrunk in the dryer: something that served me for a time, but no longer fit. What do you call someone who believes that things are likely better than they appear, and thinks that in light of this fact we should just be kinder to one another? Someone who suspects things are more mysterious than they seem, and more connected than we'll ever know? (…) I almost dare not label these things, lest I become an idolater. All I know is that the felt sense of them keeps me company, even when I am alone. (Petersen 2019)

If this were the case, it would affect the status of some psychedelic epistemic transformations, especially the road-to-Damascus types situated toward the end of the epistemic risk spectrum.

To make this more systematic, there are three dimensions along which a psychedelic epistemic transformation might be graded. One pertains to the degree to which the new/modified beliefs *cohere with* the person's pre-psychedelic belief set. Returning to our Mary case, note that her new beliefs do not need to conflict strongly with her previous convictions. If following her experience, Mary endorsed a Buddhist-style no-self view, then this belief would presumably fit her preexisting physicalist outlook better than the belief in telepathic entities. The second dimension of grading the transformation emerges if we opt to think not in terms of categorical beliefs but rather in terms of *credences*. Suppose that Mary's epistemic transformation consists of her no longer being as confident in her physicalist/naturalist convictions and becoming more open to other possibilities. To illustrate, imagine that Mary's confidence in the proposition that physicalism is true drops from 0.95 to 0.6 (versus, say, to 0.1), and her confidence that some form of panpsychism is true shifts from 0.01 to 0.3 (versus, say, to 0.9). The third way of grading psychedelic epistemic transformations pertains to the *attitude* one takes with respect to relevant propositions. For example, assume that following her mystical experience, Mary becomes inclined to endorse theism. However, instead of acquiring a fully crystalized belief in God, Mary's is a fuzzy case of in-between believing (Schwitzgebel 2001). Some of Mary's behavioral and inferential dispositions are consistent with attributing to her a theistic belief, but others are not (for similar proposals, see Flanagan and Graham 2017; Letheby 2021, pp. 76–79).

Before moving on, it is also worth noticing that psychedelic epistemic transformations plausibly involve comprehensive and systematic changes to one's belief-updating policies. Mary's psychedelic transformation does not have to make her epistemically incoherent. Again, let me point out three ways in which these sorts of global epistemic changes may take place. One is related to overarching *epistemic goals*. For example, following her psychedelic experience, Mary may be more willing to take epistemic risks because she begins to value believing what is true more than avoiding believing what is false (James 2006/1896; Pettigrew 2022). Second, a psychedelic epistemic transformation may introduce a change in a person's *epistemic style*, that is, a unified way a person interacts with evidence (Flores 2021). For example, before her psychedelic experience, Mary may have considered drug-induced conscious states largely rationally irrelevant. Post-experience, she is much more inclined to treat the first-hand experience of such states as evidentially valuable. Third, a psychedelic epistemic transformation may affect belief-updating by targeting beliefs that constitute a *"hard core"* of one's web of beliefs (see also Letheby 2021, p. 119). For example, suppose Mary's previous deep-seated belief in naturalism is diminished or discarded following her psychedelic state. In that case, she is no longer committed to updating her beliefs in a way that must remain consistent with naturalism.

3 Psychedelics and Exploring by (Relaxed) Believing

3.1 Bayesian Brains, REBUS, and Epistemic Exploration

Treating psychedelic experiences as epistemically transformative provides a usefully nuanced conception of how psychedelics induce metaphysical beliefs. However, as it stands, this conception is incomplete as an account of how psychedelic states can have a rational bearing on one's view of the ultimate structure of reality.

For starters, notice that the epistemic transformations that Paul discusses in her seminal work involve the acquisition of beliefs about *what it is like* to have certain conscious experiences (like the ones associated with becoming a parent, participating in a war, or acquiring a new perceptual modality; see Paul 2014). These are beliefs about subjective facts, that is, facts that become accessible by virtue of a person having the relevant sort of experience. But Mary's case is different. What she purportedly learns from her psilocybin session is *not* exhausted by the knowledge about what it is like to undergo a mystical experience. Mary claims to have learned something about the structure of reality. A question arises about whether and how a psychedelic state could be world-revealing.

Relatedly, notice that the mere fact that one had a purportedly world-revealing transformative experience does not by itself establish a rational link between this experience and the resulting beliefs. Consider two ways in which an epistemically transformative experience might fail to provide good reasons for belief. One involves brute transformations that bypass cognitive processes of the sort that could be epistemically relevant. Think of acquiring beliefs through participating in a *Clockwork Orange*-style aversion therapy. The second type of faulty transformation is epistemically evaluable but also epistemically bad. Think of Bayesian models of delusions acquired during a psychotic episode (see Fletcher, Firth 2009 for a classic exposition). These models take delusions to be formed and maintained through (unconscious) Bayesian inference. However, this inference is epistemically compromised by systematically overestimating the evidential value of sensory data.[4]

I want to argue that psychedelic epistemic transformations are neither brute nor bad in this way. For this purpose, I will take an epistemologically-oriented look at the dominant scientific model of the cognitive underpinnings of psychedelic states. This is the REBUS model, dubbed using an acronym for "RElaxed Beliefs Under pSychedelics" (Carhart-Harris and Friston 2019). I will argue that under the REBUS model, psychedelic states are capable of being rationally integrated with one's belief system.[5] Later, in Sect. 3.2 and 3.3, I will return to the question of psychedelic experiences as world-revealers.

[4] Notice that I am using this view of delusions for illustrative purposes, without necessarily endorsing it.

[5] It is worth noting that REBUS is largely continuous with historically preceding theories of psychedelic states (Swanson 2018). For example, it might be read as a Bayesian incarnation of Aldous Huxley's seminal idea that psychedelics act by opening up a cerebral filter or "reducing valve" (Huxley 2009/1954). Although I will not pursue this here, the present epistemological proposal may generalize to cover alternative approaches that can be interpreted as variants of the "filter approach" (like the thalamocortical model and the claustral model; see Doss et al. 2022 for review). Hence, the argument on offer here should be of value even for readers who do not subscribe to the REBUS model. However, from the point of view

The REBUS model employs the idea that the brain is a Bayesian system. In particular, it rests on the predictive processing (henceforth PP) variant of the Bayesian brain hypothesis. PP has become a major theoretical framework within contemporary cognitive (neuro)science, so I expect many readers to be at least cursorily familiar with it. Because of this, and due to the limitations of space, what I provide below is a bare-bones summary of PP (for up-to-date introductions, see Hohwy 2020; Parr et al. 2022).

The crux of PP lies in treating the brain as an "inference machine" comprised of hierarchically organized information processing mechanisms. Each level of the hierarchy has one computational goal: to calculate the posterior probability of some hypothesis or estimate, $p(h|d)$, through approximating Bayesian inference. The hierarchy encodes a generative model whose function is to capture the world's nested causal structure and the way this structure produces patterns of sensory data in the organism. At each level, the generative model encodes a joint probability of hypotheses and lower-level (ultimately, sensory) data, $p(h, d)$, which is factorizable into a product of the prior, $p(h)$, and the likelihood, $p(d|h)$. Equipped with the generative model, the brain is thought to engage in a variational approximation of Bayesian inference. Roughly, the brain uses the model to generate an estimate of the state of the world and iteratively brings this estimate closer to a true posterior that an exact Bayesian inference would yield (under the model). Mechanistically, this is realized by a bidirectional information flow comprised of top-down prediction signals and bottom-up prediction error signals. The task is to minimize average prediction errors across the hierarchy.

I will assume that to the degree that PP captures the computational structure of cognition, it renders this structure epistemically evaluable. This is because, under PP, the causal evolution of the brain's representation of the world approximately conforms to a rational rule of inference (Gładziejewski 2021). For example, in the case of perception, PP renders perceptual states epistemically evaluable by treating them as inferentially derived from prior and likelihood distributions encoded in the generative model (Gładziejewski 2021). This effectively constitutes a variant of Susana Siegel's "rationality of perception" approach (Siegel 2017).

I will further assume that the processing invoked by PP can remain epistemically relevant under a systematic disruption. Imagine a procedure that tinkers with the values of priors encoded in the generative model or with how precise or reliable these priors are estimated to be. Now, as long as the processing remains approximately Bayesian under this procedure, it remains epistemically appraisable (hence, epistemically relevant).

One last assumption I will be making going forward is that the processes invoked by PP remain epistemically appraisable when they run off-line. I am referring here to the idea that the generative model can become decoupled from the current sensory stream and used to run internal simulations of non-actual scenarios (Williams 2021). Such processing is rationally constrained by the generative model because the flow of the simulation is determined by the model-encoded assumptions: the brain samples

of this paper, REBUS's attraction lies in its direct connection to the Bayesian view of brain functioning, which renders it especially amenable to epistemological analysis.

sequences most-likely-under-a-model. I take this type of off-line processing to be epistemically evaluable (for a view that makes a similar case independently of PP, see e.g. Myers 2021).

Now, onto the REBUS model. PP treats perception as an interplay between top-down cognitive structures and bottom-up error signals. The relative degree to which processing is determined by prior knowledge and the incoming error signals is flexibly determined through precision estimation. Precision measures the inverse variance of priors and error signals, effectively tracking their relative reliability. Now, by acting on the serotonin 5-HT_{2A} receptors of deep pyramidal cells in the cortex (which are thought to encode priors), psychedelic compounds decrease the precision of priors (Carhart-Harris and Friston 2019). This way, they reduce the degree to which prior cognitive structures normally constrain and regulate cognitive activity. Hence, the priors harbored in the brain's generative model become "relaxed".

Relaxing of priors affects the inferential processes involved in perception, inducing perceptual distortions and hallucinations. However, perceptual changes are a relatively unimportant feature of the phenomenology of acute psychedelic states that I am interested in here. (Quite literally, in therapeutic settings, participants in psychedelic sessions are usually invited to go through the experience with their eyes closed). These states are dream-like or imagination-like in the sense of being endogenously constructed. So, in light of REBUS, deep psychedelic states consist of a free flow of internal constructions that is unconstrained, or significantly less constrained, by the priors that regulate regular cognitive operation. Upon entering a psychedelic state, the dysregulated generative model runs off-line in a way that allows it to venture into new regions of the representational state space.

Importantly, this psychedelic-induced cognitive disruption does not undermine the status of the brain as a Bayesian system. On the REBUS model, the psychedelic-influenced brain remains Bayesian, even if it runs on relaxed priors. Arguably, this is reflected in phenomenology: far from being experienced as a random mess of cognitive junk, deep psychedelic states are usually described as possessing an intelligible (even if ineffable) structure. So, given all these considerations, I propose that psychedelic epistemic transformations are not brute. They reside within the domain of normatively relevant cognition.

My further claim is that acute psychedelic states possess the power to improve the epistemic standing of one's beliefs. What does this rational role consist of, exactly? The claim is *not* that it can be pinned down to a particular chain of justification-conferring inference. I do not mean to suggest that the processing that takes place in a psychedelic state is somehow epistemically better than normal cognitive operation. Rather, to understand my point, we need to take a wider perspective on the epistemic life of a person, and on how transient, psychedelic-induced disruptions of this life can be epistemically beneficial in the long run. To a first approximation, the claim is that one can improve the epistemic standing of one's conception of oneself and the world by occasionally wandering off the beaten cognitive path to consider previously unconceived alternatives. Psychedelics constitute a tool that reliably elicits such an exploratory mode of cognition.

This view of psychedelic epistemology is already discernible in how Robin Carhart-Harris and Karl Friston (2019) frame the REBUS model. To explain how psy-

chedelic states can be therapeutic, they point to the process of relaxing the precision of "pathologically overweighted" priors. These are psychologically harmful priors so deeply entrenched into one's internal model that they become resistant to revision. When this is the case, one's model of the world gets stuck in a local minimum of the free energy (prediction error) landscape. By relaxing the precision of priors, which is equivalent to opening or flattening the free energy landscape, psychedelics allow the model to get instantly "unstuck". To give a high-level example, in a psychedelic state, priors that underlie one's narrative self-conception may become relaxed. This, according to REBUS, allows one to turn a debilitating self-narrative into a new, positive, but also more realistic conception of who one is. Now, there is an *epistemic* side to this story as it is plausible to assume that the newly gained self-conception is more well-grounded or accurate than the previous one (see Letheby 2021, Ch. 8).[6]

Let me now try to make this more systematic. Aronowitz (2021) has recently proposed that we should extend our concept of epistemic rationality in a way that takes into account an exploration/exploitation trade-off that arises at the level of belief. The exploration/exploitation problem originally applies to adaptive action. Roughly, the idea is that organisms face the choice between exploiting existing strategies for successful behavior and trying out new strategies that could, with some small but non-negligible probability, prove even more successful in the long run (should I order my favorite dish at a restaurant or risk trying out something new?). In a complex world where the reward function cannot be known in advance, it is good to adopt a strategy that mixes exploitation with some degree of exploration.

Aronowitz proposes that we face a structurally similar exploration/exploitation dilemma in our epistemic lives, in the following sense. At least some of the things we learn, or novel hypotheses we invent, we arrive at through the exercise of mental simulation. How these internal simulations unfold is guided by our beliefs (see also Myers 2021; Williams 2021). However, beliefs not only guide but also limit our mental constructions in epistemically relevant ways. That is, we may miss some relevant possibilities because they are obstructed by the beliefs that constrain, often implicitly, the imaginative search. So, it may be epistemically valuable to sometimes adopt other beliefs even if they are undersupported by current evidence.[7] This is because the different beliefs may guide the simulation-based inquiry to previously hidden truths. The upshot: to create an opportunity for learning and accuracy in the long run, it is rational to mix some degree of epistemic exploration into one's epistemic life.

I propose that this story can be extended to cover the rational role of psychedelic states. According to the REBUS model, psychedelics elicit a sort of exploratory simulation in which alternate priors are flexibly "tried out". In fact, it may be said that

[6] Note that by defending the idea of "relaxed" Bayesian processing as providing an epistemically relevant etiology for beliefs, I am offering an explicit rationale for an approach that seems to be already implicitly endorsed (and applied to non-metaphysical beliefs) by Letheby (2021). In many ways, this paper aims to detail how this sort of approach can be successfully extended to defend an optimistic view of psychedelic epistemology of metaphysical beliefs, which is the opposite of the view that Letheby himself endorses.

[7] As an example, Aronowitz invites us to consider the epistemic "double life" of Ivan Pavlov, who would switch between conducting scientific research during the academic year and immersing himself in spiritual literature during the summer. "We could imagine that these two ways of living came along with two ways of thinking." (Aronowitz 2021, p. 339).

psychedelic states constitute a *radical* form of exploration. Aronowitz's original proposal is oriented towards epistemic exploration insofar as it is guided and constrained by personal-level beliefs. Psychedelics reach deeper into cognitive architecture by targeting subpersonal priors. As will transpire in the next subsection, this includes priors that constrain our imagination in ways relevant to the metaphysical inquiry. In any case, the basic normative upshot of Aronowitz's proposal still stands in the present context. It is rational to sometimes put oneself in an epistemically risky position that provides an opportunity to learn something new and otherwise unavailable. The rational role of psychedelic states consists of eliciting bursts of such exploration.

Importantly, this is not to say that mere engagement in the psychedelic exploration of consciousness automatically generates warrant for metaphysical beliefs. The point is that epistemic exploration through psychedelic states is valuable even if many (perhaps most) of the new cognitions turn out of little epistemic value (see McGovern et al. 2023). For comparison, at the level of action, most acts of exploration may fail to bring about outcomes that outperform the exploitation of previously leaned policies. Still, exploration is valuable in virtue of enabling the learning of novel policies over longer time spans (and only when appropriately mixed with the exploitation of previously learned action policies). Similarly, the value of epistemic exploration through psychedelic states lies in how it enriches the inquiry process, where the subject actively seeks and evaluates evidence about a subject matter *before* stable beliefs are crystallized (Friedman 2019). For illustration, consider the physicist Carlo Rovelli, who recounts how early experiences with LSD guided his thinking about the nature of temporal passage:

> It was an extraordinarily strong experience that touched me also intellectually. Among the strange phenomena was the sense of time stopping. Things were happening in my mind but the clock was not going ahead; the flow of time was not passing anymore. It was a total subversion of the structure of reality. (…). And I thought: "Well, it's a chemical that is changing things in my brain. But how do I know that the usual perception is right, and this is wrong? If these two ways of perceiving are so different, what does it mean that one is the correct one?" (From Higgins 2018).

3.2 The Epistemic Value of Exploring Beyond the Manifest Image

Now is the time for the discussion to directly connect with metaphysics. Can exploring non-ordinary forms of conscious experience deliver results that are evidentially or justificationally relevant for beliefs regarding the ultimate structure of reality? For starters, let us recognize the evidential role that *ordinary* experience plays in metaphysical inquiry. Think in particular of visual perceptual experience which purports to reveal a world that flows through time, furnished with macroscopic objects occupying determinate positions in a three-dimensional space. Visual experience also seems to spring from a subjective point of view, arguably revealing an experiencing self that flows through time along with the experienced world. In this fairly innocent

sense, already implicitly embedded in ordinary experience, there is the commonsense metaphysics of the manifest image (Sellars 1963).

Ordinary experience serves as a source of defeasible data for metaphysical inquiry (Benovsky 2015; Goldman 2015; Paul 2012). In some cases, the relevant aspects of ordinary consciousness become *explicitly* invoked as potentially decisive in settling a metaphysical problem. Take the debate regarding the reality of the passage of time. Realists about temporal flow think that there is an evermoving, metaphysically privileged present moment and that whatever is located in the past or the future relative to this moment lacks proper existence. Why should we believe such a view? Here, realists often point to the temporal structure of ordinary experience (see Paul 2010 for discussion). According to this line of thinking, what justifies the belief that time flows is that the sense of temporal passage is a pervasive aspect of conscious experience.

Consider also how imagery can be evidentially relevant for metaphysics. Here, by "imagery", I mean the capacity to re-use our perceptual and action-guiding machinery for off-line cognition. It has been argued that such internal simulations can act as a source of *modal* knowledge (see e.g. Gregory 2020; Kung 2010; Williamson 2007). Very roughly, the idea is that (un)imaginability can be a guide to (im)possibility. Think of how a modal fact about consciousness itself might be claimed to be discovered through an exercise of imagination. I find myself utterly unable to imaginatively simulate a timeless mental state, a state that lacks the sense of temporal flow. From this, I might infer that atemporal consciousness is impossible. This, in turn, might lend additional modal weight to the "argument from experience" regarding the objective passage of time. On this construal, by pointing to our experience as evidence in favor of objective temporal flow, we are not simply pointing to a contingent fact about the conscious experience of a particular hominid species, but to a structure that any conscious experience must necessarily possess.

This is where PP and the REBUS model may reenter the picture. According to PP, the metaphysically relevant aspects of ordinary perceptual phenomenology—the self, the flowing time, the space, and the ordinary objects contained in it—are generated by the inferential machinery that underpins perception. The idea is that at the highest levels of the generative model, the brain stores abstract "hyperpriors" that put very general ("almost Kantian", Clark 2016, p. 174) constraints on the cognitive activity at lower levels. To make the discussion tractable, let me narrow the focus to the priors that are thought to underlie the sense of self and temporal passage. The sense of self has been theorized to result from a process in which the brain infers a single endogenous cause underlying short-term correlations in body-related signals (the embodied aspect of experienced selfhood) and long-term sensory patterns (the narrative aspect of self; see Hohwy and Michael 2017; Letheby 2021, Ch. 7; Letheby and Gerrans 2017). The sense of temporal passage has been argued to be grounded in a high-level prior expectation of the world and the sensory signal caused by it to be constantly changing (Hohwy et al. 2016).

Now, according to the REBUS model, psychedelic compounds target these metaphysically relevant hyperpriors, relaxing them—sometimes to the point of altogether dissolving them—without eliminating the conscious experience itself. This way, they elicit an experience that differs from everyday consciousness at the level of core organizing principles. But more pertinently for the present purposes, psychedelic states

are *evidentially* relevant because they broaden the range of data that can guide and constrain metaphysical inquiry. Insofar as we draw on the structural properties of consciousness to inform our metaphysics, psychedelics allow us to pool data that go beyond the confines of the default structure of experience.

In at least two ways, this widening of available evidence could undermine or problematize metaphysical views rooted in ordinary experience. For one thing, psychedelic states may directly falsify modal claims grounded in (ordinary) imagination. Once the prior underlying the sense of temporal flow is discarded, one can venture into a region of one's representational state space that would typically be inaccessible. Quite literally, one gains the capacity to construct an experience that lacks the sense of temporal flow. This way, psychedelic experiences of "eternity" falsify the idea that the sense of temporal passage is a necessary structure of any conscious experience. This, in turn, strips the experience-based argument for the objective temporal passage from its modal force.

Another way in which psychedelic states can play this sort of commonsense-undermining role is by providing first-person data in support of the idea that certain core features of normal experience are not innocently given but are actively (even if unconsciously) constructed (see also Letheby 2021, Ch. 7). For example, by undergoing a psychedelic state in which one's very sense of self is dissolved, one arguably learns that one's usual sense of self is a result of interpretation or (unconscious) inference. Now, the fact that the experience of being a self is somehow constructed does not by itself establish it as non-veridical. But it invites a view on which the very sense of self is *epistemically assessable* depending on the nature of processes that give rise to it. It can no longer act as an unjustified justifier of metaphysical claims about the self (Gładziejewski 2021). To generalize, following a deep psychedelic experience, it arguably becomes harder to remain foundationalist about the posits of the manifest image.

But I think that the value of psychedelic states goes beyond undermining or problematizing existing evidence. Psychedelic experiences also provide new data. Think of a role that psychedelic-induced mystical(-like) states could play when considered not in and of themselves, but when embedded in a wider discursive structure. Let us focus on a particular case of cosmopsychism, a monistic version of panpsychism or idealism. Roughly, the view is that consciousness is metaphysically fundamental and that the world, including individual human subjects, is grounded in a single, universal consciousness. Although still somewhat fringe, this view has recently gained some traction among philosophers (see e.g. Albahari 2019; Shani 2015; see also essays in Seager 2019).

Cosmopsychism is often theoretically motivated by the Hard Problem of Consciousness and advertised as a less problematic alternative to standard forms of panpsychism. But it faces hefty philosophical problems of its own. Can we make sense of consciousness that extends beyond individual subjects of experience? And how can such an all-encompassing consciousness ground or "decombine" into the multitude of apparently separate minds? Some authors have suggested that (1) universal consciousness is aperspectival, lacking any subject/object division, and (2) individual subjects of experience can be deflated as somehow illusory, whereby subjective points of view arise from relations between experiences rather than by being related

to persisting individual entities or "selves" (for relevant discussions, see Albahari 2019; Chalmers 2019).

The problem with this proposal is that it is not easy to make concrete sense of the theoretical posits it invokes (see Chalmers 2019, p. 367). Of course, a philosopher may posit on purely theoretical grounds that there are forms of consciousness that are "non-dual" with respect to subject/object division, thus constructing an internally coherent picture of ultimate reality. But there is a worry that independent reasons or data should be provided to "externally" validate the central theoretical posit, i.e., that there can be phenomenal states that transcend subject/object duality. This is where altered states of consciousness become relevant. Through dissolving priors that underpin the sense of self, psychedelic compounds can elicit states of consciousness that are unstructured along the subject/object distinction. As such, these experiences play crucial evidential roles in the present context by providing missing data. On a weaker reading, they make the very idea of cosmopsychism *conceivable*: "If consciousness can conceivably be experienced as aperspectival and unconditioned, then, being inherently experiential, it will conceivably *be* aperspectival and unconditioned" (Albahari 2019, p. 14). But on a stronger reading, these experiences can be interpreted as directly *confirming* the belief in conscious states with non-dual phenomenal character, which is a significant improvement over postulating such states based solely on theoretical considerations. This, by itself, does not establish that non-dual consciousness *grounds* individual conscious subjects. However, at least it adds support for the belief in the existence of the former, more exotic relatum of this postulated grounding.

Before I move on, let me mention one last potential epistemic benefit of psychedelic states. By their nature, revisionary positions in metaphysics tend to be misaligned with the ways in which humans spontaneously experience and conceptualize the world and themselves. For illustration, consider again the question of the existence of the self. Albahari (2014) discusses a case of a person who (1) forms, on purely theoretical grounds, a "reflective" belief that there are no selves, but (2) her conscious experience remains subjectively centered around a stable, continuous self, resulting in a cluster of cognitive, emotional and behavioral dispositions (for example, self-related anxieties) that Albahari calls "action-based" belief. Now, Albahari argues that accomplished Buddhist meditators systematically alter their consciousness to gain direct experiential insight into no-self. This, in effect, brings them into a state of increased cognitive coherence, whereby their spontaneous dispositions no longer contradict their reflective beliefs about the (purported) illusion of self. I think that psychedelic states can play a similar role by enabling a person to achieve cognitive coherence with respect to views about, for example, the non-existence of selves or the existence of non-dual forms of consciousness.[8]

[8] Because psychedelic states (including subjectively selfless or non-dual states) change consciousness only temporarily, their role will probably be much more limited in this respect as compared to meditation, which can bring about stable shifts in experience. That is, following even the most intense psychedelic experience, a person may find herself gradually drifting back into her default set of cognitive and emotional dispositions. Still, it needs to be noted that some studies suggest that acute psychedelic states can lead to more persistent changes of (*inter alia*) self-related aspects of experience (see Orłowski et al. 2022).

3.3 Revealing Truth?

There is one more issue that requires addressing—that of the *veracity* of acute psychedelic experiences. Skeptics about the metaphysical import of psychedelic states may still question the idea that these states could reveal metaphysical facts. Let me now sketch out a view of how acute psychedelic states could be on-track with respect to at least some metaphysical truths.

We need to distinguish two models of how psychedelic states could be truth-revealing. The first model—the "third eye" model—treats psychedelic states as akin to perception, construed along foundationalist lines. According to this picture, in a deep psychedelic state, one's "mystical sense" or "metaphysical truth detector" is opened to put one in direct epistemic contact with ultimate reality. Instead, I want to suggest the alternative, the "dispelling-the-illusion" model. This view rests on the assumption that at least some cognitive structures that give rise to the manifest image are systematically off-track with respect to metaphysical truth. In a psychedelic state, those truth-obstructing cognitive structures are removed, allowing one to enter a conscious state that better aligns with how the world is.

Here is how the dispelling-the-illusion model might work. We may start by noting that there is a growing body of work in philosophy and cognitive science defending the view that default perceptual and cognitive structures are off-track with respect to metaphysical truths (see e.g. Benovsky 2015; Goldman 2015; Ladyman and Ross 2007, Ch. 1; Korman 2019). This may mean that at least some aspects of default cognition are either (1) limited to selectively revealing only truths that are available from the perspective of a particular type of organism, like a human (weaker claim), or are (2) altogether off-track with respect to metaphysical truth (stronger claim). In any case, the core assumption of the dispelling-the-illusion model is not implausible. Hence, by relaxing or dissolving default priors that constrain perception, psychedelics could remove cognitive structures that normally obstruct the truth. Instead of opening up a new, mystical perceptual-like "modality", psychedelics disrupt the existing inferential machinery in a way that may bring it closer to accurately capturing (the relevant parts of) reality.

For illustration, consider again two aspects of ordinary experience that shape the manifest image: (1) the experience of temporal flow and (2) the experience of one's perception and conscious thoughts as originating from and being centered around a persisting individual self. Three types of considerations favor viewing those aspects of ordinary consciousness as illusory. First, and weakest, general evolutionary considerations dissociate the adaptiveness of a cognitive mechanism from its ability to represent the world truthfully. Second, there are plausible explanations of the subjective sense of temporal flow and selfhood on which neither of those experiences reflects the way the world really is. The experience of temporal flow can be explained by positing a temporally ordered (but not temporally flowing) sequence of states whose phenomenal character generates the illusion of temporal passage (Paul 2010; Price 1996, pp. 14–15; Le Poidevin 2007). The experience of being a persisting self can be explained by a binding process that gives rise to the sense of being a simple, substance-like entity without actually tracking any such entity (Letheby and Gerrans 2017). Third, and perhaps most importantly, there are strong scientific and philo-

sophical cases to be made against temporal passage and selves. For temporal flow, forceful arguments have been proposed in favor of eternalism, the view that denies that the world fades and becomes as the window of "now" moves. Eternalism has been repeatedly argued to find support in "block" models of the universe rooted in relativistic physics (Barbour 1999; Carroll 2010; Price 1996). For self, philosophical arguments have been put forward against the view that there exists an enduring entity that corresponds to what people ordinarily identify as selves (Nāgārjuna 1995; Parfit 1995).

So the point is that by relaxing or dissolving default priors that constrain perception, psychedelics could be removing cognitive structures that normally obstruct the truth. Instead of opening up a new, mystical perceptual-like "modality," psychedelics disrupt the existing inferential machinery in a way that may bring it closer to accurately capturing (relevant parts of) reality. I think there are strong reasons to think that this happens during the psychedelic-induced dissolution of the experience of being an individual self, flowing through time. Under eternalist and no-self metaphysics, selfless and timeless experiences are better aligned with ground truth about reality.

However, even given all of this, there is a substantial worry to be raised here[9]. Why not simply claim that psychedelic states make people bump from one illusion to another (see McGovern et al. 2023)? Or, why regard such states as somehow geared towards revealing truth rather than, at most, moving us from falsehood to truth by sheer epistemic luck?

Here, I do not have a definitive answer. However, let me sketch out three possible directions in which one could proceed here (two of which arguably revert us to the third-eye model of psychedelic epistemology). One would be to sidestep truth as such and go foundationalist in one's epistemology, arguably sticking in this respect to mystical traditions. Roughly, the claim would be that, ultimately, any knowledge can only be grounded in conscious experience, like perception or intellectual intuition. In the present context, this basic view could be combined with the notion that what imbues specific experiences with the power to act as an ultimate source of justification, absent defeaters, is their noetic quality or subjective sense of truthiness (see e.g. Tucker 2011). Then, one could claim that the noetic quality of mystical experience outstrips other forms of experience (save perhaps mathematical intuition). Hence, *all else being equal* (i.e., absent defeaters), mystical-experience-based beliefs have a stronger claim for being true or constituting knowledge than (almost) any other type of belief. The downside of this option is the problematic epistemology it is based on (Gładziejewski 2022), including the crucial issue of whether all else is indeed equal in the particular case of mystical states.

Another option would be to construct a more comprehensive philosophical view of mystical experience that renders it more "truthy" than ordinary experience. If reality is ultimately grounded in universal consciousness, then experiences of cosmic unity could count as cases in which an (apparent) individual discovers her deeper, true identity – by dissolving into it. On such a picture, there seems to be a direct connection between the content of the mystical experience and the ground truth about

[9] I thank the reviewers for pressing me to explicitly address this issue.

reality. The obvious drawback here is that such treatment *presupposes* the very metaphysical view that the non-ordinary experience is supposed to support.

However, I think that the most reasonable option is much more modest. No experience, ordinary or altered, carries with itself the information about its ultimate origins (consider radical skeptical scenarios). To ascertain the value of experience, we need to look beyond it. Whatever evidence non-ordinary experience delivers, the move from experience to belief should be mediated by consulting other strains of evidence. That is, the crucial question is whether what psychedelic states deliver converges with what is supported by other lines of inquiry. By "other lines of inquiry". I mean science and philosophical arguments of a more a priori kind. In some cases (e.g., the self, temporal passage, and arguably also cosmopsychism), we may find non-trivial convergence. In other cases, there is no discernible convergence, and the move from experience to belief requires one to overrule other forms of evidence or reach beyond what they allow to rationally believe. Trans-dimensional DMT entities may belong to this category. Such latter cases are epistemically problematic in a way that the former are not. For those who are compelled by the idea of psychedelic-induced mystical states as epistemically privileged gateways to ultimate truths, this sort of coherentist picture might sound too minimalistic. However, I suspect this is as much as we can get as long as we remain loyal to considerations of epistemic rationality.

4 Psychedelic Therapy Without Metaphysics?

At the outset of this article, I related the issue at hand to the Comforting Delusion Objection (henceforth, CDO) to psychedelic therapy. If the present argument is on the right track, we have reasons to think more charitably about (many) psychedelic-induced metaphysical beliefs, which arguably neutralizes the objection to some degree. Before closing, I want to address a possible concern that there are more parsimonious ways of dismantling the CDO, ones that bypass any commitments regarding the psychedelic states as possible revealers of metaphysical truths. A view of this sort has recently been advanced by Chris Letheby.

Letheby (2021, Ch. 8) carefully examines existing empirical evidence and concludes that it fails to support the idea that the acquisition of metaphysical beliefs constitutes the main mechanism of psychedelic therapy. Instead, the therapeutic work is done by psychedelic-induced cognitive states that do *not* come burdened with metaphysical baggage. These states include insights into one's personal life and relationships, new ways of apprehending previously known facts, increased flexible control of attention, and acquaintance with the fact that one's repertoire of conscious states outstrips what is available in everyday experience.

I think that Letheby's proposal goes a long way toward neutralizing the CDO. Nonetheless, there are reasons to doubt that his view lets psychedelic therapy *completely* off the epistemic hook.

First, note that optimal therapeutic outcomes are achieved with the administration of a large dose of a psychedelic substance.[10] But larger doses are also more likely to produce intense, mystical types of experience. Thus, metaphysical beliefs may constitute a natural side-effect of psychedelic therapy. Relatedly, it seems plausible that the "personal-insight-related" and the more "metaphysical" aspects of a psychedelic state—even if they can be dissociated using psychometric tools (see Letheby 2021, Ch. 4)—tend to be tightly interwoven within the *experience itself*.

Second, there is substantial empirical evidence for the claim that mystical experiences *do* have beneficial effects on mental well-being (see Ko et al. 2022; Sjöstedt-Hughes 2023). This is corroborated by self-assessments, as people often attribute the therapeutic benefits to their newfound outlooks on the nature of reality (see Richards 2018). In many cases, it is easily understandable how metaphysical beliefs could affect one's psychological well-being. Think of the belief in a benevolent God or the belief that (universal) consciousness transcends death. But, contrary to a view expressed by Letheby (2021, p. 122), this point might even apply to more abstract metaphysical aspects of psychedelic states. For example, take depression, which involves the experience of being stuck in time and particularly being cut off from one's future (see Ratcliffe 2015; Whitteley 2021). Speculatively, the temporal phenomenology of acute psychedelic states might allow depressed patients to obtain a *sub specie aeternitatis* look at the world and their lives, perhaps allowing them to espouse a more positive conception of their own temporality.

Importantly, the observations made above apply even if Letheby is right in claiming that metaphysical insights are not the main therapeutic mechanism of the psychedelic experience. In this sense, my approach to addressing the CDO in this paper is complementary to Letheby's. However, I think that the metaphysically relevant aspects of the psychedelic state either play *some* nontrivial supplementary therapeutic role or are reliably correlated with the therapeutically relevant aspects of the experience. Either way, there is no escaping metaphysics in the context of psychedelic therapy (see also Sjöstedt-Hughes 2023).

5 Conclusions

The following view of the epistemology of deep psychedelic experiences emerges from this paper. Some of our metaphysical convictions—both for the folk and professional metaphysicians—are evidentially rooted in ordinary ways of experiencing the world. Psychedelics allow epistemic subjects to obtain experiences whose core structure differs from that of ordinary states of consciousness. Thus, deep psychedelic states radically and transiently broaden the range of cognitions that could inform metaphysical inquiry. As such, psychedelic states can provide otherwise unachievable epistemic benefits: (1) they can block arguments "from experience" in favor of certain commonsense metaphysical claims, (2) they can directly challenge (in)

[10] For example, a recommended therapeutic dose of psilocybin is 25 mg (see Garcia-Romeu et al. 2021). This is roughly the amount of psilocybin contained in 3–4 g of dried *psilocybe cubensis* mushrooms, which lies around the range of a "strong" dose (https://erowid.org/plants/mushrooms/mushrooms_dose.shtml).

conceivability claims (and their purported modal consequences) that figure in metaphysical debates, and (3) provide extra support for metaphysical projects that rely on positing conscious states with non-ordinary phenomenal character (like non-dual or selfless forms of consciousness). Furthermore, because our default ways of experiencing can be either limited or altogether off-track with respect to metaphysical truth, at least some aspects of psychedelic experiences can count as dissolving such truth-obstructing cognitive structures. Crucially, psychedelic states can be transformative in a way that does not impede one's epistemic autonomy, allowing one (when in a sober state) to rationally integrate the fruits of psychedelic exploration with one's preexisting beliefs. Given all this, I think it is reasonable to think that at least some instances of psychedelic-induced metaphysical beliefs are more than comfortable delusions.

There is an urgent caveat, however. What I have provided in this paper is unlikely to succeed as a *complete* epistemology of psychedelic-induced metaphysical beliefs. There are potential limitations to the optimistic picture I have painted. Before closing, let me briefly discuss them.

One concern regards the simple fact that psychedelic states can involve epistemologically relevant impairments of cognitive function (see Bayne and Carter 2018; Letheby 2017, Ch. 8; McGovern et al. 2023). This is most obvious in the case of visual perception, which gets heavily distorted in the psychedelic state. Deep feelings of having true insights into the nature of reality may simply stem from faulty metacognitive reality monitoring (see Simons et al. 2017) or from psychedelic-induced shifts in precision weighting of priors (McGovern et al. 2023). So, for example, the experience of encountering telepathic entities may plausibly be an extremely immersive hallucination associated with a deep but erroneous sense of reality. Thus, while some aspects of the psychedelic state may be useful in rationally guiding the metaphysical inquiry, others may turn out epistemically useless or misleading.

The second problem relates to the question of the degree to which psychedelic-induced metaphysical beliefs are convergent across individuals. What should we make of a situation in which, following a psychedelic experience, some people end up believing p, while others not-p? If this doxastic difference is explained by the different experiences people have, then we would need to decide whose experiences are more credible. Or perhaps the experiences themselves are similar across individuals, but they are interpreted differently? But then we face the problem of deciding which interpretation is right or closer to being right.

The third concern, closely related to the one described above, pertains to the fact that psychedelic experiences are likely influenced by preexisting expectations (McGovern et al. 2022). For example, coming into the psychedelic experience expecting to obtain profound insights into the nature of God might shape the experience in such a way that it seemingly delivers profound insights into the nature of God. This looks like an epistemically bad sort of cognitive penetrability. Furthermore, differing expectations might explain the possible divergences in psychedelic-induced beliefs. This mirrors the concerns that have been historically raised over the evidential status of mystical experiences in general (Katz 1978). It remains to be seen whether, or to what extent this sort of cognitive penetration explains the psychedelic-based acquisition and revision of metaphysical beliefs.

Acknowledgements I would like to thank the three referees for this journal for their helpful suggestions and insightful criticisms of the previous version of this paper. I benefitted a lot from the feedback I received when presenting this work at the Centre for Consciousness and Contemplative Studies at Monash University (special thanks to Jakob Hohwy). Work on this paper was supported by the Polish National Science Centre under the Opus grant 2019/33/B/HS1/00677.

Declarations

Conflict of Interest None.

References

Albahari, M. 2014. Insight knowledge of no self in Buddhism: an epistemic analysis. *Philosophers' Imprint* 14(21): 1–30.

Albahari, M. 2019. Perennial idealism: a mystical solution to the mind-body problem. *Philosophers' Imprint* 19(44): 1–37.

Alston, W. 1991. *Perceiving God: the Epistemology of Religious Experience*. Ithaca: Cornell University Press.

Aronowitz, S. 2021. Exploring by believing. *The Philosophical Review* 130(3): 339–383.

Barbour, J. 1999. *The end of Time: the Next Revolution in Physics*. Oxford: Oxford University Press.

Bayne, T., and O. Carter. 2018. Dimensions of consciousness and the psychedelic state. *Neuroscience of Consciousness* 2018(1): niy008.

Benovsky, J. 2015. From experience to metaphysics: on experience-based intuitions and their role in metaphysics. *Noûs* 49(3): 684–697.

Bortolotti, L. 2005. Delusions and the background of rationality. *Mind & Language* 20(2): 189–208.

Broad, C. D. 1939. Arguments for the existence of God, II. *Journal of Theological Studies* XL(1): 156–167.

Carhart-Harris, R. L., and K. J. Friston. 2019. REBUS and the anarchic brain: toward a unified model of the brain action of psychedelics. *Pharmacological Reviews* 71(3): 316–344.

Carhart-Harris, R. L., M. Kaelen, M. Bolstridge, T. M. Williams, L. T. Williams, R. Underwood, A. Feilding, and D. J. Nutt. 2016. The paradoxical psychological effects of lysergic acid diethylamide (LSD). *Psychological Medicine* 46(7): 1379–1390.

Carroll, S. 2010. *From eternity to Here: the Quest for the Ultimate Theory of Time*. New York, NY: Dutton.

Chalmers, D. 2019. Idealism and the mind-body problem. In *The Routledge Handbook of Panpsychism*, ed. W. Seager. 353–373. New York: Routledge.

Clark, A. 2016. *Surfing uncertainty: Prediction, Action and the embodied mind*. Oxford: Oxford University Press.

Coltheart, M., R. Langdon, and R. McKay. 2011. Delusional belief. *Annual Review of Psychology* 62: 271–298.

Davis, A. K., J. M. Clifton, E. G. Weaver, E. S. Hurwitz, M. W. Johnson, and R. R. Griffiths. 2020. Survey of entity encounter experiences occasioned by inhaled N,N-dimethyltryptamine: Phenomenology, interpretation, and enduring effects. *Journal of Psychopharmacology* 34(9): 1008–1020.

Doss, M. K., M. B. Madden, A. Gaddis, M. B. Nebel, R. R. Griffiths, B. N. Mathur, and F. S. Barrett. 2022. Models of psychedelic drug action: modulation of cortical-subcortical circuits. *Brain* 145(2): 441–456.

Flanagan, O., and G. Graham. 2017. Truth and sanity: positive illusions, spiritual delusions, and metaphysical hallucinations. In *Extraordinary Science and Psychiatry: responses to the Crisis in Mental Health Research*, eds. J. Poland, and S. Tekin. 293–313. Cambridge, MA: The MIT Press.
Flores, C. 2021. Epistemic styles. *Philosophical Topics* 49(2): 35–55.
Friedman, J. 2019. Inquiry and belief. *Noûs* 53(2): 296–315.
Garcia-Romeu, A., F. S. Barrett, T. M. Carbonaro, M. W. Johnson, and R. R. Griffiths. 2021. Optimal dosing for psilocybin pharmacotherapy: considering weight-adjusted and fixed dosing approaches. *Journal of Psychopharmacology* 35(4): 353–361.
Gładziejewski, P. 2021. Perceptual justification in the bayesian brain: a foundherentist account. *Synthese* 199: 11397–11421.
Gładziejewski, P. 2022. The rational role of the perceptual sense of reality. *Mind & Language*. https://doi.org/10.1111/mila.12445
Goldman, A. I. 2015. Naturalizing metaphysics with the help of cognitive science. *Oxford Studies in Metaphysics* 9: 171–212.
Gregory, D. I. 2020. Imagery and possibility. *Noûs* 54(4): 755–773.
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 (Berl)* 187(3): 268–283.
Griffiths, R. R., E. S. Hurwitz, A. K. Davis, M. W. Johnson, and R. Jesse. 2019. Survey of subjective 'God encounter experiences': comparisons among naturally occurring experiences and those occasioned by the classic psychedelics psilocybin, LSD, ayahuasca, or DMT. *Plos One* 14: 1–26.
Higgins, Ch. 2018. "There is no such thing as past or future": Physicist Carlo Rovelli on changing how we think about time. *The Guardian*. Available at: https://www.theguardian.com/books/2018/apr/14/carlo-rovelli-exploding-commonsense-notions-order-of-time-interview
Hohwy, J. 2020. New directions in predictive processing. *Mind & Language* 35(2): 209–223.
Hohwy, J., and J. Michael. 2017. Why should any body have a self? In *The subject's matter: self-consciousness and the body*, eds. F. de Vignemont, and A. Alsmith. 363–392. Cambridge (MA): The MIT Press.
Hohwy, J., B. Paton, and C. Palmer. 2016. Distrusting the present. *Phenomenology and the Cognitive Sciences* 15(3): 315–335.
Huxley, A. 2009/1954. *The Doors of Perception*. New York (NY): Harper Collins.
James, W. 2006/1896. *The Will to believe and other Essays in Popular Philosophy*. New York: Cosimo Classics.
James, W. 2008/1902. *The Varieties of Religious Experience*. London: Routledge.
Jones, R., and J. Gellman. 2022. Mysticism. In: E.N. Zalta (ed.). *The Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/entries/mysticism/
Katz, S. T. 1978. Language, epistemology, and mysticism. In *Mysticism and philosophical analysis*, ed. S. T. Katz. 22–74. New York: Oxford University Press.
Ko, K., G. Knight, J. J. Rucker, and A. J. Cleare. 2022. Psychedelics, mystical experience, and therapeutic efficacy: a systematic review. *Frontiers in Psychiatry* 13: 917199.
Korman, D. Z. 2019. Debunking arguments. *Philosophy Compass* 14: e12638.
Kung, P. 2010. Imagining as a guide to possibility. *Philosophy and Phenomenological Research* 81(3): 620–663.
Ladyman, D., and D. Ross. 2007. *Everything must go: Metaphysics Naturalized*. Oxford: Oxford University Press.
Le Poidevin, R. 2007. *The images of Time*. Oxford: Oxford University Press.
Letheby, C. 2021. *Philosophy of Psychedelics*. Oxford: Oxford University Press.
Letheby, C., and P. Gerrans. 2017. Self unbound: Ego dissolution in psychedelic experience. *Neuroscience of Consciousness* 2017(1): 1–11.
McGovern, H. T., P. Leptourgos, B. T. Hutchinson, and P. R. Corlett. 2022. Do psychedelics change beliefs? *Psychopharmacology (Berl)* 239(6): 1809–1821.
McGovern, H. T., H. J. Grimmer, M. Doss, B. Hutchinson, C. TimMermann, A. Lyon, P. R. Corlett, and R. Laukkonen. 2023. The power of insight: How psychedelics solicit false beliefs. PsyArXiv preprint: https://psyarxiv.com/97gjw/
Myers, J. 2021. The epistemic status of the imagination. *Philosophical Studies* 178(10): 3251–3270.
Nāgārjuna. 1995. *The fundamental wisdom of the Middle Way: Nāgārjuna's Mulamadhyamakakarika. Translation and commentary by Jay Garfield*. Oxford: Oxford University Press.

Nayak, S. M., M. Singh, D. B. Yaden, and R. R. Griffiths. 2023. Belief changes associated with psychedelic use. *Journal of Psychopharmacology* 37(1): 80–92.
Orłowski, P., A. Ruban, J. Szczypiński, J. Hobot, M. Bielecki, and M. Bola. 2022. Naturalistic use of psychedelics is related to emotional reactivity and self-consciousness: the mediating role of ego-dissolution and mystical experiences. *Journal of Psychopharmacology* 36(8): 987–1000.
Parfit, D. 1995. The unimportance of identity. In *Identity*, ed. H. Harris. 13–45. Oxford: Oxford University.
Parr, T., G. Pezzulo, and K. J. Friston. 2022. *Active inference: the Free Energy Principle in mind, brain, and Behavior*. Cambridge (MA): The MIT Press.
Paul, L. A. 2010. Temporal experience. *Journal of Philosophy* 107(7): 333–359.
Paul, L. A. 2012. Metaphysics as modelling: the handmaiden's tale. *Philosophical Studies* 160(1): 1–29.
Paul, L. A. 2014. *Transformative experience*. Oxford: Oxford University Press.
Paul, L. A. 2021. The paradox of empathy. *Episteme* 18(3): 347–366.
Petersen, R. 2019. Taking mushrooms for depression cured me of my atheism. *The Outline*. Available at: https://theoutline.com/post/7367/taking-mushrooms-for-depression-cured-me-of-my-atheism
Pettigrew, R. 2022. *Epistemic risk and the demands of rationality*. Oxford: Oxford University Press.
Plantinga, A. 2000. *Warranted christian belief*. New York: Oxford University Press.
Price, H. 1996. *Time's Arrow and Archimedes' point*. Oxford: Oxford University Press.
Ratcliffe, M. 2015. *Experiences of Depression: a study in Phenomenology*. Oxford: Oxford University Press.
Richards, W. A. 2018. *Sacred Knowledge: Psychedelics and Religious Experiences*. New York: Columbia University Press.
Schwitzgebel, E. 2001. In-between believing. *The Philosophical Quarterly* 51(202): 76–82.
Seager, W. ed. 2019. *The Routledge Handbook of Panpsychism*. New York: Routledge.
Sellars, W. 1963. *Science, Perception, and reality*. New York: Humanities Press.
Shani, I. 2015. Cosmopsychism: a holistic approach to the metaphysics of experience. *Philosophical Papers* 44(3): 389–437.
Siegel, S. 2017. *The rationality of Perception*. Oxford: Oxford University Press.
Simons, J. S., J. R. Garrison, and M. K. Johnson. 2017. Brain mechanisms of reality monitoring. *Trends in Cognitive Sciences* 21(6): 462–473.
Sjöstedt-Hughes, P. 2023. On the need for metaphysics in psychedelic therapy and research. *Frontiers in Psychology* 14: 112858.
Stace, W. T. 1960. *Mysticism and philosophy*. London: Macmillan.
Swanson, L. R. 2018. Unifying theories of psychedelic drug effects. *Frontiers in Pharmacology* 9: 172.
TimMermann, C., H. Kettner, C. Letheby, L. Roseman, F. E. Rosas, and R. L. Carhart-Harris. 2021. Psychedelics alter metaphysical beliefs. *Scientific Reports* 11(1): 22166.
Tucker, C. 2011. Phenomenal conservatism and evidentialism in religious epistemology. In *Evidence and religious belief*, eds. K. J. Clark, and R. J. VanArragon. 52–73. Oxford: Oxford University Press.
Whitteley, C. 2021. Depression as a disorder of consciousness. *British Journal for the Philosophy of Science*. https://doi.org/10.1086/716838
Williams, D. 2021. Imaginative constraints and generative models. *Australasian Journal of Philosophy* 99(1): 68–82.
Williamson, T. 2007. *The philosophy of philosophy*. Oxford: Blackwell Publishing.