PhiMiSci Philosophy and the Mind Sciences

Dissolving the self Active inference, psychedelics, and ego-dissolution

George Deane^a (georgejwdeane@gmail.com)

Abstract

Psychedelic drugs such as psilocybin, LSD and DMT are known to induce powerful alterations in phenomenology. Perhaps of most philosophical and scientific interest is their capacity to disrupt and even "dissolve" one of the most primary features of normal experience: that of being a self. Such "peak" or "mystical" experiences are of increasing interest for their potentially transformative therapeutic value. While empirical research is underway, a theoretical conception of the mechanisms underpinning these experiences remains elusive. In the following paper, psychedelic-induced ego-dissolution is accounted for within an active inference framework, as a collapse in the "temporal thickness" of an agent's deep temporal model, as a result of lowered precision on high-level priors. The argument here is composed of three moves: first, a view of the self-model is proposed as arising within a temporally deep generative model of an embodied organism navigating an affordance landscape in the service of allostasis. Next, a view of the action of psychedelics as lowering the precision of high-level priors within the generative model is unpacked in terms of a high Bayesian learning rate. Finally, the relaxation of high-level priors is argued to cause a "collapse" in the temporal thickness of the generative model, resulting in a collapse in the self-model and a loss of the ordinary sense of being a self. This account has implications for our understanding of ordinary self-consciousness and disruptions in self-consciousness present in psychosis, autism, depression, and dissociative disorders. The philosophical, theoretical and therapeutic implications of this account are touched upon.

Keywords

Active inference \cdot Consciousness \cdot Ego-dissolution \cdot Predictive processing \cdot Psychedelics \cdot Selfmodel

This article is part of a special issue on "Radical disruptions of self-consciousness", edited by Thomas Metzinger and Raphaël Millière.

1 Introduction

Psychedelic ("mind-manifesting") drugs are known to occasion radically altered states of consciousness, including profound changes in sensory perception, emo-

^aUniversity of Edinburgh

tion, cognition, time perception, and self-consciousness (Preller & Vollenweider, 2016). One of the most interesting of all of these effects is the experience of egodissolution. Although the experience is notoriously difficult to articulate and even considered ineffable, psychedelic researcher Stanislas Grof, who considers egodissolution the "main objective" of psychedelic therapy, describes it as "an ecstatic state, characterized by the loss of boundaries between the subject and the objective world, with ensuing feelings of unity with other people, nature, the entire Universe, and God" (Grof, 1980, p. 79).¹ Ego-dissolution is of considerable philosophical and theoretical value for understanding selfhood and the nature of consciousness (Letheby & Gerrans, 2017; Millière, 2017; Nour & Carhart-Harris, 2017).² It is also considered to be central to the therapeutic potential of psychedelics (Carhart-Harris & Goodwin, 2017; see also Letheby, 2020, Limanowski & Friston, 2020, Sebastián, 2020, all in this special issue). Despite this, very little is known about the mechanisms underpinning psychedelic-induced ego-dissolution.

"Predictive processing" theories of brain function (Clark, 2013; Friston, 2010; Wiese & Metzinger, 2017) have recently taken precedence in cognitive science, affording a novel theoretical framework to approach cognitive phenomena. In this paper I propose a novel account of ego-dissolution within an active inference framework. To this end, I initially furnish an account of self-modelling within active inference, where pre-reflective self-consciousness emerges in organisms as a consequence of "temporal thickness", the need to model the consequences of potential actions over time (Friston, 2018). I then give an account of the action of psychedelics within a predictive processing framework, unpacking the view that psychedelics "relax" high-level priors (Carhart-Harris, 2019; Carhart-Harris & Friston, 2019) in terms of a high Bayesian learning rate (Hohwy, 2017; Mathys et al., 2014). Finally, I argue that low precision at high-levels of the inferential hierarchy results in a collapse of the temporal thickness of the generative model and the corresponding self-model, leading to the phenomenon known as ego-dissolution (see also Limanowski & Friston, 2020, this issue).

2 The free energy principle

The Free Energy Principle (FEP) has the most ambitious explanatory scope of all "predictive processing" style theoretical frameworks (Friston, 2010). It combines, subsumes and links to several brain theories, including the Bayesian brain hypothesis (Knill & Pouget, 2004), predictive coding (Mumford, 1992; Rao & Ballard, 1999), efficient coding (Barlow, 2001) and reinforcement learning (Dayan & Daw, 2008). The mathematics of the theory are complex and beyond the scope of this paper (for a review see Buckley, Kim, McGregor, & Seth, 2017). According to the FEP,

¹See https://www.erowid.org for over 20,000 "trip reports" of various psychoactive drugs, including many descriptions of ego-dissolution.

²No particular theory of consciousness is assumed in this paper, but the proposal is considered consistent with predictive processing approaches to consciousness (Clark, 2019; Hohwy, 2012).

Deane, G. (2020). Dissolving the self: Active inference, psychedelics, and ego-dissolution. *Philosophy and the Mind Sciences*, *1*(I), 2. https://doi.org/10.33735/phimisci.2020.I.38

simply in virtue of existing, all organisms tend to minimise the entropy or dispersion of their states. This much is intuitive: the conditions that are viable for an organism are fairly narrow – deviation from homeostatic bounds, such as having a body temperature of 50 degrees centigrade, is incompatible with continued existence. Organisms that fail to stay within their "species-specific window of viability" (Clark, 2013 p. 13) simply cease to exist. Life, on this account, resists the tendency towards disorder imposed by the second law of thermodynamics, and this principle applies at all levels – "from their gross morphology to fine details of cortical microcircuitry as well as at timescales from the neuronal to the phylogenetic" (Seth & Tsakiris, 2018, p. 973). Organisms, then, must resist entropy, the long-term average of (information-theoretic) surprise. Because this quantity is beyond direct epistemic access to an organism, according to the FEP, organisms minimise a proxy variable or upper bound – dubbed (variational) free energy. Free energy (under some simplifying assumptions) is equivalent to precision-weighted prediction error in predictive processing.

2.1 Predictive processing

On the predictive processing view, the brain has stored prior beliefs (in the form of probability distributions) about the causes of sensory inputs in the world (Clark, 2013; Wiese & Metzinger, 2017). Prior beliefs are hierarchically organised, where higher-levels encode predictions about representations at lower levels. Prediction errors, arising from the discrepancy between the low-level predictions and incoming sensory signals, are passed up the hierarchy, where higher-level predictions are updated to minimise further prediction errors. Perception, then, both exteroceptive and interoceptive, is the product of (approximate) Bayesian inference, whereby the influence of prior beliefs and sensory evidence are weighted according to "expected precision", e.g. confidence in the given context, to generate a posterior. Inference in these schemes is thought to occur across a hierarchy of inferred causes, where higher levels encode regularities that occur at larger spatial and temporal scales (Kiebel, Daunizeau, & Friston, 2008). In perceptual inference, sensory prediction errors can be minimised by tweaking the parameters of the generative model - that is, generating predictions to quash the influx of prediction error. Prediction error can also be minimised though action by changing the incoming sensory data to fit a prediction - for instance, I can move my eyes to bring my coffee cup into view, to fulfil the prediction of a coffee cup. Actions can be thought of as the fulfilment of proprioceptive (or oculomotor) predictions - an intended movement occurs as a result of predicting the proprioceptive consequences (Friston et al., 2010; Shadmehr, Smith, & Krakauer, 2010). There are detailed accounts of the neural implementation of these schemes available (Bastos et al., 2012; Keller & Mrsic-Flogel, 2018; Shipp, 2016).

2.2 Precision-weighting

A key feature of predictive processing schemes is the contextual flexibility afforded by precision-weighting (Clark, 2013; Feldman & Friston, 2010). Precision regulates the interaction between top-down and bottom-up signals, through the synaptic gain on neuronal populations signalling prediction error, in order to approximate optimal inference over time. Precision can be thought of as tracking both the reliability and relevance of the incoming sensory information, where weighting by reliability is analogous to assigning greater weight to more reliable information when updating a belief. Prediction error signals with high precision (inverse variance) have greater influence in updating the top-down predictions. Precision itself has to be inferred, both by the empirical variance in the sensory data itself, and according to prior expectations about precision. The optimisation of precision weighting, through updating of the precision expectations (precision-related priors), is frequently equated to attention within predictive processing (Clark, 2013; Feldman & Friston, 2010). Importantly for the current treatment, precision is thought to mediate both sensory attenuation - the top-down filtering out of afferent information, and affordances, where affordances refers to the latent possibilities for action given the capabilities of the agent (Cisek, 2007; Friston, Shiner, et al., 2012).

2.3 Control-oriented inference

In mandating that existence necessitates maintaining oneself within a limited repertoire of states via *control-oriented predictive regulation* (instrumental active inference) (Seth & Tsakiris, 2018), the FEP aligns itself with precursors of this view, cybernetic theories that build on control, feedback and predictive modelling (e.g., the "good regulator theorem") (Conant & Ashby, 1970). Note that while a purely Helmholtzian view of the brain might cast it in terms of inferring hidden causes in the world, casting the predictive machinery in terms of being *for* ensuring continued existence means that the generative model is not constrained to veridicality. Rather than faithfully reconstructing the world, perception is "ultimately geared towards driving actions that preserve [the] physiological integrity of the organism. In other words, we do not perceive the world (and self) as it is, but as it is useful to do so" (Seth & Tsakiris, 2018, p. 975).

2.4 Homeostasis

Homeostasis refers to the tendency of living systems to keep an "internal balance" despite changes in the surrounding environment (Cannon, 1929). This has long been described in terms of control theoretic and cybernetic mechanisms, and more recently this homeostatic control is thought to involve interoceptive signals that report current physiological states (e.g., heart rate, or blood-bound glucose levels) (Craig, 2002; Damasio & Carvalho, 2013). One way to restore bodily conditions to favourable states is to engage autonomic reflexes – for example, a hyperthermic

animal can perspire to cool down. Of course, autonomic regulation alone is not sufficient to ensure continued existence – to avoid hunger or thirst the animal must engage actions, such as seeking out food and water. Collectively, these actions are termed *allostasis*, the process via which the brain regulates the needs of the body (Corcoran & Hohwy, 2018; Corcoran, Pezzulo, & Hohwy, 2019; Schulkin & Sterling, 2019). Crucially, to stay viable on longer timescales, this action must be *anticipatory* – avoiding dyshomeostatic conditions before they arise (Pezzulo, Rigoli, & Friston, 2015; Sterling, 2012).

2.5 Active inference

The FEP regards homeostasis and allostasis as the central aspects of organic life, thus the autopoietic principles at the basis of the FEP act as a kind of "first prior" (Allen & Tsakiris, 2018). In other words, "[t]he brain is in the game of predicting the world, but only as a means to the end of embodied self-preservation" (Allen & Friston, 2018, p. 12). In so doing, the free energy principle collapses expected utility (instrumental value) and information gain (epistemic value) under a single quantity. On this approach, action planning is itself a form of inference, where preferences and goals are framed in terms of prior beliefs, such that these priors are fulfilled by action (Botvinick & Toussaint, 2012). Casting value and utility purely as inferential problems may at first appear unintuitive - if an agent finds itself in consistently adverse circumstances, then such adverse circumstances should, at first pass, seem to have high probability. However, "[t]he critical step in this logic is the assumption that evolution has equipped us with the belief that low utility states are low probability, due to the fact that if our ancestors spent a lot of time in those states they would be less likely to reproduce" (Gershman, 2019, p. 7). The so-called "first prior", that of maintaining existence via homeostatic and allostatic regulatory behaviour, ensures that organisms seek to actively maintain internal and external conditions conducive to their own persistence.

Active inference refers to the process by which agents actively sample states of the world so as to reduce uncertainty and realise prior preferences, rendering the action selection process itself an inference problem. This arbitration occurs according to priors pertaining to *expected* free energy over a given course of action, or *policy* (Friston et al., 2018; Pezzulo et al., 2015). Expected free energy is the free energy an agent predicts itself to average in opting to pursue a particular course of action. Intuitively, some courses of action are more likely than others to lead to "expected" or desirable outcomes. A policy that has lower expected free energy is going to have a higher prior probability than a policy with higher expected free energy, because agents equipped with prior beliefs about their continued existence will pursue policies that reduce expected free energy (Friston et al., 2015, 2018; Kaplan & Friston, 2018). Crucially, agents engaging in active inference do not merely restrict themselves to the states they expect; rather they anticipate in order to minimise uncertainty about potential future outcomes (Friston et al., 2015,

2018; Pezzulo et al., 2016). This prospective form of control relies on the contextualization provided by higher levels in the inferential hierarchy, which anticipate the downstream consequences of actions and select policies accordingly (Friston, 2010; Pezzulo et al., 2015). Contextualisation here depends on the relative precision at various hierarchical levels, where "precision dynamics subsume the role of arbitration" (Pezzulo et al., 2015, p. 27). This approach bears similarities to other control-theoretic approaches, such as the *affordance competition hypothesis* (Cisek, 2007; Pezzulo & Cisek, 2016), where an affordance is a potential for action that avails itself to an organism in its action-oriented perception of environmental features.³ On this view, perceived affordances jostle for precedence and are arbitrated on the basis of the desirability of their predicted outcomes.

3 The self in active inference

This section outlines an account of how *pre-reflective self-consciousness* – an implicit sense of being a subject present in all experience – is structured within an active inference framework. Here, the self-model is underpinned by the same inferential Bayesian schemes that are increasingly being used to describe perception and action. This predictive-modelling approach to selfhood has roots in Thomas Metzinger's work on conscious and unconscious self-models, and the "self-model theory of subjectivity" (Blanke & Metzinger, 2009; Metzinger, 2003, 2009) where "[a] self-model, an inner image of the organism as a whole [is] built into the world-model, and this is how the consciously experienced first-person perspective develop[s]" (Metzinger, 2009, p. 64).

The account presented here follows the increasing focus on the embodied nature of selfhood, where "being" or "having" a body is thought to be one of the most basic aspects of the experience of being a self (Allen & Friston, 2018; Apps & Tsakiris, 2014; Blanke & Metzinger, 2009; Limanowski & Blankenburg, 2013). A growing number of researchers seek to ground selfhood and emotion in interoceptive processes, particularly in their functional relation to allostatic regulation (Barrett & Simmons, 2015; Seth, 2015; Seth & Friston, 2016). A key reason for this is that interoceptive inference is apt to put greater emphasis on *control* over *discovery* (Seth & Friston, 2016), due to "a priori hyper-precision of visceral channels" (Allen & Friston, 2018, p. 7), in which interoceptive signals are assigned very high precision in virtue of communicating information about key physiological variables (Seth, 2015). Grounding the self-model in *control-oriented active inference* (Seth & Tsakiris, 2018) inflects perception of the affordance landscape in terms of bodily states, an idea which is nicely expressed by Montague and King-Casas:

³It is worth noting here that the proposal put forward in this paper is considered to be compatible with both internalist and externalist conceptions of the mind. Thanks to an anonymous reviewer for suggesting this clarification.

Deane, G. (2020). Dissolving the self: Active inference, psychedelics, and ego-dissolution. *Philosophy and the Mind Sciences*, *1*(I), 2. https://doi.org/10.33735/phimisci.2020.I.38

A sated and comfortable lioness looking at two antelopes sees two unthreatening creatures against the normal backdrop of the temperate savanna. The same lioness, when hungry, sees only one thing – the most immediate prey. In another circumstance, in which the lioness may be inordinately hot, the distant, shaded tree becomes the prominent visual object in the field of view.

(Montague & King-Casas, 2007, p. 519)

This forms the basis for the view that will be unpacked in more detail in what follows, that the self-model can be understood as an "allostatic control model", arising from the system's sense of control of the temporally deep consequences of actions for allostasis. On this view, pre-reflective self-consciousness is underpinned by the inference about endogenous control of the sensory consequences of actions within deep goal hierarchies, where goals and preferences are framed in terms of prior beliefs, such that goals are fulfilled by actions (Botvinick & Toussaint, 2012; Pezzulo et al., 2015). Recall, action allows an organism to change the sensory input in order to conform to its generative model, as opposed to perceptual inference that involves revising model parameters to conform to the sensory input. In order to act, then, the system implicitly infers its own ability to bring about the intended sensory consequences - it is in this sense that "implicit in a model of sampling is a representation or sense of agency" (Friston, 2012a, p. 173), which is closely related to what has been called the "primacy of the 'I can'" (Bruineberg, 2017). Crucially, organisms with deep temporal models have "temporal thickness" - expectations regarding the sensory consequences of actions on multiple interlocking timescales. The following sections will unpack this conception of the self-model in terms of hierarchically deep allostatic control, starting with the notion of temporal thickness, and then moving to how motivated control hierarchies "attune" organisms to action opportunities on multiple timescales, for both proximal goals, for instance, pain motivating an organism to act so as to fulfil a "healthy body condition" prior; and distal goals, for example emotions motivating a change of circumstances pertaining to longer timescales such as moving to a different city. The discussion will then move to how deep self-models allow organisms to arbitrate between different policies and trade off outcomes on different timescales.

3.1 Temporal thickness

To successfully navigate the world over longer timescales, and select policies that result in survival – and not dispersion or non-existence – organisms must possess models of the future; in other words, they require *deep temporal models* (Friston et al., 2018). The generative models that endow organisms with the capability of inferring the consequences of future actions must have the property of temporal thickness, which allows the organism to anticipate the downstream consequences of potential actions, conferring the ability to select policies or action scripts that are favourable to the organism's continued existence (Friston, 2018). The minimi-

sation of surprise through active inference on the FEP involves acting so as to reduce uncertainty, and to do this the system must model itself across time and counterfactuals – that is, it must model what kind of agent it is at varying degrees of temporal depth. Self-modelling, then, emerges as a natural consequence of prospective action selection (Friston, 2018), where the principal function of a counterfactually rich self-model is to facilitate navigation of the affordance land-scape and action selection across multiple interlocking timescales – for example expectations of what an agent can do on shorter timescales inform expectations of what the agent can do over longer timescales. The functional role of having a rich self-model, then, is that it enables the organism to predict outcomes across diverse policies, and endows the organism with "what if?" capabilities (Friston, 2018), which puts this picture into contact with mental time travel and offline simulation (Buckner & Carroll, 2007; Schacter, Addis, & Buckner, 2008).

3.2 Attuning to the world

Conceiving of the self-model through an active inference framework, a hierarchically deep self-model guides policy selection over various timescales in service of minimising expected free energy. In what follows, pain perception, viewed as arising through the violation of the prior of "healthy body condition" (Ongaro & Kaptchuk, 2019), will be used to illustrate how inferences about the self "attune" an organism to adaptive action opportunities.

One key advantage that the active inference account of self-modelling has over strictly Bayesian approaches is that it is goal-directed (Moutoussis et al., 2014). Classical models of pain perception as the consequence of physiological dysfunction are challenged by the efficacy of placebo treatments in relieving pain (Anchisi & Zanon, 2015), and cases in which pain is experienced without physiological disruption, as is often the case in chronic pain. Instead, there is evidence to suggest that affectively charged percepts, such as pain, are best understood as resting on the same inferential mechanisms as are assumed to underpin perception and action under a predictive processing framework (Büchel, Geuter, Sprenger, & Eippert, 2014). Bayesian models of pain perception (Morton, El-Deredy, Watson, & Jones, 2010) indicate that prior beliefs about the generation of painful percepts are integrated with current sensory data to infer the posterior or hidden worldly cause (the painful percept). Crucially, these pain percepts incorporate the "weight" or precision of past experiences when computing the current painful percept (Morton et al., 2010). On their own, however, these models of pain perception are silent on the functional role of pain as a motivator to an embodied organism (Moutoussis et al., 2014). Optimal inference about pain to the allostatically concerned organism is heavily dependent on the context, as anyone who has felt the pain of an injury only after danger is averted can attest to. In this way, pain perception is allostatically "tuned": "organisms can tune their own pain perception according to both their prior beliefs and the specific biological goals they believe are attainable in that context" (Moutoussis et al., 2014, p. 70).

A Bayesian framework of pain perception, therefore, needs to represent the agency and aims of the organism. This is precisely what is afforded by conceiving of the self-model within an active inference framework (Friston, 2012b) – as this provides the necessary context to study the self-model, across multiple hierarchical levels. Like physical pain, and sharing the neural underpinnings of physical pain (Eisenberger, 2012), social pain is similarly understood in inferential terms, and does not scale with "damage" per se (for instance, social rejection), as evidenced by the wide range of sensitivity people have to the same physical manipulation (Eisenberger, Jarcho, Lieberman, & Naliboff, 2006). Accordingly, there is evidence to suggest that appropriately "tuning" emotional responses in social contexts allows for agents to approximate Bayesian inference in policy selection given bounded cognitive capacity and rationality. For example, on a "stag hunt" game,⁴ agents with "prosocial" preferences can outperform agents of similar cognitive sophistication that lack social biases (Yoshida, Dolan, & Friston, 2008).

3.3 Emotion

Conceptualising emotions in terms of a contextualisation of bodily states has historical roots dating back to the James-Lange theory of emotion (Cannon, 1927) and two-factor theory of emotion (Schachter & Singer, 1963). Lisa Feldman-Barrett has developed this approach specifically within the active inference framework as the "theory of constructed emotion" (Barrett, 2017). According to the theory of constructed emotion, emotions are constructed in the same manner as percepts, where priors are recruited according to context to make a "best guess" at the hidden causes of (interoceptive) sensory signals. On Barrett's view, emotions arise through a context-sensitive inferential categorisation of interoceptive states. For this reason, emotions on this view are "constructions" - there are no neural or physiological signatures that reliably discriminate any emotional state (Siegel et al., 2018; Wager et al., 2015). Rather, physiological reactions in the body occur in order to prepare it for action, and these are categorised as emotions only contextually through the predictive models recruited to explain away the incoming afferent interoceptive signals. For example, heart rate increases or decreases depending only on an anticipated action - e.g., fight or flight - and given an emotional ascription only contextually - e.g., the same bodily state could be categorised as fear in one context and anger in another. Interoceptive inference is experienced as emotion in service of producing allostatic action (Barrett, Quigley, & Hamilton, 2016; Barrett & Simmons, 2015; Wilkinson, Deane, Nave, & Clark, 2019). In viewing the self-model in terms of hierarchical allostatic control, interoceptive inference on the hidden causes of bodily states pertaining to longer timescales tunes perception to the world and affordances differently, such that more abstract emo-

⁴A stag hunt game is a cooperative hunting game, where in order to hunt a stag (higher-payoff), an agent must have the cooperation of a partner. An agent can catch a rabbit (lower-payoff) without cooperating.

Deane, G. (2020). Dissolving the self: Active inference, psychedelics, and ego-dissolution. *Philosophy and the Mind Sciences*, *1*(I), 2. https://doi.org/10.33735/phimisci.2020.I.38

tions might track regularities over longer time scales, informing policy selection thereon (Pezzulo, 2014), and allowing for more abstract and distal outcomes to be motivationally salient.

3.4 Hierarchically deep self-models

Viewing the self-model in terms of allostatic control renders selfhood fundamentally affective and action-oriented, such that different aspects of the self in a given context - precision on goals and preferences at different levels of the hierarchy motivate behaviour and arbitrate between policies. On this view, the self-model inflects perception of possible actions in the world and mediates salience to facilitate the selection of policies with minimal expected free energy. Critical to this picture is the notion that these various models are associated with varying degrees of temporal depth (Pezzulo, Rigoli, & Friston, 2018). Deep generative models capture increasingly distal relations between actions and outcomes within hierarchical active inference, allowing for the coordination of behaviour across different hierarchical levels, enabling goals to become prioritised relative to current context (Pezzulo et al., 2018). The result is an inferential framework of hierarchically nested contextual complexity, in which lower levels track basic (and sometimes evolutionarily hard-wired) motivations or affordances, while higher levels track motivations and plans over deeper timescales. In this way, higher-level contextualization of lower sensorimotor functions optimises expected actions in terms of both long-range consequences of actions and anticipated future affordances. Goals at different levels of abstraction may, of course, conflict - for instance, resolving proximal interoceptive prediction error by eating chocolate cake might conflict with the longer-term goal of sticking to a diet (Pezzulo et al., 2018). Alternatively, temporary deviation from homeostatic set points at lower levels may be elicited to maintain higher level set points – such as a temporary change in blood pressure and adrenaline levels to engage fight-flight behaviour, with the goal of reaching safety and maintaining physiological integrity on a longer timescale. On the view of self-modelling in terms of allostatic control described, dimensions of the self at higher-levels constrain the self at lower-levels in that the self-model "actively shapes itself over time to align with those higher level regularities" (Hohwy & Michael, 2017, p. 370), for example long-term goals can be decomposed into intermediate short term-goals.

This section has explored how the self-model arises as a consequence of a system engaged in temporally deep active inference, as prior probabilities over particular policies depend on knowledge about what and where the system finds itself, and what actions are available to it (Friston et al., 2013; Moutoussis et al., 2014). Through active inference, agents can use their self-model to inform their goal and policy selection in order to arrive at high probability outcomes. This could entail assigning low probability to the self occupying states that are aversive, either physically or socially – with different hierarchical levels of the self-model contributing to different goal states. In this way, the hierarchical self-model determines salience – where "salience is literally defined by whatever has the most (or least) impact on visceral and autonomic homeostasis" (Allen & Tsakiris, 2018, p. 7), at increasingly deep spatiotemporal scales and levels of abstraction.

4 Psychedelics

One of the most striking and philosophically interesting effects of psychedelics is the radical disruptions of self-consciousness they can occasion (Huxley, 2010; Leary, Metzner, & Alpert, 1964), including apparently "selfless states" (Lebedev et al., 2015; Nour, Evans, Nutt, & Carhart-Harris, 2016). These states, instances of "Drug-Induced Ego-Dissolution" (DIED) are characterised by an experienced loss of self and/or loss of self/world boundary (Millière, 2017; Millière, Carhart-Harris, Roseman, Trautwein, & Berkovich-Ohana, 2018). DIED occurs most reliably under high doses of "classical" psychedelic drugs (5-HT_{2A} receptor agonists), such as dimethyltryptamine (DMT), lysergic acid diethylamide (LSD), and psilocybin. Ego-dissolution appears to be induced more reliably under psychedelics than meditation, in a dose-dependent manner, and prompted most reliably by high-doses (Nour et al., 2016). Recent theoretical work has explored the phenomenological and neurophysiological similarities and differences of ego-dissolution induced by drugs and meditation (Millière et al., 2018; see also Limanowski & Friston, 2020, Millière, 2020, Sebastián, 2020, all in this special issue).

4.1 **Psychedelic therapy**

Recent years have seen a resurgence of interest in the therapeutic potential of psychedelics. Several studies have found preliminary evidence that with administration in controlled circumstances psychedelics can be both safe and therapeutic, with an emphasis on the importance of context in achieving therapeutic outcomes (Carhart-Harris et al., 2018). Interestingly, the positive therapeutic effects seem to scale with "peak" or mystical experience in the psychedelic state (Roseman, Nutt, & Carhart-Harris, 2018). Psychedelics have been shown to be effective in treating depression (Lyons & Carhart-Harris, 2018; Palhano-Fontes et al., 2019), obsessivecompulsive disorder (Moreno, Wiegand, Taitano, & Delgado, 2006), end of life existential distress (Griffiths et al., 2016), and have even been proposed as a potential treatment for disorders of consciousness such as the vegetative state and the minimally conscious state (Scott & Carhart-Harris, 2019). Carhart-Harris interprets the therapeutic effects as a result of "relaxing" high-level beliefs, allowing for a revision of pathological beliefs that have become overly dominant and resistant to revision, coined the "TIghtened BEliefs in Response to uncertainty" (TIBER) model (Carhart-Harris, 2019). The basic tenet is that under conditions of uncertainty the model falls back on "tightened" belief structures as a defence mechanism against intolerable stress and uncertainty. This fits with what we might expect under the FEP, as adopting shallow policies (such as addictive behaviours) may appear adaptive in the short term, rather than risking policies with greater expected free energy due to low precision or uncertainty. It has recently been proposed that psychedelics "relax" high-level priors in the generative model, allowing for the (context-dependent) revision of pathological high-level beliefs (Carhart-Harris & Friston, 2019). Both psychological insight and peak-experience in the psychedelic state appear to be predictors of long-term positive prognoses (Roseman et al., 2018).

4.2 Psychedelics in the predictive brain

The REBUS – "RElaxed Beliefs Under pSychedelics" – model of psychedelic function, offers a preliminary but promising model of psychedelic action where psychedelics, through 5-HT_{2A} agonism, "relax" high-level priors or beliefs (Carhart-Harris & Friston, 2019). Here, the focus will be on how this mechanism may be cast under the hierarchical predictive processing framework as modulating precision-weighting. To bring this into focus, this section will review how precision-weighting sets a variable *Bayesian learning rate* in order to highlight certain features relevant to understanding the effects of psychedelics within this framework. Christoph Mathys and colleagues have recently developed a mathematical tool for modelling Bayesian inference modulated by expectations of volatility known as the hierarchical Gaussian filter (Mathys, Daunizeau, Friston, & Stephan, 2011; Mathys et al., 2014). The hierarchical Gaussian filter Mathys posits allows a system to optimally balance the influence of prediction errors in changing environments – in other words, to adjust its learning rate.

4.3 Bayesian learning rate

To recap, the predictive processing framework asserts that the brain instantiates "generative models" of the causes of incoming sensory data, iteratively updating these predictive models in light of incoming "prediction error" (Clark, 2013). This predictive inference is thought to occur across a hierarchy of inferred causes, where high levels track causes and regularities operating over deeper spatial and temporal scales, and lower levels track regularities over shallower spatial and temporal scales (Kiebel et al., 2008). The picture of the living or cognitive system as one which needs to optimise its own learning rate emerges out of the operationalization of Bayesian inference in predictive processing, namely in terms of predictions and precision-weighted prediction errors (Mathys et al., 2014). According to predictive processing the prediction is given by the prior probability (which itself comes from the previous posterior) and the prediction error is given by the difference between the prediction and the incoming sensory evidence. Prediction error is weighted according to the relative precisions of the prior and the prediction error (where precision is equivalent to the inverse variance of each probability distribution). Intuitively, highly precise prediction error will drag

the posterior closer to the distribution of the sensory evidence and further from the prior, and in cases of low precision weighting of the prediction error, the inference relies more on the prior. This determines the learning rate:

The more certain we are that the prior hypothesis is correct, the less we should be influenced by the prediction error (the evidence), which means that the learning rate is low. Conversely, the better the precision on the prediction error, the higher the learning rate; that is, the more we trust the quality of the evidence the more we should learn from it (Hohwy, 2017, p. 76)

In other words, the lower the learning rate, the greater the influence of top-down modulation from priors; the higher the learning rate, the greater the influence of the sensory evidence on the resulting posterior. Here, precision-weighting is the key mechanism - heavily weighted prediction errors drive a higher learning rate. In order to approximate Bayesian inference over time, it is essential for sensory systems to balance the learning rate appropriately. Over-reliance on priors means the system will fail to learn from experience, whereas over-reliance on sensory evidence (which may be noisy) will lead the system to "overfit". On this picture, Bayesian perceptual inference that minimises prediction error on the appropriate timescale - that is, not overfitting or underfitting the model - needs to have a means of regulating the learning rate (Mathys et al., 2014). This is implemented by building models of precision, or expected uncertainty, where higher-level priors track longer-term regularities that inform the relative precisions of more basic priors (Hohwy, 2017). Optimising the learning rate, and in-so-doing minimising prediction error over time, is a critical challenge the brain faces. This is equivalent to selecting a time frame over which to minimise prediction error. Minimising prediction error over too short a timescale - overfitting - runs the risk of increasing prediction error in the long run. Conversely, failing to accommodate new evidence will lead to underfitting, a failure to update predictions in light of new sensory evidence.

4.4 Psychedelic action as high Bayesian learning rate

In line with the REBUS model (Carhart-Harris & Friston, 2019), the relaxing of high-level priors under classical (serotonergic) psychedelics⁵ means the system adopts a very high Bayesian learning rate – that is, it is in a highly plastic state, in accordance with research showing an increase in plasticity under psychedelics (Ly et al., 2018). This picture casts the perceptual effects of psychedelics – "tripping" – as rampant overfitting of the sensory data, resulting from a loss of the usual con-

⁵This account is not thought to generalise to dissociative anaesthetics, like ketamine or phencyclidine (PCP), which are thought to have quite different mechanisms of action, acting primarily at NMDA glutamate receptors. Comparative approaches of both the mechanisms of action of other drugs as compared to psychedelics, and the mechanisms underpinning different experiences of ego-dissolution, would be a fruitful avenue for future research.

straint exerted by higher-levels on lower-levels of the inferential hierarchy. This "rampant overfitting", resulting from diminished influence from contextualising high-level priors tracking regularities on longer timescales, means the model fits a very short temporal scale, rapidly cycling through candidate models to account for the incoming sensory signal. It is worth highlighting a compatibility of the high Bayesian learning rate approach with other accounts of the mechanism of action of psychedelics in the predictive brain. The REBUS model posits the mechanism of action of psychedelics as reduced precision at high levels rather than increased precision at the sensory peripheries, as psychedelics appear to disrupt functioning via stimulation of 5-HT_{2A} receptors on deep pyramidal neurons, thought to encode high level priors or beliefs (Beliveau et al., 2017; Carhart-Harris & Friston, 2019; Jakab & Goldman-Rakic, 1998). In contrast, Philip Corlett and colleagues have suggested that psychedelics *preserve* normal priors and act by increasing sensory noise through enhanced AMPA signalling (Corlett, Frith, & Fletcher, 2009; Corlett, Honey, & Fletcher, 2016). On this approach, if the relaxation of high-level priors is indeed an effect of psychedelics, it could be understood to be the *result* of the fact that "the persistence and strength of the sensory signal suggest that there is something to be explained" (Corlett et al., 2009, p. 521). Arbitrating between these two mechanistic accounts and disentangling causation – whether the relaxation of high-level priors causes the reduction in sensory gating, or reduction in sensory gating eventually lowers precision at high levels – becomes very difficult here, and it is not clear a simplistic causal account is the right approach. While identifying the mechanisms of action is a key empirical and theoretical project, one potential advantage of the high Bayesian learning rate hypothesis is that it doesn't distinguish between high precision at low levels and low precision at high levels, and as such remains agnostic over the mechanism of action while preserving the useful theoretical features of both accounts that will inform the theoretical account of ego-dissolution that follows.

4.5 Evidence for the high Bayesian learning rate hypothesis

A high Bayesian learning rate is concordant with the enhanced neural plasticity observed in individuals in a psychedelic state (Barre et al., 2016; Berthoux, Barre, Bockaert, Marin, & Bécamel, 2018; Ly et al., 2018). While an impairment to high-level cognition is found under psychedelics (Bayne & Carter, 2018), in line with the high Bayesian learning rate hypothesis, low-level learning (including extinction learning) and processing appears to be unaffected or enhanced in the psychedelic state (Carhart-Harris & Nutt, 2017; King, Martin, & Seymour, 1972; Romano et al., 2010). Further evidence for a high Bayesian learning rate under psychedelics is provided by a study looking at the effect of psilocybin on Kanisza triangles – perceptual objects where the brain "fills in" illusory contours using prior expectations – which found reduced filling in and a reduction in the related evoked potentials (Kometer, Cahn, Andel, Carter, & Vollenweider, 2011), concordant with

the fact that a high Bayesian learning rate will reduce the effect of sensory history on current perception. In binocular rivalry studies – where different images are presented to each eye simultaneously, and are typically experienced as switching from one percept to the other – reduced switch rates and increased likelihood of the percept being a fusion of the two images has been observed under psilocybin (Carter et al., 2007, 2005), suggestive of less influence of priors on constraining current perception. Oddball paradigms are also suggestive of a weakened influence of priors on perception under psychedelics. In a sequence of tones, an "oddball" tone (unexpected given prior experience and context) generates a "mismatch negativity", an evoked brain response which has been interpreted in predictive coding terms as prediction error violating the expectations of the sequence (Garrido, Kilner, Stephan, & Friston, 2009). Under LSD, the surprise response to oddball stimuli is blunted, suggestive of a weakened influence of prior expectations (Timmermann et al., 2018).

Arguably, there is also phenomenological evidence for the high Bayesian learning rate hypothesis. Perhaps most eloquently articulated by Aldous Huxley: "Visual impressions are greatly intensified and the eye recovers some of the perceptual innocence of childhood, when the sensum was not immediately and automatically subordinated to the concept" (Huxley, 2010, p. 12).

This observation lends itself to a straightforward translation into the terms of predictive processing, where "subordinated to the concept" can be understood as "constrained by higher-level priors". More generally, psychedelic phenomenology such as dynamic distortions of spatial dimensions, where things change dramatically in size and shape can be understood as a failure of high-level priors to canalise and constrain lower level predictions.

5 Psychedelic-induced ego-dissolution in active inference

Given this picture of the action of psychedelics within a predictive processing framework, and the characterisation of self-models in terms of allostatic control, how should states of psychedelic-induced ego dissolution be conceptualised? The proposal here is that a loss of precision on high-level priors results in a flattening of temporal depth of the affordance landscape for the organism – precisely because it is high-level priors tracking longer timescales that structure temporally deep generative models. Recall, under active inference, lower and higher hierarchical levels encode regularities that unfold at faster and slower timescales respectively (Kiebel et al., 2008), such as the expected consequences of action both for proximal and distal goals (Pezzulo et al., 2015, 2018). Adopting a high Bayesian learning rate is equivalent to changing the time frame over which prediction error is minimised to fit very short timescales. As a result, the deep temporal models that typically guide action and policy selection collapse, and the faster timescales correspond-

ing to lower levels are modelled in a much finer degree of detail (Pink-Hashkes, Rooij, & Kwisthout, 2017). On the account presented in this paper, the self-model is constructed and bolstered in relation to affordances in the environment on several interlocked timescales, where high-levels contextualise and canalise the levels below and allow for motivational orientation to action opportunities pertaining to distal outcomes. Under psychedelics, the relaxation of high-level priors and the corresponding high Bayesian learning rate results in a collapse in the temporal thickness of deep generative models, and a collapse in the temporal depth of the corresponding self-model, which is understood as being is bolstered according to counterfactually rich expectations of the consequences of action on multiple timescales.

The collapse in temporal thickness can be understood as occurring due to a failure of sensory attenuation, occurring due to low precision at high-levels and a correspondingly high Bayesian learning rate. Similar stories about aberrant precision at high-levels of the hierarchy corresponding to inferences about affordances and agency have been proposed to underpin hallucinations and delusions in psychosis (Adams, Stephan, Brown, Frith, & Friston, 2013; Fletcher & Frith, 2009; Sterzer et al., 2018). Distinguishing between endogenous and exogenous causes - that is, distinguishing between perceptual inputs caused by oneself and those caused by the world – is vital for an agent to be able to effectively move through action space. Corollary discharges - predictions about the sensory consequences of actions - allow the system to do this by withdrawing precision from self-generated movements, and are thought to underpin experienced agency of actions (Crapse & Sommer, 2008; Friston, 2012b). The failure to predict the consequences of movement due to a failure of sensory attenuation is thought to result in an inability to attribute agency (Adams et al., 2013; Brown, Adams, Parees, Edwards, & Friston, 2013); for instance, a failure of corollary discharge has been thought to cause the attribution of inner speech to an external source in voice-hearing (Ford, Gray, Faustman, Roach, & Mathalon, 2007; Ford & Mathalon, 2005; Heinks-Maldonado et al., 2007). Importantly, for present purposes, corollary discharge can be understood as a kind of prior (Friston, 2010), and low-precision priors have been associated with the severity in psychotic symptoms and disturbances of agency in people with schizophrenia (Rösler et al., 2015). A reduction of precision on highlevel priors in the psychedelic state means that the corollary discharges that would usually cancel out the expected consequences of actions fail to do so, generating an increase in prediction error at lower levels. These unexpected consequences are then attributed to external rather than internal causes, as the more prediction error is generated, the more likely an action (or thought) has exogenous rather than endogenous causes (Frith, 2003). This echoes similar themes in the autism literature. In autism, the failure of sensory attenuation "leads to the hypervigilant attention to sensory detail at the expense of a hierarchically deep explanation for sensations" (Picard & Friston, 2014, p. 1116) leading to what has been termed a "loss of central coherence" (Frith, 2003). Attribution to exogenous rather than en-

Deane, G. (2020). Dissolving the self: Active inference, psychedelics, and ego-dissolution. Philosophy and the Mind Sciences, 1(I), 2. https://doi.org/10.33735/phimisci.2020.I.38 © © © The author(s). https://philosophymindscience.org ISSN: 2699-0369

dogenous causes could result in a loss of "perceptual mineness" – the background feeling that my experiences are "mine" – if, as has been argued, perceptual mineness is underpinned by anticipation of changes in perceptual inputs in relation to movements (Hohwy, 2007).⁶

Ego-dissolution is not, however, confined to a loss of agentive control over immediate action outcomes, but may be characterised by a more profound dissolution of the sense of being a self or "I" distinct from the outside world. On the view presented in this paper, pre-reflective self-consciousness arises not just through modelling control over the most immediate sensory consequences of actions, but is bolstered by inferences about endogenous control over the distal sensory consequences of allostatic action and action policies. Under a high dose of a psychedelic, the temporary suspension on the gating mechanism on incoming sensory data, described in this paper in terms of a high Bayesian learning rate, render both the proximal and distal sensory consequences of actions highly unpredictable, and the system ceases to have the sense of their being an agent which can (and should) be controlling sensory outcomes. Several authors have emphasised the psychedelic experience is a dynamic process as opposed to a firmly designated state (Masters & Houston, 1966; Preller & Vollenweider, 2016), and different types of ego-dissolution might occur both over the course of the experience and at different dosages. For example, inferences on the boundaries of the body (Blanke & Metzinger, 2009) might be increasingly blurred due to a failure to attenuate the flurry of low-level prediction error. Aspects of the self-model corresponding to longer timescales may break down due to a sustained failure of high levels to attenuate prediction error from low levels due to highly volatile prediction errors, consistent with the fact that bodily ego-dissolution tends to precede dissolution of narrative self (Savage, 1955). This fact is also perhaps suggestive, in opposition to the high-levels posited by the REBUS model, that ego-dissolution could be seen as the *result* of the high-levels failing to contextualise the upsurge of prediction error from across the cortex. The fact that the highest level of the self-model are "increasingly abstract, complex and invariant" (Limanowski & Friston, 2018, p. 5), may explain why higher levels of the self-model are going to be less perturbed by prediction error and perhaps only reliably altered at high dosages. Empirical exploration of these possibilities might be a fruitful avenue for future work, in particular through bridging the neurocomputational mechanisms posited here to both the dynamics of the experience as uncovered through "microphenomenological" interviews (Millière, 2017; Petitmengin, 2006), and to the underlying neural correlates of the experience (Timmermann et al., 2019).

The account of ego-dissolution in terms of a collapse in the temporal thickness

⁶It is worth noting that while some comparisons can be drawn between the psychedelic state and predictive processing accounts of psychosis (Sterzer et al., 2018) and autism (Palmer, Lawson, & Hohwy, 2017), a key difference here is that in the psychedelic state the system typically begins "from a baseline of stable high-level priors (e.g., a stable ego) to which it returns as drug effects subside" (Carhart-Harris & Friston, 2019, p. 328).

Deane, G. (2020). Dissolving the self: Active inference, psychedelics, and ego-dissolution. *Philosophy and the Mind Sciences*, 1(I), 2. https://doi.org/10.33735/phimisci.2020.I.38

of the affordance landscape presented here should also apply to the concept of a "cognitive affordance" landscape, where the "central function of autonomous activity in the mind wandering network is to create a constant stream of affordances for cognitive agency, a continuing internal competition among possible cognitive actions" (Metzinger, 2017, p. 2). Metzinger argues that mental actions – such as the volitional control of endogenous attention, or retrieval of an episodic memory - have epistemic rather that pragmatic goal states. On the allostatic control model of selfhood, the self-model would be constructed and bolstered relative not only to the expectations of the control of the sensory consequences of actions, but also the consequences of mental actions, where the consequences of a mental action might be epistemic and also interoceptive (consider a case where a memory triggers an autonomic response which subsequently acts as the afferent input to an interoceptive inference underpinning a felt emotion). Under psychedelics, loss of control of the expected outcomes of mental actions (as well as a loss of the pragmatic concerns usually driving which epistemic actions to take) might then also be fundamental to the experience of ego-dissolution. This idea is consistent with the fact that under psychedelics mental phenomena "take on the character of objective reality" (Savage, 1955, p. 12), where the ownership of mental phenomena seems to subside and "the individual may feel like a bystander watching the mental activity of another person" (Girn & Christoff, 2018, p. 145).

It is worth mentioning a potential implication of this view for consciousness science more broadly. The psychedelic experience and ego-dissolution are often described as an "expansion" of consciousness. Friston (2018) argues that not only self-consciousness, but consciousness itself, is underpinned by temporal thickness: "consciousness is nothing more than inference about my future; namely, the self-evidencing consequences of what I could do" (Friston, 2018, p. 1). States of ego-dissolution, understood as collapse in the temporal thickness of the generative model, suggest that while temporal thickness very much structures our normal waking experience, it is not clear that temporal thickness ought to be equated with consciousness per se (see also Metzinger, 2020, this issue; Sebastián, 2020, this issue).

5.1 Ecstatic ego-dissolution and challenging experiences

The question remains as to why the hypothesised collapse in the temporal thickness of the self-model under psychedelics can be both ecstatic and of enduring therapeutic value. To bring this into focus, it's worth recapitulating core features of the self-model provided earlier. Recall, interoceptive inference on states of the embodied self "attunes" organisms to their affordance landscape, where inferences about the state of the embodied self (e.g. hunger) prescribe certain prediction error minimising policies (e.g. finding food). Inferences pertaining to allostatic consequences on longer timescales may mean higher-level imperatives trump lowerlevel drives, such as choosing to abstain from chocolate cake to stay healthy (Pezzulo et al., 2018). In the case of basic bodily needs, as described, these variables are controlled (Seth, 2015) through action - active inference is deployed to bring the world into line with predictions, rather than adjusting predictions (via perceptual inference) to conform to the world - for instance eating when hungry (Pezzulo et al., 2015). In just the same way that a hungry organism can act so as to harvest confirmatory evidence for the hypothesis "I am sated", hypotheses relating to higher-levels of the self-model geared towards control of outcomes on longer timescales act to constrain action in the present to bring downstream outcomes closer in line with the prior expectation. Overly precise priors driving action on a long timescale which are failing to be fulfilled, on this view, would be a persistent cause of suffering, due to the system consistently failing to meet (or align actions towards) the goal state (Hesp, Smith, Allen, Friston, & Ramstead, 2019). Under the model of psychedelic-induced ego-dissolution proposed, the high-precision highlevel priors geared towards control on multiple timescales cease to exert influence on the system due to the proposed lowering of precision of high-level priors under psychedelics. If action ordinarily arises from a process of minimising deviations between the organism's actual (inferred) and desired trajectory (Friston et al., 2010), the loss of precision on high-level priors means that, instead of driving action policies, they lose influence on the rest of the system and cease to structure pre-reflective self-consciousness to orient to action opportunities favouring their fulfilment. As these prior beliefs are relaxed, they instead become amenable to perceptual revision from the influx of (highly precise) interoceptive and exteroceptive information. The collapse in temporal depth in the psychedelic state is therefore not experienced as a loss of allostatic control, precisely for the reason that the priors pertaining to longer timescales are no longer asserting an influence on the system and constraining action (and perception) in their usual manner. This picture seems to align well with phenomenological reports of ego-dissolution: "It felt as if 'I' did no longer exist. There was purely my sensory perception of my environment, but sensory input was not translated into needs, feelings, or acting by 'me'" (unpublished online survey data quoted in Millière et al., 2018, p. 7). Peak experiences under psychedelics, then, could be understood as absence of prediction errors relating to allostasis due to a flattening of the temporal depth of the affordance landscape, resulting in the feeling of "oceanic boundlessness" - a sense of immense well-being and peace. Here, the "itinerant strategies" to stay within our "species-specific window of viability" (Clark, 2013, p. 13), are no longer necessary as the "first prior" - the expectation or imperative for existence - is being met without conditions.

Following the TIBER model, many psychopathologies may be due to high precision on high-level priors (Carhart-Harris, 2019; Clark, Watson, & Friston, 2018). Peak psychedelic experience may act as a "reset" allowing for revision of entrenched high-level beliefs that structure pre-reflective self-consciousness (and, accordingly, the affordance landscape) – opening up new domains of salience and possibility for meaningful engagement with the world, through revised and

retuned self-models. Increased bottom-up information flow (particularly from the limbic system), through a high Bayesian learning rate, may make entrenched high-level priors amenable to revision via perceptual inference rather than driving control via active inference. This lays the theoretical groundwork for why psychedelics may effectively treat depression: if depression is underpinned by a high precision prior of low allostatic self-efficacy (Stephan et al., 2016), it follows that relaxation and revision of this prior should alleviate depressive symptoms. Finally, (and speculatively), if the account of "retuning" of self-models under psychedelics presented here generalises to the bodily self (which the experiential changes in bodily selfhood would suggest) this account is suggestive of a potential role for psychedelics in the treatment of chronic pain, and for phantom limb pain – for which there has already been promising results (Fanciullacci, Bene, Franchi, & Sicuteri, 1977; Ramachandran, Chunharas, Marcus, Furnish, & Lin, 2018).

The primary focus so far has been on "peak" experiences, due to the growing number of papers indicating they are central to positive long-term therapeutic outcomes (Roseman et al., 2018). However, while generally psychedelics are thought to be very low risk (Nutt, King, & Phillips, 2010), and there is evidence to suggest they are protective against mental health problems (Hendricks, Thorne, Clark, Coombs, & Johnson, 2015), acute and occasionally persistent adverse psychological reactions do sometimes occur (Strassman, 1984). While "complete" ego-dissolution is described as a "state of complete surrender, associated bliss, and union with all things" (Carhart-Harris & Friston, 2019, p. 321), "incomplete" ego-dissolution due to psychological resistance or an insufficient dose – can be characterised by intense fear, anxiety, or distress. On the account presented in this paper, this can be understood as resulting from psychological resistance, where psychological resistance here may be conceptualised as a high-precision prior on being able to control the experience, that is maintained though fear-driven endogenous attention. Failure to control the experience, in violating the highly precise prior for the goal state of control, is then experienced as a loss of allostatic control, bringing with it feelings of intense fear or distress. In therapeutic contexts, encouraging users to "let go" and "surrender" to the experience (Richards, 2015), could be understood in these terms, as discouraging the user from putting high (endogenous) precision on a prior for control that could result in adverse experiences when unfulfilled. These considerations highlight the essential importance of context in achieving therapeutic outcomes (Carhart-Harris et al., 2018).

6 Conclusion

Psychedelics are known for their ability to profoundly alter consciousness and occasion so-called "mystical" experiences (Huxley, 2010). The renaissance in psychedelic research in the past decade is beginning to shed light on the mechanisms underpinning the extraordinary states of consciousness induced by psychedelics (Carhart-Harris & Goodwin, 2017). Within psychedelic phenomenology, experiences of ego-dissolution are of particular phenomenological, philosophical and therapeutic interest (Letheby & Gerrans, 2017; Millière, 2017; Nour & Carhart-Harris, 2017). This paper has given a preliminary account of how ego-dissolution under psychedelics can be understood in terms of predictive processing and active inference. The hypothesis here is that the action of psychedelics within the predictive processing framework is best understood as a "relaxation of high-level beliefs" (Carhart-Harris, 2019), and this can be unpacked in terms of a high Bayesian learning rate (Hohwy, 2017; Mathys et al., 2014). Psychedelic-induced ego-dissolution, then, results in a collapse in temporal thickness (Friston, 2018) of the self-model as conceived within an active inference framework. The therapeutic effects of ego-dissolution, then, can be understood in terms of the relaxing and retuning of entrenched self-models, or a "resetting" or "opening" of the affordance landscape, allowing for the possibility of new modes of engagement with the world, oneself, and other people.

Acknowledgments

Thanks very much to Mel Andrews, Andrew Corcoran, and Andy Clark for invaluable critical feedback and discussion in the production of this paper. Work on this article was supported by the European Research Council (ERC) Advanced Grant XSPECT - DLV-692739.

References

- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., & Friston, K. (2013). The computational anatomy of psychosis. Frontiers in Psychiatry, 4, 47. https://doi.org/10.3389/fpsyt.2013.00047
- Allen, M., & Friston, K. (2018). From cognitivism to autopoiesis: Towards a computational framework for the embodied mind. Synthese, 195(6), 2459–2482. https://doi.org/10.1007/s11229-016-1288-5
- Allen, M., & Tsakiris, M. (2018). The body as first prior: Interoceptive predictive processing and the primacy. In *The Interoceptive Mind: From Homeostasis to Awareness* (pp. 27–45). https://doi.org/10.1093/oso/9780198811930.003.0002
- Anchisi, D., & Zanon, M. (2015). A Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia. *PloS One*, 10(2), e0117270. https://doi.org/10.1371/journal.pone.0117270
- Apps, M. A., & Tsakiris, M. (2014). The free-energy self: A predictive coding account of self-recognition. Neuroscience & Biobehavioral Reviews, 41, 85–97. https://doi.org/10.1016/j.neubiorev.2013.01.029
- Barlow, H. (2001). Redundancy reduction revisited. Network: Computation in Neural Systems, 12(3), 241–253. https://doi.org/10.1088/0954-898x/12/3/301
- Barre, A., Berthoux, C., De Bundel, D., Valjent, E., Bockaert, J., Marin, P., & Bécamel, C. (2016). Presynaptic serotonin 2A receptors modulate thalamocortical plasticity and associative learning. *Proceedings of the National Academy of Sciences*, 113(10), E1382–E1391. https://doi.org/10.1073/pnas.1525586113
- Barrett, L. F. (2017). The theory of constructed emotion: An active inference account of interoception and categorization. Social Cognitive and Affective Neuroscience, 12(1), 1–23. https://doi.org/10.1093/scan/nsw154
- Barrett, L. F., Quigley, K. S., & Hamilton, P. (2016). An active inference theory of allostasis and interoception in depression. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1708), 20160011. https://doi.org/10.1098/rstb.2016.0011
- Barrett, L. F., & Simmons, W. K. (2015). Interoceptive predictions in the brain. Nature Reviews Neuroscience, 16(7), 419–429. https://doi.org/10.1038/nrn3950
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4), 695–711. https://doi.org/10.1016/j.neuron.2012.10.038
- Bayne, T., & Carter, O. (2018). Dimensions of consciousness and the psychedelic state. Neuroscience of Consciousness, 2018(1), niy008. https://doi.org/10.1093/nc/niy008

- Beliveau, V., Ganz, M., Feng, L., Ozenne, B., Højgaard, L., Fisher, P. M., et al. (2017). A high-resolution in vivo atlas of the human brain's serotonin system. *Journal of Neuroscience*, *37*(1), 120–128. https://doi.org/10.1038/sdata.2018.63
- Berthoux, C., Barre, A., Bockaert, J., Marin, P., & Bécamel, C. (2018). Sustained activation of postsynaptic 5-HT2A receptors gates plasticity at prefrontal cortex synapses. *Cerebral Cortex*, 29(4), 1659–1669. https://doi.org/10.1093/cercor/bhy064
- Blanke, O., & Metzinger, T. (2009). Full-body illusions and minimal phenomenal selfhood. Trends in Cognitive Sciences, 13(1), 7–13. https://doi.org/10.1016/j.tics.2008.10.003
- Botvinick, M., & Toussaint, M. (2012). Planning as inference. *Trends in Cognitive Sciences*, 16(10), 485–488. https://doi.org/10.1016/j.tics.2012.08.006
- Brown, H., Adams, R. A., Parees, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. Cognitive Processing, 14(4), 411–427. https://doi.org/10.1007/s10339-013-0571-3
- Bruineberg, J. (2017). Active inference and the primacy of the 'I can'. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing*. Frankfurt am Main: MIND Group.
- Buckley, C. L., Kim, C. S., McGregor, S., & Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81, 55–79. https://doi.org/10.1016/j.jmp.2017.09.004
- Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends in Cognitive Sciences*, 11(2), 49–57. https://doi.org/10.1016/j.tics.2006.11.004
- Büchel, C., Geuter, S., Sprenger, C., & Eippert, F. (2014). Placebo analgesia: A predictive coding perspective. Neuron, 81(6), 1223–1239. https://doi.org/10.1016/j.neuron.2014.02.042
- Cannon, W. B. (1927). The James-Lange theory of emotions: A critical examination and an alternative theory. *The American Journal of Psychology*, *39*(1/4), 106–124. https://doi.org/10.2307/1415404
- Cannon, W. B. (1929). Organization for physiological homeostasis. *Physiological Reviews*, 9(3), 399–431. https://doi.org/10.1152/physrev.1929.9.3.399
- Carhart-Harris, R. L. (2019). How do psychedelics work? *Current Opinion in Psychiatry*, 32(1), 16–21. https://doi.org/10.1097/yco.00000000000467
- Carhart-Harris, R. L., & Friston, K. (2019). REBUS and the anarchic brain: Toward a unified model of the brain action of psychedelics. *Pharmacological Reviews*, 71(3), 316–344. https://doi.org/10.1124/pr.118.017160
- Carhart-Harris, R. L., & Goodwin, G. M. (2017). The therapeutic potential of psychedelic drugs: Past, present, and future. *Neuropsychopharmacology*, 42(11), 2105. https://doi.org/10.1038/npp.2017.84
- Carhart-Harris, R. L., & Nutt, D. (2017). Serotonin and brain function: A tale of two receptors. Journal of Psychopharmacology, 31(9), 1091–1120. https://doi.org/10.1177/0269881117725915
- Carhart-Harris, R. L., Roseman, L., Haijen, E., Erritzoe, D., Watts, R., Branchi, I., & Kaelen, M. (2018). Psychedelics and the essential importance of context. *Journal of Psychopharmacology*, 32(7), 725–731. https://doi.org/10.1177/0269881118754710
- Carter, O. L., Hasler, F., Pettigrew, J. D., Wallis, G. M., Liu, G. B., & Vollenweider, F. X. (2007). Psilocybin links binocular rivalry switch rate to attention and subjective arousal levels in humans. *Psychopharmacology*, 195(3), 415–424. https://doi.org/10.1007/s00213-007-0930-9
- Carter, O. L., Pettigrew, J. D., Hasler, F., Wallis, G. M., Liu, G. B., Hell, D., & Vollenweider, F. X. (2005). Modulating the rate and rhythmicity of perceptual rivalry alternations with the mixed 5-HT 2A and 5-HT 1A agonist psilocybin. *Neuropsychopharmacology*, 30(6), 1154. https://doi.org/10.1038/sj.npp.1300621
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485), 1585–1599. https://doi.org/10.1017/cbo9780511731525.015
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behavioral and Brain Sciences, 36(3), 181–204. https://doi.org/10.1017/s0140525x12000477
- Clark, A. (2019). Consciousness as generative entanglement. *The Journal of Philosophy*, *116*, 645–662. https://doi.org/10.5840/jphil20191161241
- Clark, J. E., Watson, S., & Friston, K. (2018). What is mood? A computational perspective. *Psychological Medicine*, 48(14), 2277–2284. https://doi.org/10.1017/s0033291718000430
- Conant, R. C., & Ashby, R. W. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97. https://doi.org/10.1080/00207727008920220
- Corcoran, A., & Hohwy, J. (2018). Allostasis, interoception, and the free energy principle: Feeling our way forward. https://doi.org/10.1093/oso/9780198811930.003.0015
- Corcoran, A. W., Pezzulo, G., & Hohwy, J. (2019). From allostatic agents to counterfactual cognisers: Active inference, biological regulation, and the origins of cognition. https://doi.org/10.20944/preprints201911.0083.v1

- Corlett, P. R., Frith, C. D., & Fletcher, P. C. (2009). From drugs to deprivation: A Bayesian framework for understanding models of psychosis. *Psychopharmacology*, 206(4), 515–530. https://doi.org/10.1007/s00213-009-1561-0
- Corlett, P. R., Honey, G. D., & Fletcher, P. C. (2016). Prediction error, ketamine and psychosis: An updated model. Journal of Psychopharmacology, 30(11), 1145–1155. https://doi.org/10.1177/0269881116650087
- Craig, A. D. (2002). How do you feel? Interoception: The sense of the physiological condition of the body. Nature Reviews Neuroscience, 3(8), 655. https://doi.org/10.1038/nrn894
- Crapse, T. B., & Sommer, M. A. (2008). Corollary discharge across the animal kingdom. Nature Reviews Neuroscience, 9(8), 587. https://doi.org/10.1038/nrn2457
- Damasio, A., & Carvalho, G. B. (2013). The nature of feelings: Evolutionary and neurobiological origins. Nature Reviews Neuroscience, 14(2), 143. https://doi.org/10.1038/nrn3403
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. Cognitive, Affective, & Behavioral Neuroscience, 8(4), 429–453. https://doi.org/10.3758/cabn.8.4.429
- Eisenberger, N. I. (2012). The pain of social disconnection: Examining the shared neural underpinnings of physical and social pain. *Nature Reviews Neuroscience*, *13*(6), 421. https://doi.org/10.1038/nrn3231
- Eisenberger, N. I., Jarcho, J. M., Lieberman, M. D., & Naliboff, B. D. (2006). An experimental study of shared sensitivity to physical pain and social rejection. *Pain*, *126*(1-3), 132–138. https://doi.org/10.1016/j.pain.2006.06.024
- Fanciullacci, M., Bene, E. D., Franchi, G., & Sicuteri, F. (1977). Phantom limb pain: Sub-hallucinogenic treatment with lysergic acid diethylamide (LSD-25). *Headache: The Journal of Head and Face Pain, 17*(3), 118–119. https://doi.org/10.1111/j.1526-4610.1977.hed1703118.x
- Feldman, H., & Friston, K. (2010). Attention, uncertainty, and free-energy. Frontiers in Human Neuroscience, 4, 215. https://doi.org/10.3389/fnhum.2010.00215
- Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: A Bayesian approach to explaining the positive symptoms of schizophrenia. Nature Reviews Neuroscience, 10(1), 48. https://doi.org/10.1038/nrn2536
- Ford, J. M., Gray, M., Faustman, W. O., Roach, B. J., & Mathalon, D. H. (2007). Dissecting corollary discharge dysfunction in schizophrenia. *Psychophysiology*, 44(4), 522–529. https://doi.org/10.1111/j.1469-8986.2007.00533.x
- Ford, J. M., & Mathalon, D. H. (2005). Corollary discharge dysfunction in schizophrenia: Can it explain auditory hallucinations? *International Journal of Psychophysiology*, 58(2-3), 179–189. https://doi.org/10.1016/j.ijpsycho.2005.01.014
- Friston, K. (2010). The free-energy principle: A unified brain theory? Nature Reviews Neuroscience, 11(2), 127. https://doi.org/10.1038/nrn2787
- Friston, K. (2012a). Embodied inference and spatial cognition. Cognitive Processing, 13(1), 171–177. https://doi.org/10.1007/s10339-012-0519-z
- Friston, K. (2012b). Prediction, perception and agency. International Journal of Psychophysiology, 83(2), 248–252. https://doi.org/10.1016/j.ijpsycho.2011.11.014
- Friston, K. (2018). Am I self-conscious?(Or does self-organization entail self-consciousness?). Frontiers in Psychology, 9, 579. https://doi.org/10.3389/fpsyg.2018.00579
- Friston, K., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: A free-energy formulation. *Biological Cybernetics*, 102(3), 227–260. https://doi.org/10.1007/s00422-010-0364-z
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. Cognitive Neuroscience, 6(4), 187–214. https://doi.org/10.1080/17588928.2015.1020053
- Friston, K., Rosch, R., Parr, T., Price, C., & Bowman, H. (2018). Deep temporal models and active inference. Neuroscience & Biobehavioral Reviews, 90, 486–501. https://doi.org/10.1016/j.neubiorev.2017.04.009
- Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2013). The anatomy of choice: Active inference and agency. *Frontiers in Human Neuroscience*, 7, 598. https://doi.org/10.3389/fnhum.2013.00598
- Friston, K., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., et al. (2012). Dopamine, affordance and active inference. PLoS Computational Biology, 8(1), e1002327. https://doi.org/10.1371/journal.pcbi.1002327
- Frith, U. (2003). Autism: Explaining the enigma. Oxford: Blackwell Publishing.
- Garrido, M. I., Kilner, J. M., Stephan, K. E., & Friston, K. (2009). The mismatch negativity: A review of underlying mechanisms. *Clinical Neurophysiology*, 120(3), 453–463. https://doi.org/10.1016/j.clinph.2008.11.029
- Gershman, S. J. (2019). What does the free energy principle tell us about the brain? Neurons, Behavior, Data Analysis, and Theory, 4(1), 1-10.
- Girn, M., & Christoff, K. (2018). Journal of Consciousness Studies, 25(11-12), 131–154. Retrieved from https://www.ingentaconnect.com/content/imp/jcs/2018/00000025/f0020011/art00008
- Griffiths, R. R., Johnson, M. W., Carducci, M. A., Umbricht, A., Richards, W. A., Richards, B. D., et al. (2016). Psilocybin produces substantial and sustained decreases in depression and anxiety in patients with life-threatening cancer: A randomized double-blind trial. *Journal of Psychopharmacology*, 30(12), 1181–1197. https://doi.org/10.1177/0269881116675513

Grof, S. (1980). LSD psychotherapy. Pomona, CA: Hunter House.

- Heinks-Maldonado, T. H., Mathalon, D. H., Houde, J. F., Gray, M., Faustman, W. O., & Ford, J. M. (2007). Relationship of imprecise corollary discharge in schizophrenia to auditory hallucinations. *Archives of General Psychiatry*, 64(3), 286–296. https://doi.org/10.1001/archpsyc.64.3.286
- Hendricks, P. S., Thorne, C. B., Clark, C. B., Coombs, D. W., & Johnson, M. W. (2015). Classic psychedelic use is associated with reduced psychological distress and suicidality in the united states adult population. *Journal of Psychopharmacology*, 29(3), 280–288. https://doi.org/10.1177/0269881114565653
- Hesp, C., Smith, R., Allen, M., Friston, K., & Ramstead, M. (2019). Deeply felt affect: The emergence of valence in deep active inference. *PsyArXiv*. https://doi.org/10.31234/osf.io/62pfd

Hohwy, J. (2007). The sense of self in the phenomenology of agency and perception. Psyche, 13(1), 1-20.

- Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, *3*, 96. https://doi.org/10.3389/fpsyg.2012.00096
- Hohwy, J. (2017). Priors in perception: Top-down modulation, Bayesian perceptual learning rate, and prediction error minimization. Consciousness and Cognition, 47, 75–85. https://doi.org/10.1016/j.concog.2016.09.004
- Hohwy, J., & Michael, J. (2017). Why should any body have a self? In F. de Vignemont & A. Alsmith (Eds.), The subject's matter: Self-consciousness and the body (pp. 363–391). Cambridge, MA: MIT Press.
- Huxley, A. (2010). The doors of perception: And heaven and hell. New York, NY: Random House.
- Jakab, R. L., & Goldman-Rakic, P. S. (1998). Proceedings of the National Academy of Sciences, 95(2), 735–740. https://doi.org/10.1073/pnas.95.2.735
- Kaplan, R., & Friston, K. (2018). Planning and navigation as active inference. Biological Cybernetics, 112(4), 323–343. https://doi.org/10.1007/s00422-018-0753-2
- Keller, G. B., & Mrsic-Flogel, T. D. (2018). Predictive processing: A canonical cortical computation. Neuron, 100(2), 424–435. https://doi.org/10.1016/j.neuron.2018.10.003
- Kiebel, S. J., Daunizeau, J., & Friston, K. (2008). A hierarchy of time-scales and the brain. PLoS Computational Biology, 4(11), e1000209. https://doi.org/10.1371/journal.pcbi.1000209
- King, A., Martin, I., & Seymour, K. (1972). Reversal learning facilitated by a single injection of lysergic acid diethylamide (LSD 25) in the rat. British Journal of Pharmacology, 45(1), 161P. https://doi.org/10.1111/j.1476-5381.1974.tb08611.x
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. Trends in Neurosciences, 27(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007
- Kometer, M., Cahn, B. R., Andel, D., Carter, O. L., & Vollenweider, F. X. (2011). The 5-HT2A/1A agonist psilocybin disrupts modal object completion associated with visual hallucinations. *Biological Psychiatry*, 69(5), 399–406. https://doi.org/10.1016/j.biopsych.2010.10.002
- Leary, T. F., Metzner, R., & Alpert, R. (1964). The psychedelic experience: A manual based on the Tibetan book of the dead.
- Lebedev, A. V., Lövdén, M., Rosenthal, G., Feilding, A., Nutt, D. J., & Carhart-Harris, R. L. (2015). Finding the self by losing the self: Neural correlates of ego-dissolution under psilocybin. *Human Brain Mapping*, 36(8), 3137–3153. https://doi.org/10.1002/hbm.22833
- Letheby, C. (2020). Being for no-one: Psychedelic experience and minimal subjectivity. *Philosophy and the Mind Sciences*, *1*(I), 5. https://doi.org/10.33735/phimisci.2020.I.47
- Letheby, C., & Gerrans, P. (2017). Self unbound: Ego dissolution in psychedelic experience. *Neuroscience of Consciousness*, 3(1), nix016. https://doi.org/10.1093/nc/nix016
- Limanowski, J., & Blankenburg, F. (2013). Minimal self-models and the free energy principle. Frontiers in Human Neuroscience, 7, 547. https://doi.org/10.3389/fnhum.2013.00547
- Limanowski, J., & Friston, K. (2018). "Seeing the dark": Grounding phenomenal transparency and opacity in precision estimation for active inference. *Frontiers in Psychology*, 9, 643. https://doi.org/10.3389/fpsyg.2018.00643
- Limanowski, J., & Friston, K. (2020). Attenuating oneself: An active inference perspective on "selfless" experiences. *Philosophy and the Mind Sciences*, 1(I), 6. https://doi.org/10.33735/phimisci.2020.I.35
- Ly, C., Greb, A. C., Cameron, L. P., Wong, J. M., Barragan, E. V., Wilson, P. C., et al. (2018). Psychedelics promote structural and functional neural plasticity. *Cell Reports*, 23(11), 3170–3182. https://doi.org/10.1016/j.celrep.2018.05.022
- Lyons, T., & Carhart-Harris, R. L. (2018). More realistic forecasting of future life events after psilocybin for treatment-resistant depression. *Frontiers in Psychology*, 9, 1721. https://doi.org/10.3389/fpsyg.2018.01721
- Masters, R. E., & Houston, J. (1966). *The varieties of psychedelic experience* (Vol. 9289). New York, Chicago, San Francisco: Holt, Rinehart; Winston New York.
- Mathys, C., Daunizeau, J., Friston, K., & Stephan, K. E. (2011). A Bayesian foundation for individual learning under uncertainty. Frontiers in Human Neuroscience, 5, 39. https://doi.org/10.3389/fnhum.2011.00039

- Mathys, C., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K., & Stephan, K. E. (2014). Uncertainty in perception and the hierarchical gaussian filter. *Frontiers in Human Neuroscience*, 8, 825. https://doi.org/10.3389/fnhum.2014.00825
- Metzinger, T. (2003). Being no one: The self-model theory of subjectivity. Cambridge, MA: MIT Press.
- Metzinger, T. (2009). The ego tunnel: The science of the mind and the myth of the self. New York: Basic Books.
- Metzinger, T. (2017). The problem of mental action. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing*. Frankfurt am Mainz: MIND Group.
- Metzinger, T. (2020). Minimal phenomenal experience: Meditation, tonic alertness, and the phenomenology of "pure" consciousness. *Philosophy and the Mind Sciences*, 1(I), 7. https://doi.org/10.33735/phimisci.2020.I.46
- Millière, R. (2017). Looking for the self: Phenomenology, neurophysiology and philosophical significance of drug-induced ego dissolution. *Frontiers in Human Neuroscience*, *11*, 245. https://doi.org/10.3389/fnhum.2017.00245
- Millière, R. (2020). The varieties of selflessness. *Philosophy and the Mind Sciences*, 1(I), 8. https://doi.org/10.33735/phimisci.2020.I.48
- Millière, R., Carhart-Harris, R. L., Roseman, L., Trautwein, F.-M., & Berkovich-Ohana, A. (2018). Psychedelics, meditation, and self-consciousness. Frontiers in Psychology, 9, 1–29. https://doi.org/10.3389/fpsyg.2018.01475
- Montague, P. R., & King-Casas, B. (2007). Efficient statistics, common currencies and the problem of reward-harvesting. Trends in Cognitive Sciences, 11(12), 514–519. https://doi.org/10.1016/j.tics.2007.10.002
- Moreno, F. A., Wiegand, C. B., Taitano, E. K., & Delgado, P. L. (2006). Safety, tolerability, and efficacy of psilocybin in 9 patients with obsessive-compulsive disorder. *Journal of Clinical Psychiatry*, 67(11), 1735–1740. https://doi.org/10.4088/jcp.v67n1110
- Morton, D. L., El-Deredy, W., Watson, A., & Jones, A. K. (2010). Placebo analgesia as a case of a cognitive style driven by prior expectation. Brain Research, 1359, 137–141. https://doi.org/10.1016/j.brainres.2010.08.046
- Moutoussis, M., Fearon, P., El-Deredy, W., Dolan, R. J., & Friston, K. (2014). Bayesian inferences about the self (and others): A review. *Consciousness and Cognition*, 25, 67-76. https://doi.org/10.1016/j.concog.2014.01.009
- Mumford, D. (1992). On the computational architecture of the neocortex. *Biological Cybernetics*, 66(3), 241–251. https://doi.org/10.1007/bf00198477
- Nour, M. M., & Carhart-Harris, R. L. (2017). Psychedelics and the science of self-experience. The British Journal of Psychiatry, 210(3), 177–179. https://doi.org/10.1192/bjp.bp.116.194738
- Nour, M. M., Evans, L., Nutt, D., & Carhart-Harris, R. L. (2016). Ego-dissolution and psychedelics: Validation of the ego-dissolution inventory (EDI). *Frontiers in Human Neuroscience*, *10*, 269. https://doi.org/10.3389/fnhum.2016.00269
- Nutt, D. J., King, L. A., & Phillips, L. D. (2010). Drug harms in the UK: A multicriteria decision analysis. *The Lancet*, 376(9752), 1558–1565. https://doi.org/10.1016/s0140-6736(10)61462-6
- Ongaro, G., & Kaptchuk, T. J. (2019). Symptom perception, placebo effects, and the Bayesian brain. *Pain*, *160*(1), 1. https://doi.org/10.1097/j.pain.00000000001367
- Palhano-Fontes, F., Barreto, D., Onias, H., Andrade, K. C., Novaes, M. M., Pessoa, J. A., et al. (2019). Rapid antidepressant effects of the psychedelic ayahuasca in treatment-resistant depression: A randomized placebo-controlled trial. *Psychological Medicine*, 49(4), 655–663. https://doi.org/10.1017/s0033291718001356
- Palmer, C. J., Lawson, R. P., & Hohwy, J. (2017). Bayesian approaches to autism: Towards volatility, action, and behavior. Psychological Bulletin, 143(5), 521. https://doi.org/10.1037/bul0000097
- Petitmengin, C. (2006). Describing one's subjective experience in the second person: An interview method for the science of consciousness. *Phenomenology and the Cognitive Sciences*, 5(3-4), 229–269. https://doi.org/10.1007/s11097-006-9022-2
- Pezzulo, G. (2014). Why do you fear the bogeyman? An embodied predictive coding model of perceptual inference. Cognitive, Affective, & Behavioral Neuroscience, 14(3), 902–911. https://doi.org/10.3758/s13415-013-0227-x
- Pezzulo, G., Cartoni, E., Rigoli, F., Pio-Lopez, L., & Friston, K. (2016). Active inference, epistemic value, and vicarious trial and error. *Learning & Memory*, 23(7), 322–338. https://doi.org/10.1101/lm.041780.116
- Pezzulo, G., & Cisek, P. (2016). Navigating the affordance landscape: Feedback control as a process model of behavior and cognition. *Trends in Cognitive Sciences*, 20(6), 414–424. https://doi.org/10.1016/j.tics.2016.03.013
- Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active inference, homeostatic regulation and adaptive behavioural control. Progress in Neurobiology, 134, 17–35. https://doi.org/10.1016/j.pneurobio.2015.09.001
- Pezzulo, G., Rigoli, F., & Friston, K. (2018). Hierarchical active inference: A theory of motivated control. Trends in Cognitive Sciences, 22(4), 294–306. https://doi.org/10.1016/j.tics.2018.01.009
- Picard, F., & Friston, K. (2014). Predictions, perception, and a sense of self. Neurology, 83(12), 1112–1118. https://doi.org/10.1212/wnl.00000000000798

- Pink-Hashkes, S., Rooij, I. van, & Kwisthout, J. (2017). Perception is in the details: A predictive coding account of the psychedelic phenomenon. Proceedings of the 39th annual meeting of the cognitive science society, 2907–2912.
- Preller, K. H., & Vollenweider, F. X. (2016). Phenomenology, structure, and dynamic of psychedelic states. In A. L. Halberstadt, F. X. Vollenweider, & D. E. Nichols (Eds.), *Behavioral neurobiology of psychedelic drugs* (pp. 221–256). https://doi.org/10.1007/7854_2016_459
- Ramachandran, V., Chunharas, C., Marcus, Z., Furnish, T., & Lin, A. (2018). Relief from intractable phantom pain by combining psilocybin and mirror visual-feedback (MVF). *Neurocase*, 24(2), 105–110. https://doi.org/10.1080/13554794.2018.1468469
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79. https://doi.org/10.1038/4580
- Richards, W. A. (2015). Sacred knowledge: Psychedelics and religious experiences. New York: Columbia University Press.
- Romano, A. G., Quinn, J. L., Li, L., Dave, K. D., Schindler, E. A., Aloyo, V. J., & Harvey, J. A. (2010). Intrahippocampal LSD accelerates learning and desensitizes the 5-ht 2A receptor in the rabbit, romano et al. *Psychopharmacology*, 212(3), 441–448. https://doi.org/10.1007/s00213-010-2004-7
- Roseman, L., Nutt, D. J., & Carhart-Harris, R. L. (2018). Quality of acute psychedelic experience predicts therapeutic efficacy of psilocybin for treatment-resistant depression. *Frontiers in Pharmacology*, *8*, 974. https://doi.org/10.3389/fphar.2017.00974
- Rösler, L., Rolfs, M., Van der Stigchel, S., Neggers, S. F., Cahn, W., Kahn, R. S., & Thakkar, K. N. (2015). Failure to use corollary discharge to remap visual target locations is associated with psychotic symptom severity in schizophrenia. *Journal of Neurophysiology*, 114(2), 1129–1136. https://doi.org/10.1152/jn.00155.2015
- Savage, C. (1955). Variations in ego feeling induced by D-lysergic acid diethylamide (LSD-25). *Psychoanalytic Review*, 42(1), 1–16.
- Schachter, S., & Singer, J. (1963). Cognitive, social, and physiological determinants of emotional state. Psychological Review, 70(1), 121–122. https://doi.org/10.1037/h0038845
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2008). Episodic simulation of future events: Concepts, data, and applications. Annals of the New York Academy of Sciences, 1124(1), 39–60. https://doi.org/10.1196/annals.1440.001
- Schulkin, J., & Sterling, P. (2019). Allostasis: A brain-centered, predictive mode of physiological regulation. Trends in Neurosciences, 42(10), 740–752. https://doi.org/10.1016/j.tins.2019.07.010
- Scott, G., & Carhart-Harris, R. L. (2019). Psychedelics as a treatment for disorders of consciousness. Neuroscience of Consciousness, 2019(1), niz003. https://doi.org/10.1093/nc/niz003
- Sebastián, M. Á. (2020). Perspectival self-consciousness and ego-dissolution: An analysis of (some) altered states of consciousness. *Philosophy and the Mind Sciences*, 1(1), 9. https://doi.org/10.33735/phimisci.2020.I.44
- Seth, A. (2015). The cybernetic Bayesian brain: From interoceptive inference to sensorimotor contingencies. In T. Metzinger & J. Windt (Eds.), *Open MIND* (pp. 1–24). Frankfurt am Main: MIND Group.
- Seth, A. K., & Friston, K. (2016). Active interoceptive inference and the emotional brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1708), 20160007. https://doi.org/10.1098/rstb.2016.0007
- Seth, A. K., & Tsakiris, M. (2018). Being a beast machine: The somatic basis of selfhood. Trends in Cognitive Sciences, 22(11), 969–981. https://doi.org/10.1016/j.tics.2018.08.008
- Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. Annual Review of Neuroscience, 33, 89–108. https://doi.org/10.1146/annurev-neuro-060909-153135
- Shipp, S. (2016). Neural elements for predictive coding. *Frontiers in Psychology*, 7, 1792. https://doi.org/10.3389/fpsyg.2016.01792
- Siegel, E. H., Sands, M. K., Van den Noortgate, W., Condon, P., Chang, Y., Dy, J., et al. (2018). Emotion fingerprints or emotion populations? A meta-analytic investigation of autonomic features of emotion categories. *Psychological Bulletin*, 144(4), 343. https://doi.org/10.1037/bul0000128
- Stephan, K. E., Manjaly, Z. M., Mathys, C., Weber, L. A., Paliwal, S., Gard, T., et al. (2016). Allostatic self-efficacy: A metacognitive theory of dyshomeostasis-induced fatigue and depression. *Frontiers in Human Neuroscience*, 10, 550. https://doi.org/10.3389/fnhum.2016.00550
- Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology & Behavior*, 106(1), 5–15. https://doi.org/10.1016/j.physbeh.2011.06.004
- Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., et al. (2018). The predictive coding account of psychosis. *Biological Psychiatry*, 84(9), 634–643. https://doi.org/10.1016/j.biopsych.2018.05.015
- Strassman, R. J. (1984). Adverse reactions to psychedelic drugs. A review of the literature. Journal of Nervous and Mental Disease, 172(10), 577–595. https://doi.org/10.1097/00005053-198410000-00001
- Timmermann, C., Roseman, L., Schartner, M., Millière, R., Williams, L., Erritzoe, D., et al. (2019). Neural correlates of the DMT experience as assessed via multivariate EEG. *bioRxiv*, 706283. https://doi.org/10.1101/706283

- Timmermann, C., Spriggs, M. J., Kaelen, M., Leech, R., Nutt, D. J., Moran, R. J., et al. (2018). LSD modulates effective connectivity and neural adaptation mechanisms in an auditory oddball paradigm. *Neuropharmacology*, 142, 251–262. https://doi.org/10.1016/j.neuropharm.2017.10.039
- Wager, T. D., Kang, J., Johnson, T. D., Nichols, T. E., Satpute, A. B., & Barrett, L. F. (2015). A Bayesian model of category-specific emotional brain responses. *PLoS Computational Biology*, 11(4), e1004066. https://doi.org/10.1371/journal.pcbi.1004066
- Wiese, W., & Metzinger, T. (2017). Vanilla PP for philosophers: A primer on predictive processing. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing*. https://doi.org/10.15502/9783958573024
- Wilkinson, S., Deane, G., Nave, K., & Clark, A. (2019). Getting warmer: Predictive processing and the nature of emotion. In L. Candiotto (Ed.), *The value of emotions for knowledge* (pp. 101–119). https://doi.org/10.1007/978-3-030-15667-1
- Yoshida, W., Dolan, R. J., & Friston, K. (2008). Game theory of mind. *PLoS Computational Biology*, 4(12), e1000254. https://doi.org/10.1371/journal.pcbi.1000254

Open Access

This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.