From Allostatic Agents to Counterfactual Cognisers: Active Inference, Biological Regulation, and The Origins of Cognition

Andrew W. Corcoran[1]

Giovanni Pezzulo[2]

Jakob Hohwy[1]

[1]Cognition & Philosophy Laboratory, School of Philosophical, Historical, & International Studies, Monash University, Australia

[2]Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche, Roma, Italy

Correspondence: andrew.corcoran1@monash.edu

Abstract

What is the function of cognition? On one influential account, cognition evolved to co-ordinate behaviour with environmental change or complexity (Godfrey-Smith, 1996). Liberal interpretations of this view ascribe cognition to an extraordinarily broad set of biological systems – even bacteria, which modulate their activity in response to salient external cues, would seem to qualify as cognitive agents. However, equating cognition with adaptive flexibility *per se* glosses over important distinctions in the way biological organisms deal with environmental complexity. Drawing on contemporary advances in theoretical biology and computational neuroscience, we cash these distinctions out in terms of the representation and resolution of different varieties of uncertainty. This analysis leads us to propose a formal criterion for delineating cognition from other, more pervasive forms of adaptive plasticity. On this view, biological cognition is rooted in a particular kind of functional organisation; namely, one that enables the agent to detach from the present and engage in counterfactual (active) inference.

Keywords: complexity; uncertainty; cognition; allostasis; homeostasis; free energy principle; active inference; environmental complexity thesis

1 Introduction

What is cognition? What is it *for*? While the former question is a perennial source of philosophical dispute, the latter seems to attract rather less controversy. Cognition – whatever it consists in and however realised – is ultimately functional to adaptive success. It enables the organism to register information about the state of its environment, and to exploit such information in the service of adaptive behaviour. Cognition, in short, is for *action*.

As benign as this characterisation might appear on first blush, a host of thornier questions lie in wait: Are *all* varieties of adaptive behaviour mediated by cognition, or only a select few? If the former, does this notion of behaviour extend to artificial and multi-agent systems, or is it limited to individual organisms? If the latter, what properties distinguish cognitive from non-cognitive modes of behaviour (assuming there *is* a clear distinction to be made)? And what of those cognitive processes that seem entirely encapsulated from one's present transactions with the world – how do they fit into the picture?

This paper attempts to approach some of these difficult questions indirectly, via an analysis of the principles by which cognition might have evolved. This broadly *teleonomic* strategy – whereby cognitive processes are understood in terms of their fitness-enhancing properties – draws inspiration from Peter Godfrey-Smith's (1996) *environmental complexity thesis*. On this view, cognition evolved to co-ordinate organismic behaviour with certain complex (i.e. heterogeneous or variable) properties of the eco-niche. Thus construed, cognition functions to generate flexible patterns of behaviour in response to fluctuating environmental conditions.

We shall not dwell on the details of the environmental complexity thesis here. What interests us, rather, is how the general shape of Godfrey-Smith's explanatory framework – taken in conjunction with more recent advances in theoretical biology, computational neuroscience, and related disciplines – can inform contemporary philosophical debates about the nature of (biological) cognition. Drawing on insights afforded by these fields, we analyse complexity into different kinds of *uncertainty*, and show how distinctive profiles of adaptive plasticity might emerge as a function of uncertainty reduction. This analysis suggests behavioural flexibility *per se* is not sufficient to determine the cognitive status of an adaptive organism. Rather, we propose a narrower conception of cognition as a process rooted in a particular kind of functional organisation; namely, one that affords the capacity to model and interrogate counterfactual possibilities. We construe such cognitive architectures as adaptations to higher-order forms of environmental uncertainty.

This paper is structured as follows: Section 2 begins by considering the homeostatic challenges posed by fluctuating environments. We approach this topic from the perspective of the *free energy principle* (Friston, 2010), a formal account of the autopoietic processes by which biological systems organise and sustain themselves as adaptive agents. Section 3 outlines how the theoretical resources of the free energy principle extend to predictive (i.e. allostatic) forms of biological regulation. We focus on two complementary formulations of allostasis, highlighting how these hierarchical control schemes inform fundamental questions about learning, planning, and adaptive behaviour. Section 4 examines the relation between environmental and biological complexity via an analysis of uncertainty. We sketch out three scenarios designed to illustrate how sensitivity towards different sorts of uncertainty promotes distinctive kinds of regulatory activity. Finally, Section 5 elaborates some of the key implications of this analysis for the concept of biological cognition. We argue that cognition does not simply coincide with adaptive biological activity (allostatic or otherwise), but inheres rather in the agent's capacity to disengage from the present and entertain counterfactual states of affairs.

2 Homeostasis and the free energy principle

The free energy principle provides a mathematical framework explaining how adaptive organisms come to exist, persist, and thrive – at least for a while – by resisting what Schrödinger described as "the natural tendency of things to go over into disorder" (1992, p. 68). In this section, we sketch a relatively non-technical overview of this perspective, and show how it relates to familiar notions of homeostasis and adaptive behaviour.[1]

---

[1] For broader philosophical discussion of these ideas in the context of predictive processing, see Clark (2016), Hohwy (2013), and Wiese and Metzinger (2017). For more technical explications of the free energy principle and its corollaries, see Bogacz (2017), Buckley et al. (2017), and Friston et al. (2017a).

2.1 Life, formalised: Thermodynamics, attracting sets, and (un)certainty

The free energy principle starts with the simple (but fundamental) premise that organisms must maintain the stability of their internal dynamics in order to survive (Bernard, 1974; Cannon, 1929; Friston, 2012a). This is to say that living systems must act to preserve their structural and functional integrity in the face of environmental perturbation (cf. autopoiesis; Maturana and Varela 1980), thereby resisting the tendency to disorder, dispersal, or *thermodynamic entropy* alluded to by Schrödinger (Friston, 2013; Nicolis and Prigogine, 1977).[2] Reformulated in the language of statistical mechanics: Living systems *live* in virtue of their capacity to keep themselves within a state of (local) thermodynamic equilibrium, while occupying a far from (global) thermodynamic equilibrium region of state-space.[3]

It follows from this postulate that any entity *qua* adaptive biological system can be expected to frequent a relatively small number of *attracting states*; namely those which compose its *attracting set* (Friston, 2012a, 2013). In dynamical systems theoretic terms, this set of states corresponds to a *random dynamical attractor*, the invariant set towards which the system inevitably evolves over time (Crauel and Flandoli, 1994). The existence of this invariant set means that the probability of finding the system in any given state can be summarised by a distribution (technically, an *ergodic density*), which can be interpreted in terms of its *information-theoretic entropy* or *uncertainty* (Shannon, 1948).

The upshot of this picture is that any biotic (random dynamical) system which endures over time must do so in virtue of maintaining a low-entropy distribution over its attracting set (Friston, 2012a; Friston and Ao, 2012). This is tantamount to saying there is a high degree of certainty concerning the state of the system at any given moment in its lifetime, and that such attracting states will correspond to the conditions of the organism's homeostatic integrity. Conversely, there is a low probability of finding the system occupying a state outside of its attracting set, since such states are incompatible with the system's (long-term) existence. It follows that the repertoire of attracting states in which the system is typically located is constitutive of that

---

[2] Technically, living systems appear to violate *fluctuation theorems* that generalise the second law of thermodynamics to nonequilibrium systems (Evans and Searles, 1994, 2002; Seifert, 2012).

[3] See Linson et al. (2018), for a lucid explication of the deep continuities between thermodynamics and the free energy principle. For a more technical exposition, see Sengupta et al. 2013.

agent's *phenotype* (Friston et al., 2009, 2010a), insofar as the phenotype is simply a description of the organism's characteristic (i.e. typically-observed) states.

2.2 Surprise and free energy minimisation

According to this framework, then, homeostasis amounts to the task of keeping the organism within the bounds of its attracting set (or, equivalently, of maintaining a *low conditional entropy* over its internal states). How might biological agents realise this outcome?

To answer this question, we must invoke another information-theoretic term: *surprise* (Shannon, 1948). Surprise (i.e. 'surprisal' or self-information) quantifies the improbability (i.e. negative log-probability) of some outcome. In the present context, the outcome in question refers to some sensory state induced in any part of the system receptive to external perturbation. Obvious realisers of sensory states include the sensory epithelia (e.g., retinal photoreceptor cells), but also extend to ion channel receptors in cell membranes, photosensitive receptors in plants, and so on. These receptive surfaces can be construed as states embedded within a (statistical) boundary or interface (technically, a *Markov blanket*; Pearl 1988) separating (i.e. 'shielding' or 'screening-off') system-internal from system-external conditions (see Friston 2013; Friston and Ao 2012; Hohwy 2017a).[4]

Importantly, the quantity of surprise associated with any given sensory state is not absolute, but depends on the kind of system the organism embodies (i.e. its phenotype or internal configuration; Friston and Stephan 2007). The fish that finds itself on dry land (i.e. well beyond the bounds of its attracting set) experiences a high degree of surprise, and will perish unless something is done (quickly!) to reinstate its usual milieu. Conversely, this very same state will elicit relatively little surprise in land-dwelling creatures. It turns out that *minimising* or *suppressing* the surprise evoked by sensory states – that is, by avoiding surprising states and favouring unsurprising ones – the agent will tend to keep the (conditional) entropy of its states low, since entropy (almost certainly) converges with the long-term time average of surprise (Birkhoff, 1931; Friston and Ao, 2012).

---

[4] Note that complex organisms may be composed of multiple, hierarchically-nested Markov blankets (for recent discussion, see Allen and Friston 2018; Clark 2017; Kirchhoff et al. 2018; Palacios et al. 2017; Ramstead et al. 2018).

In other words, by avoiding surprising interactions with their environment, biological systems keep themselves within the neighbourhood of attracting states that are conducive to their ongoing existence. Indeed, as a random dynamical system that repeatedly revisits its attracting set over time, the agent thereby realises itself as its own random dynamical attractor – and by extension, its own 'existence proof' (Friston 2018; more on which shortly).

There is, however, an important complication in this story: The organism "cannot know whether its sensations are surprising and could not avoid them even if it did" (Friston 2010, p. 128). Surprise is computationally intractable, since its direct evaluation would require the agent to possess exhaustive knowledge of the external dynamics responsible for its sensory experiences (Friston, 2009). This is where the concept of *free energy minimisation* comes in.

Variational free energy is an information-theoretic quantity developed to finesse difficult integration problems in quantum statistical mechanics (Feynman, 1972).[5] In the present context, free energy serves as a proxy for the amount of surprise elicited by sensory inputs (Friston, 2010, 2011). As free energy is a function of the agent's sensory and internal states (i.e. two sources of information available to the agent), and can be minimised to form a tight (upper) bound on sensory surprise, free energy minimisation enables the agent to indirectly evaluate the surprise associated with its sensory states (Friston and Stephan, 2007). Moreover, since the agent is also capable of evaluating how free energy is likely to *change* in response to state transitions (Friston et al., 2012d), it will appear to select (or 'sample') actions that reduce surprise (Friston et al., 2015b).[6] The free energy principle thus implies that biological systems will tend to avoid (or suppress) surprising observations over the long-run, thereby restricting themselves within the neighbourhood of their invariant (attracting) set.

Naturally, this explanation raises yet further questions: How does the agent minimise free energy to a 'tight bound' on surprise? How can simple organisms 'expect'

---

[5] Variational inference techniques are also widely used in machine learning to approximate density functions through optimisation (see Blei et al. 2017).

[6] Of course, just because a system can be *described* as acting in a way that minimises variational free energy (maximises Bayesian model evidence, approximates Bayesian inference, etc.) does not guarantee that it *actually* implements any such computation. The extent to which the free energy principle should be construed as a useful heuristic for describing and predicting adaptive behaviour (a kind of *intentional stance*; Dennett 1987), versus a more substantive ontological claim, remains an open question. That said, recent progress has been made towards casting the free energy principle as a process theory of considerable explanatory ambition (Friston et al., 2017a).

to occupy certain states, or be said to 'prefer' these states over others? In order to address such questions, we first need to elaborate a notion of the agent as a *generative model*.

2.3 Existence implies inference: Agents as generative, self-evidencing models

According to the free energy principle, adaptive biological agents embody a *probabilistic, generative model* of their environment (Calvo and Friston, 2017; Friston, 2011, 2012a; Kirchhoff et al., 2018; Ramstead et al., 2018). As we shall see, this is a rather bold claim that moves us far beyond conventional accounts of homeostatic regulation[7] and their reformulation in the language of statistical mechanics and dynamical systems theory.

Roughly, the system's form and internal configuration are said to parameterise a probabilistic mapping between the agent's sensory states and the external (hidden) causes of such states. This is to say that organisms interact with their eco-niche in ways that distill and recapitulate its causal structure, meaning that biological agents constitute (embody) a statistical model encoding conditional expectations about environmental dynamics (Allen and Friston, 2018; Friston, 2011; Kirchhoff et al., 2018).[8] Indeed, according to the free energy principle, the very existence of the organism over time implies that it must optimise a generative model of the external causes of its sensory flows. This follows from the observation that optimising a model of the hidden dynamics impinging on one's sensory surfaces will give rise to (free-energy minimising) exchanges with the environment, which manifest as adaptive responses to evolving external conditions (Friston et al., 2006; Friston and Stephan, 2007).

Under this account, then, even such simple biological agents as unicellular organisms will 'expect' (abstractly and nonconsciously) to find themselves in certain (unsurprising) states, according to the model they embody. Moreover, such agents will strive

---

[7] Note that we interpret the notion of regulation rather broadly here. For philosophical arguments distinguishing regulation from related concepts such as feedback control and homeostasis, see Bich et al. (2016). On this view, regulatory control consists in a special kind of functional organisation characterised in terms of *second-order control*. This formulation seems broadly in line with our understanding of allostasis (see Section 3).

[8] Note that the organism's morphology and internal organisation impose constraints on the way it models and represents environmental dynamics (e.g., Parr and Friston 2018a) – a point we shall elaborate in Section 4.

to sample (i.e. bring about) those attracting, free energy minimising states they expect to occupy – or risk perishing (Friston et al., 2006; Friston and Stephan, 2007).

In Bayesian terms, this activity of expectation-fulfilment (or maximisation) – where expectations correspond to prior probability distributions parameterised by the agent's internal states – is tantamount to maximising the evidence for the agent's model (and by extension, their own existence; Friston 2010, 2013), a process known as *self-evidencing* (Hohwy, 2016). Hence, under the free energy principle, adaptive biological systems conserve their own integrity through free energy minimising interactions which, over the long-term time average, minimise entropy (i.e. resolve uncertainty) and maximise self-evidence.[9] The process by which they accomplish this feat is *active inference.*

2.4 Active inference: Closing the perception-action loop

The scheme outlined above implies that biological agents conserve their morphology and internal dynamics (and in turn, the generative model these characteristics embody) by acting to offset the dispersive effects of random environmental fluctuations. But why should the agent sustain its model through such adaptive exchanges, rather than allowing its model to change in line with evolving environmental dynamics? As it turns out, the free energy principle supports both of these possibilities: agent and environment are locked in a perpetual cycle of reciprocal influence. This dialectical interplay, which emphasises the inherent *circular causality* at the heart of adaptive behaviour, is formalised under the active inference process theory (Friston et al., 2017a).

Active inference comprises two basic processes that play out at the agent-environment interface: perception and action.[10] Here, perception is construed as the process of changing ('updating') one's internal states in response to external perturbations, and over longer timescales corresponds to learning (i.e. Bayesian updating of time-invariant model parameters; FitzGerald et al. 2015; Friston et al. 2016, 2017a).[11] In

---

[9] See Parr and Friston (2018b) for a mathematical explanation of the (inverse) relationship between variational free energy and model evidence.

[10] While active inference is sometimes narrowly construed as the active or behavioural component of the perception-action loop, the term was originally introduced to characterise the reciprocal interplay between perception and action (e.g., Friston et al. 2009, p. 4). This broader interpretation emphasises the deep continuity of the (Bayesian inferential) processes underwriting perception, learning, planning, and action under the free energy principle (Friston et al., 2017a).

[11] This general understanding of perception need not entail conscious experience of sensations, just as learning can occur through entirely unconscious – and even artificial – mechanisms. Rather, what

other words, perceptual (state) inference describes how the agent updates its representation of environmental dynamics to resolve uncertainty about the hidden causes of its sensory fluctuations. A prevalent neurocomputational implementation of this scheme is *predictive coding* (Elias 1955; Lee and Mumford 2003; Rao and Ballard 1999; Srinivasan et al. 1982; Huang and Rao 2011; Spratling 2017; for some variational free energy treatments, see Barrett and Simmons 2015; Bastos et al. 2012; Friston and Kiebel 2009; Kanai et al. 2015; Pezzulo 2014; Seth et al. 2012; Shipp et al. 2013; Shipp 2016).

Action, on the other hand, involves the activation of effector mechanisms (e.g., motor reflexes, cell migration; Friston et al. 2015a) in order to bring about new sensory states (Adams et al., 2013; Friston et al., 2010a). Different states can be sampled either through actions that directly intervene on the environment (e.g., turning off a bright light), or alter the relationship between the agent's sensory surfaces and external states (e.g., looking away from a bright light). In either case, free energy is affected by the sensory consequences of the agent's actions, where expectations about the modifiability of sensory flows are conditioned on a model of hidden states and their time-evolving trajectories (Friston and Ao, 2012).[12] Active inference thus recalls the cybernetic adage that organisms "control what they *sense*, not... what they *do*" (Powers 1973, p. 355, emphasis in original).

Although we shall have more to say about the role of action under active inference in later sections, these cursory remarks are sufficient to motivate the basic claim that adaptive agents recruit effector systems in order to propel themselves towards the sensory states they expect to inhabit.

Superficially at least, the inferential dynamics underwriting perception and action seem to pull in opposite directions (i.e. *change the model to reflect the world* vs. *change the world to reflect the model*). Under the active inference scheme, however, these two processes are complementary and deeply interwoven. This is because perception can only minimise free energy (or, under certain simplifying assumptions, *prediction error*;

---

is at stake here is the statistical notion of *Bayesian belief*, where probability distributions encode the conditional probability that sensory observation Y was caused by hidden state X.

[12] Technically, actions are physical, real-world states that are not represented within the agent's generative model (Attias, 2003). Rather, the agent infers (fictive) 'control' states that explain the (sensory) consequences of its actions (Friston et al., 2012a,d). Action selection (or decision-making) thus amounts to the optimisation of posterior beliefs about the control states that determine hidden state transitions (Friston et al., 2013, 2015b).

Friston 2009; Friston et al. 2007) to a tight (upper) bound on surprise, whereas action suppresses surprise by invoking new sensory states that conform to (expectations prescribed by) the agent's phenotype. Consequently, perception serves to optimise the agent's model of environmental conditions, such that the agent has adequate information to choose actions that engender low sensory entropy (Friston et al., 2010a).[13]

Although perceptual inference might seem to imply that agents ought to adapt their internal organisation to reflect environmental fluctuations as accurately as possible, unrestricted acquiescence to such dynamics would result in a precarious (and in many cases, rather brief) existence. Rather, the exigencies of homeostatic control dictate that biological systems preserve the *conditional independence* of their internal and external states (Ramstead et al., 2018). This is to say that the biological agent must maintain a boundary (i.e. Markov blanket) that separates (and insulates) its internal dynamics from external conditions.[14] Consequently, the free energy minimising agent must exploit inferences about the state of the world beyond its Markov blanket in order to act in ways that keep it within the neighbourhood of its attracting states (Friston, 2013).

The agent's capacity to maintain the integrity of its Markov blanket is aided by prior beliefs about the sorts of conditions it expects to encounter. Many such expectations are directly functional to homeostasis (Pezzulo et al., 2015), having been shaped and refined through generations of natural selection (Allen and Friston, 2018; de Vries and Friston, 2017; Friston, 2010). Pushing this logic one step further, we can say that the agent embodies a deeply-engrained expectation to survive (i.e. to remain within the confines of its attracting set – and thus to maintain its homeostatic integrity over time); this is simply the expectation to minimise average surprise over the long-run (Allen and Tsakiris, 2018; Seth, 2015). This remark highlights the point that not all beliefs are equally amenable to model updating. Rather, certain strongly-held or *high-precision* beliefs (e.g., those pertaining to homeostatic stability) will be stubbornly defended through actions that seek to substitute conflicting sensory evidence with

---

[13] Although one might be tempted to subordinate perceptual inference to free energy minimising action, we interpret perception and action as mutually dependent moments within a unified dynamical loop (cf. the perception-action cycle; Fuster 2001, 2004). Ultimately, *both* modes of active inference are in the service of uncertainty reduction: Percepts without actions are idle; actions without percepts are blind.

[14] Formally speaking, the sensory and active states that compose the Markov blanket render the probability distributions over internal and external states statistically independent of one another (see Pearl 1988). In other words, internal and external states provide no additional information about one another once the Markov blanket's active and sensory states are known.

input that conforms more closely to prior expectations (Yon et al., 2019).

In sum, perception and action work in concert to achieve free energy minimisation, ensuring that the biological system maintains itself in an invariant relationship with its environment over time. Critically, this formulation explains how apparently teleological or purposive behaviours emerge as a consequence of free energy minimising sensory sampling, without resorting to additional concepts such as 'value' or 'reward' (Friston et al., 2009, 2010a). Rather, value and reward simply fall out of the active inference process, as what is inherently valuable or rewarding for any particular organism is prescribed by the attracting states that compose its phenotype (i.e. those states the agent expects itself to occupy; Friston and Ao 2012). Simply put, unsurprising (i.e. expected) states are valuable; hence, minimising free energy corresponds to maximising value (Friston et al., 2012a).[15]

3 Beyond homeostasis: Allostasis and hierarchical generative models

The free energy principle is founded on the premise that biological systems maintain their homeostatic (i.e. local thermodynamic) equilibrium in the face of random environmental perturbations. However, the question of how adaptive organisms secure the homeostatic integrity of their internal milieu had been largely neglected until relatively recently. A growing number of researchers are now applying the theoretical resources of predictive coding and active inference to explain how nervous systems monitor internal bodily sensations (i.e. perceptual inference in the *interoceptive* domain; Barrett and Simmons 2015; Seth 2013; Pezzulo 2014) and modulate physiological conditions (for recent reviews, see Khalsa et al. 2018; Owens et al. 2018; Quadt et al. 2018).

An important conceptual development within this line of work has been the move beyond traditional notions of homeostatic stability to more modern accounts of allostatic variability. The concept of allostasis ("stability through change") was first introduced by Sterling and Eyer (1988), who criticised conventional homeostatic control theory as overly restrictive and reactive in character.[16] By contrast, allostasis was

---

[15] Note that value here is not equivalent to expected utility, but rather a composite of utility (*extrinsic value*) and information gain (*epistemic value*; see Friston et al. 2015b; Schwartenbeck et al. 2015).

[16] Although we focus here on allostasis, numerous other concepts emphasising the dynamic nature

doi:10.20944/preprints201911.0083.v1

intended to replace setpoint defence with a more flexible scheme of parameter variation, and to supersede local feedback loops with centrally co-ordinated feedforward mechanisms (e.g., *central command*; Dampney 2016; Goodwin et al. 1972; Krogh and Lindhard 1913). Allostasis was thus posited to account for a wide variety of anticipatory physiological activity that could not be explained in terms of closed-loop control.

Despite controversy over the theoretical merits and conceptual scope of allostasis (see Corcoran and Hohwy 2018, for a recent overview), there is ample evidence that biological regulation consists in both anticipatory and reactive modes of compensation (see for e.g., Burdakov 2019; Ramsay and Woods 2016; Schulkin and Sterling 2019).[17] These complementary mechanisms are easily accommodated within the active inference framework, mapping neatly onto the hierarchically-stratified models posited under the free energy principle (Friston, 2008). Moreover, we believe that mature versions of allostatic theory are enriched and invigorated by active inference, insofar as the latter furnishes precisely the kind of inferential machinery required to underwrite effective forms of prospective control across various timescales (Corcoran and Hohwy, 2018; Kiebel et al., 2008; Friston et al., 2017d; Pezzulo et al., 2018).

The remainder of this section briefly outlines two recent attempts to integrate homeostatic and allostatic mechanisms within the broader scheme of active inference. Although these perspectives assume a rather complex, neurally-implemented control architecture, we shall argue in Section 4 that the basic principles underwriting such schemes can be generalised to much simpler biological systems with relative ease.

3.1 Allostasis under active inference

Stephan and colleagues (2016, see also Petzschner et al. 2017) developed an active inference-based account of allostasis that maps interoception and physiological regulation onto a three-layer neural hierarchy. At the lowest level of this hierarchy are *homeostatic reflex arcs*, which operate much like classical feedback loops (i.e. deviation of an essential variable beyond certain limits elicits an error signal, which in

---

of biological regulation have been proposed in an effort to extend (or transcend) classical notions of homeostatic setpoint control (see for e.g., Bauman 2000; Berntson and Cacioppo 2000, and references therein).

[17] Indeed, evidence of anticipatory physiological regulation antedates Walter B. Cannon's influential work – Ivan Pavlov's (1902) Nobel prize-winning research on the digestive system demonstrated that gastric and pancreatic enzymes are secreted *before* nutrient ingestion (see Smith 2000; Teff 2011).

turn triggers a countervailing effector response; see Ashby 1956, Ch. 12; Wiener 1961, Ch. 4). Critically, however, the range of states an essential variable may occupy is prescribed by intermediate-level *allostatic circuits*. This formulation thus recasts essential variable setpoints as (probabilistic) prior expectations (or equivalently, top-down model-based predictions) about the likely states of interoceptors (cf. Penny and Stephan 2014), with deviations from expected states provoking interoceptive prediction error.[18]

Two important features of this account are that (1) prior expectations about essential variables encode a distribution over states (rather than a singular ideal reference value), and that (2) the sufficient statistics which specify this distribution – its mean and precision (inverse variance) – are free to vary (cf. Ainley et al. 2016). On this view, such classic allostatic phenomena as diurnal patterns of body temperature (Kräuchi and Wirz-Justice 1994) and blood pressure variation (Degaute et al. 1991) emerge as a consequence of the cyclical modulation of the priors over these physiological states (cf. Sterling 2004, 2012). Likewise, phasic increases or decreases in the stability of such variables correspond to periodic shifts between more- or less-precise distributions, respectively.[19]

Subordinating homeostatic reflex arcs to allostatic circuits transforms traditional conceptions of physiological control as setpoint defence into a far more dynamic and context-sensitive process. Access to perceptual and cognitive representations (e.g., via the anterior insular and cingulate cortices; Barrett and Simmons 2015; Craig 2009; Gu et al. 2013; Menon and Uddin 2010; Paulus and Stein 2006) enables allostatic circuitry to harness multiple streams of information such that homeostatic parameters may be deftly altered in preparation for expected environmental changes (Ginty et al. 2017; Peters et al. 2017). Not only does this arrangement enable the system to anticipate periodic nonstationarities in essential variable dynamics (such as the circadian oscillations in body temperature and blood pressure mentioned above), it also confers potentially vital adaptive advantages under unexpected and uncertain conditions.

---

[18] This formulation is congruent with contemporary efforts to finesse traditional notions of setpoint rigidity with more dynamic accounts of homeostatic control (e.g., Cabanac 2006; Ramsay and Woods 2014; cf. Ashby 1940). It also seems more felicitous to Cannon's original conception of homeostatic control (see for e.g., Cannon 1939, p. 39).

[19] Note that certain physiological expectations are likely to be held with greater precision (e.g., core temperature and blood pH) – and thus restricted to a narrower range of attracting states – than others (e.g., blood pressure and heart rate; see Allen and Tsakiris 2018; Seth and Friston 2016; Yon et al. 2019).

As a brief illustration, consider the case of an animal that detects the presence of a nearby predator. Registering its perilous situation, the brain triggers a cascade of autonomic activity – the 'fight-or-flight' response famously characterised by Cannon (1914; 1915). On Stephan and colleagues' (2016) account, these rapid physiological alterations are mediated via the allostatic enslavement of homeostatic reflex loops. This generative model-based scheme explains why physiological parameters should change so dramatically in the *absence* of any immediate homeostatic disturbance: Predictions (or 'forecasts'; Petzschner et al. 2017) about the likely evolution of external conditions mandate the adoption of atypical, metabolically expensive states in preparation for evasive action (cf. Requin et al. 1991).

Notice that the physiological states realised via allostatic modulation of homeostatic loops might themselves constitute surprising departures from the organism's typically-expected states. Since these deviations cannot be locally resolved on account of the higher-order imperative to mobilise metabolic resources for impending action, interoceptive prediction error propagates up the neural hierarchy, possibly manifesting as the suite of sensations associated with acute stress (Peters et al., 2017). Such prediction error is tolerated to the extent that these emergency measures are expected to expedite the reinstatement of a more hospitable environment (namely, one in which there is no immediate threat of predation). In other words, allostatic regimes of interoceptive active inference are functional to the agent's deeply-held expectation to survive, insofar as they serve to minimise uncertainty and maximise self-evidence *over the long-run.*[20]

Stephan and colleagues (2016) crown their hierarchical framework with a *metacognitive* layer that monitors the efficacy of one's control systems. This processing level is posited to explain the emergence of higher-order beliefs about one's ability to adaptively respond to homeostatic perturbation. Persistent failure to suppress interoceptive surprise – either as a consequence of harbouring inaccurate allostatic expectations, or one's inability to realise free energy minimising actions – results in a state of *dyshomeostasis* (cf. *allostatic load*; McEwen and Stellar 1993; Peters et al. 2017), ex-

---

[20] One might protest that all we have done here is pivot from one sort of reactive homeostatic mechanism to another; albeit, one involving responses to an external (rather than internal) threat. Nevertheless, we consider this simple scenario as exemplary of the fundamental principle of allostatic regulation; namely, the modulation of physiological states in anticipation of future conditions, and in the absence of any immediate homeostatic perturbation. This example can easily be extended to capture a rich assortment of allostatic dynamics that play out across increasing levels of abstraction and spatiotemporal scale.

perience of which may erode confidence in one's capacity for self-regulation. Stephan and colleagues (2016) speculate that the affective and intentional states engendered by chronic dyshomeostasis play a role in the development of major depressive disorder (cf. Badcock et al. 2017; Barrett et al. 2016; Seth and Friston 2016). Although such psychopathological implications are beyond the scope of this paper, the basic idea that the brain's homeostatic/allostatic architecture is reciprocally coupled with higher-order inferential processing will be explored further in Section 4.

In sum, the hierarchical regulatory scheme proposed by Stephan and colleagues (2016) provides a promising formal description of the inferential loops underwriting both reactive (homeostatic) and prospective (allostatic) modes of biological regulation, and their interaction with higher-order beliefs. This framework accommodates a rich variety of allostatic phenomena spanning multiple timescales; ranging from deeply-entrenched, slowly-unfolding regularities (e.g., circadian and circannual rhythms) to highly unpredictable, transient events (e.g., predator-prey encounters), and everything in between (e.g., meal consumption; Morville et al. 2018; Teff 2011). However, this account remains somewhat preliminary; more needs to be done to flesh out the ways in which allostatic circuits leverage prior experience (i.e. learning and memory), information about evolving state trajectories, and complex behavioural repertoires to optimise homeostatic control. We pursue some of these considerations next.

3.2 Broadening the inferential horizon: Preferences, policies, and plans

In order to fully appreciate the adaptive potential of allostatic regulation under active inference we need to address the emergence of free energy minimising *policies*. Policies are prior beliefs about the sequence of actions (or more precisely, control states; see Footnote 12) required to realise some preferred (i.e. expected, self-evidencing, and thus *valuable*) outcome (Attias, 2003; Friston et al., 2012a, 2013; Pezzulo et al., 2018).

A key question here concerns how the agent translates inferences on time-varying internal states (e.g., declining blood glucose concentration) into complex behavioural strategies (e.g., preparing a meal) intended to ameliorate – or better yet, prevent – homeostatic disturbance (Keramati and Gutkin, 2014). Finessing this slightly, one might ask why it is that agents sometimes prefer to engage in relatively complex courses of action, thereby ignoring simpler or more immediately available alternatives.[21]

Pezzulo and colleagues (2015) offer an account of allostasis that seeks to explain the gamut of behavioural control schemes acquired via associative learning from a unified active inference perspective.[22] Specifically, this account grounds the emergence of progressively more flexible and sophisticated patterns of adaptive behaviour on evolutionarily primitive control architectures (e.g., low-level circuitry akin to Stephan and colleagues' (2016) homeostatic reflex arc). From a broader ethological perspective, this scheme implies a deep continuity between the homeostatic loops underpinning simple, stereotypical response behaviour on the one hand, and the complex processes supporting goal-directed decision-making and planning on the other.

According to this view, all associative learning-based control schemes fall out of the same entropy-reducing dynamics prescribed by the free energy principle. What distinguishes these schemes under the active inference framework is their place in the model hierarchy: While rudimentary adaptive behaviours (e.g., approach/avoidance reflexes) are availed by 'shallow' architectures, more sophisticated modes of control require greater degrees of hierarchical depth. Goal-directed actions require generative models that are capable of representing the prospective evolution of hidden states over sufficiently long intervals (cf. Botvinick and Toussaint 2012; Penny et al. 2013; Solway and Botvinick 2012), while simultaneously predicting how these projected trajectories are likely to impact upon the organism. On this account, activity at higher (or deeper) hierarchical layers (e.g., prefrontal cortical networks) contextualises that of more primitive control schemes operating at lower levels of the hierarchy (see also Pezzulo and Cisek 2016; Pezzulo et al. 2018). This means that higher-level inferences about distal or remote states (and the policies most likely to realise them) inform lower-level mechanisms governing action over shorter timescales (see also Attias 2003; Badre 2008;

---

[21] An example of this might be deciding to purchase ingredients from a local supermarket and returning home to cook a meal, rather than ordering a meal from a neighbouring fast food restaurant. In both cases, the underlying homeostatic motivation driving behaviour (i.e. increasing prediction error manifesting as intensifying hunger) is identical; the interesting question is why one does not always opt for the policy that is most likely to resolve prediction error (hunger) most rapidly. Selection of the *Purchase-ingredients-and-cook* policy, which postpones the resolution of interoceptive prediction errors (and thus engenders greater free energy in the short-term), appears on first blush to contradict the free energy principle. Such choices can however be explained by recourse to the agent's superordinate expectation to minimise (expected) free energy over longer timescales (e.g., prior beliefs about the health, financial, and/or social benefits associated with domestic meal preparation; cf. Friston et al. 2015b; Pezzulo 2017; Pezzulo et al. 2018). Note that the appeal to expected free energy was also implicit in the predator example of the previous section, insofar as transient increases in homeostatic prediction error were tolerated in order to avoid a much more surprising fate – being eaten!

[22] See Moore (2004) for a thoroughgoing review of associative learning mechanisms.

Friston et al. 2016; Kaplan and Friston 2018; Pezzulo et al. 2018).

A distinctive feature of Pezzulo and colleagues' (2015) scheme is the crucial role played by the (cross- or multimodal) integration of interoceptive, proprioceptive, and exteroceptive information over time. This is to say that the emergence of nervous systems which enable their owners to envisage and pursue certain future states at the expense of others depends upon the (allostatic) capacity to track and anticipate co-evolving internal/sensory and external/active state trajectories.[23] In short, Pezzulo and colleagues (2015) posit that hierarchical generative models harness prior experience to map sensorimotor events to interoceptive fluctuations. This mapping enables the agent to learn how their interoceptive/affective states are likely to change both endogenously (e.g., *I am likely to become irritable if I forgo my morning coffee*), and in the context of external conditions (e.g., *I am likely to dehydrate if I exercise in this heat without consuming fluids*).[24]

With this (hierarchical) inferential architecture in place, it is relatively easy to see how allostatic policies take root. As alluded to above, interoceptive/homeostatic dynamics often exhibit (quasi)periodic cycles, thus facilitating the modelling and prediction of time-evolving changes in internal sensory states. Given a model of how interoceptive states typically oscillate, the agent learns how particular external perturbations (including those caused by its own actions) modulate this trajectory (cf. Allen and Tsakiris 2018). As the agent accrues experience, it progressively refines its model of the contingent relations that obtain between sensorimotor occurrences and physiological fluctuations, engendering the ability to extrapolate from sensations experienced in the past and present to those expected in the future (Friston et al., 2017a). This capacity is not only crucial for finessing the fundamental control problems posed by homeostasis (i.e. inferring the optimal policy for securing future sur-

---

[23] More precisely, this capacity depends on the ability to infer the expected free energy of the outcomes associated with various potential state trajectories, as well as the expected likelihood of outcomes under each policy (see Friston et al. 2017a,c; Parr and Friston 2017, 2018b). Such processing might be facilitated by the co-ordination of exteroceptive sampling and motor planning with periodic regimes of autonomic/interoceptive active inference (Corcoran et al., 2018).

[24] We emphasise again that the conscious, reflective character of these intuitive examples should not detract from the idea that the *possibility* of such experiences is underwritten by more basic, unconscious allostatic mechanisms. For example, the growth onset of a horse's winter coat is not assumed to represent a strategic decision on the part of the horse, but rather a physiological response to seasonal changes in photoperiod. Similarly, a rabbit might schedule her foraging bouts to balance energy gain against predation risk, even though she might not be capable of representing and evaluating these concerns explicitly (this trade-off may, for instance, be implicitly encoded within the animal's circadian rhythm – see Section 4.3).

vival and reproductive success), but also for its vital contribution in establishing the agent's understanding of itself *qua* autonomous agent (cf. Fotopoulou and Tsakiris 2017; Friston 2017). It is a relatively small step from here to the emergence of goal-directed behaviours that are ostensibly independent of (i.e detached or decoupled from) the immediate exigencies of biological regulation (e.g., purchasing food when one is not hungry; see Pezzulo and Castelfranchi 2009; Pezzulo 2017).

3.3 Interim summary

In this section, we have presented two closely-related computational perspectives on biological regulation that cast homeostasis and allostasis within the broader scheme of active inference. We believe these accounts can be productively synthesised into a comprehensive framework that explains the emergence of increasingly versatile, context-sensitive, and temporally-extended forms of allostatic regulation. This framework provides a formal account of biological regulation that eschews the conceptual limitations of setpoint invariance (see Cabanac 2006; Ramsay and Woods 2014), unifies habitual ('model-free') and goal-directed ('model-based') behaviour (Dolan and Dayan, 2013) under a single hierarchical architecture (see FitzGerald et al. 2014; Pezzulo et al. 2016), and converges with neurophysiologically-informed perspectives on mind-body integration (e.g., Critchley and Harrison 2013; Smith et al. 2017). We have also introduced the important notion of policy selection, which explains how adaptive behaviour emerges through (active) inference of beliefs about the future (cf. 'planning as inference'; Attias 2003; Botvinick and Toussaint 2012; Solway and Botvinick 2012).

From a broader perspective, the capacity of higher model levels to track the evolution of increasingly distal, temporally-extended, and abstract hidden dynamics, and to infer the likely consequences of such dynamics for the agent's own integrity and well-being, provides a compelling explanation of how allostatic control schemes could have established themselves over ontogenetic and phylogenetic timescales. Not only does this perspective provide a principled account of how allostatic mechanisms should 'know' when to initiate adaptive compensations in the absence of physiological disturbance (i.e. how the body 'acquires its wisdom'; Dworkin 1993), the embedding of such processes within an overarching hierarchical model also explains how agents are able to effectively arbitrate and trade-off multiple competing demands (a core feature of many allostatic frameworks; e.g., Sanchez-Fibla et al. 2010; Sterling 2012; Schulkin and Sterling 2019; Verschure et al. 2014).[25] In the next section of this paper, we consider *why* such allostatic regimes should have evolved.

4 Biological regulation in an uncertain world

We have argued that adaptive biological activity is underwritten by active inference, where more sophisticated (predictive or prospective) forms of biological regulation (allostasis) are supported by hierarchical generative models that extract and exploit patterned regularities in internal and external conditions. In this section, we take a closer look at how the functional organisation of the inferential architecture constrains the organism's capacity to represent time-evolving state trajectories, and the impact this has upon its ability to adapt to environmental mutability.

Our analysis draws inspiration from Godfrey-Smith's influential *environmental complexity thesis* (1996), which casts cognition as an adaptation to certain complex (i.e. heterogeneous or variable) properties of the organism's eco-niche. On this view, cognition evolved to mitigate or 'neutralise' environmental complexity by means of *behavioural complexity* – "the ability to do a lot of different things, in different conditions" (Godfrey-Smith 1996, p. 26).[26]

The concept of complexity at the core of Godfrey-Smith's analysis is deliberately broad and abstract. Environments may comprise manifold dimensions of complexity, many of which may be of no ecological relevance to their inhabitants. Patterns of variation only become biologically salient once the capacity to track and co-ordinate with them confers a selective advantage (i.e. when sensitivity to environmental variation helps the organism to solve problems – or exploit opportunities – that bear on its fitness; Godfrey-Smith 2002). Much like the notion of surprise (conditional entropy)

---

[25] See Morville et al. (2018) for discussion of the nontrivial challenges posed by high-dimensional homeostatic needs in uncertain environments. The ability to reliably navigate such complex demands speaks also to the notion of *competence* in artificial intelligence research (see Miracchi 2019).

[26] This gloss on the environmental complexity thesis is reminiscent of W. Ross Ashby's *law of requisite variety* (Ashby 1956, 1958; cf. Conant and Ashby 1970), and is clearly in line with recent neuroscientific interest in the brain's teleonomic function as a sophisticated biological regulator (for discussion, see Williams and Colling 2018). Although Godfrey-Smith (1996, pp. 76–79) briefly remarks upon the connection between cybernetic accounts of homeostatic control and cognitive function, he rejects their strong continuity on the grounds that cognition can sustain biological viability through actions that circumvent homeostatic mechanisms. We concur that non-trivial definitions of homeostasis and cognition invoke concepts that are distinct from one another, and argue below that this distinction can be cashed out in terms of their constitutive inferential architectures.

introduced in Section 2 then, the implications of environmental complexity for any given organism are determined by the latter's constitution and relation to its niche.

In what follows, we analyse the connection between environmental and behavioural complexity as mediated by increasingly elaborate schemes of active inference. Following Godfrey-Smith's observation that complexity can be cast as "disorder, in the sense of uncertainty" (Godfrey-Smith 1996, p. 24; see also pp. 153–154), we ask whether sensitivity to particular kinds of uncertainty can be mapped onto particular forms of inferential processing, and the adaptive functional profiles they avail.

4.1 Varieties of uncertainty

Uncertainty can be decomposed in various ways depending on the domain of interest (e.g., Bland and Schaefer 2012; Bradley and Drechsler 2014; Kozyreva and Hertwig 2019). Here, we focus on three kinds of uncertainty that can be defined in terms of their statistical properties: *risk*, *ambiguity*, and *volatility*.[27]

Under active inference, risk and ambiguity (sometimes collectively termed *expected uncertainty*; Yu and Dayan 2003, 2005) are formalised as components of the expected free energy of a policy (Friston et al., 2017a,b,d). Risk is defined by the relative entropy (Kullback-Leibler divergence) between preferred outcomes and those expected under a given policy. Ambiguity is defined by the (Shannon) entropy over the probabilistic mapping between observations and hidden states. More intuitively, risk can be understood as the probability of gaining some reward (e.g., finding a cookie) as a consequence of some action (e.g., reaching into a cookie jar). Ambiguity pertains to the fact that a given outcome might have come about in various different ways (e.g., the cookie in my hand might have been given to me, stolen from the jar, etc.).

Heuristically, resolving ambiguity will increase one's confidence about the process responsible for generating observations, enabling one to evaluate the risk (i.e. expected cost) associated with alternative possible actions (e.g., staking a bet on the colour of a ball drawn from an urn). This understanding of risk and ambiguity is broadly consistent with descriptions in economics (e.g., Camerer and Weber 1992; Ellsberg 1961; Kahneman and Tversky 1979; Knight 1921) and neuroscience (e.g., Daw et al.

---

[27] Risk and ambiguity are also known as irreducible uncertainty and (parameter) estimation uncertainty, respectively (de Berker et al., 2016; Payzan-LeNestour and Bossaerts, 2011).

2005; Hsu et al. 2005; Huettel et al. 2006; Levy et al. 2010; Payzan-LeNestour and Bossaerts 2011; Preuschoff et al. 2008; for a review, see Bach and Dolan 2012).

Volatility (or *unexpected uncertainty*; Yu and Dayan 2003, 2005) is a higher-order form of uncertainty that refers to changes in the probabilistic mapping between observations and hidden states.[28] While ambiguity can be conceived in terms of (inverse) sensory precision (i.e. confidence about the hidden causes of sensory states), volatility can be construed as the (inverse) precision over transition probabilities (i.e. confidence about the way hidden states evolve over time; Parr and Friston 2017; Parr et al. 2019; Sales et al. 2019; Vincent et al. 2019). This formulation suggests that volatile environments will tend to generate more surprising outcomes than stable environments, insofar as their states are apt to change in ways that are difficult to anticipate.

In the remainder of this section, we consider how the exigencies of biological regulation under conditions of uncertainty may have promoted the evolution of increasingly complex inferential architectures. We shall see how representational access to higher orders of uncertainty enables the agent to regulate policy selection in a progressively more flexible and autonomous fashion. To illustrate the qualitatively distinct control regimes that derive from increasingly elaborate forms of hierarchical organisation, we consider three instantiations of a schematic active inference agent. We model this ‘creature’ on simplified aspects of *Escherichia coli* (*E. coli*) bacteria to emphasise the generality of these schemes beyond neurally-implemented control systems.

Our *E. coli*-like creature is a unicellular organism equipped with a cell membrane (i.e. a Markov blanket separating internal from external states), a metabolic pathway (i.e. an autopoietic network that harnesses thermodynamic flows to realise and replenish the organism’s constitutive components), and a sensorimotor pathway; but at the outset nothing approximating a nervous system (actual *E. coli* is of course much more complicated than this). Cellular metabolism depends on the agent’s ability to absorb sufficient amounts of nutrient (e.g., glucose) from its immediate environment. However, the distribution of nutrient varies across the environment, meaning the agent must seek out nutrient-rich patches in order to survive. Like real *E. coli*, our creature

---

[28] Terminology is somewhat less consistent here. Bland and Schaefer (2012) distinguish volatility from unexpected uncertainty, where the latter refers to *rare* (as opposed to frequent) changes in state–outcome contingencies. Mathys et al. (2014), by contrast, decompose unexpected uncertainty into *tonic* and *phasic* volatility. In line with Behrens et al. (2007) and Payzan-LeNestour and Bossaerts (2011), we construe volatility/unexpected uncertainty as a single dimension varying from stable/reliable contingencies (low volatility) to changeable/unpredictable contingencies (high volatility).

attempts to realise this goal by alternating between two chemotactic policies: *Run* (i.e. swim along the present course) vs. *Tumble* (i.e. randomly reorient to a new course, commence swimming; see Figure 1).

4.2 Model 1: Minimal active inference

First, let us consider a 'basic' example of homeostatic conservation through 'minimal' active inference.[29] Our simplified *E. coli*-like creature embodies a model that encodes an expectation to inhabit a nutrient-rich milieu. Variation in the environment's chemical profile means that this expectation is not always satisfied – sometimes the agent finds itself in regions where chemical attractant is relatively scarce. Crucially, however, the organism can infer its progress along the nutrient gradient through periodic sampling of its chemosensory states, and acts on this information such that it tends to swim up the gradient over time.[30]

This rudimentary sensorimotor system affords the agent a very primitive picture of the world, picking out a single, salient dimension of environmental complexity (i.e. attractant rate of change). The capacity to estimate or infer this property implies a model that prescribes a fixed expectation about the kind of milieu the agent will inhabit, while also admitting some degree of uncertainty as to whether this expectation will be satisfied at any given moment. The task of the agent then is to accumulate evidence in favour of its model by sampling its policies in such a way that it ascends the nutrient gradient, thereby realising its expected sensory states and resolving uncertainty.

Although severely limited in terms of the perceptual or representational capacities at its disposal, this need not imply suboptimality *per se*. Consider the case in which various kinds of attractant are compatible with the organism's chemoreceptors. The

---

[29] See Baltieri and Buckley (2017) and McGregor et al. (2015) for alternative 'minimal' active inference formulations.

[30] In fact, real *E. coli* realise a similar 'adaptive gradient climbing' strategy by integrating chemosensory information about the ambient chemical environment over time, and modulating the probability of tumbling as a function of attractant rate of change (Berg and Brown, 1972; Falke et al., 1997). More recent work has indicated that such chemotactic activity approximates optimal Kalman filtering (Andrews et al., 2006), where hidden states are estimated on the basis of prior and present observations weighted by their uncertainty (Kalman 1960; Kalman and Bucy 1961; see Grush 2004, for discussion). As Kalman filtering constitutes a special case of Bayesian filtering (one that is equivalent to predictive coding; Bastos et al. 2012; Friston et al. 2018, 2010b), chemotaxis can be cast as a gradient descent on variational free energy. Notice that our model is deliberately simpler than this scheme, since sensory prediction errors are not modulated by an uncertainty (precision) parameter.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 November 2019 doi:10.20944/preprints201911.0083.v1

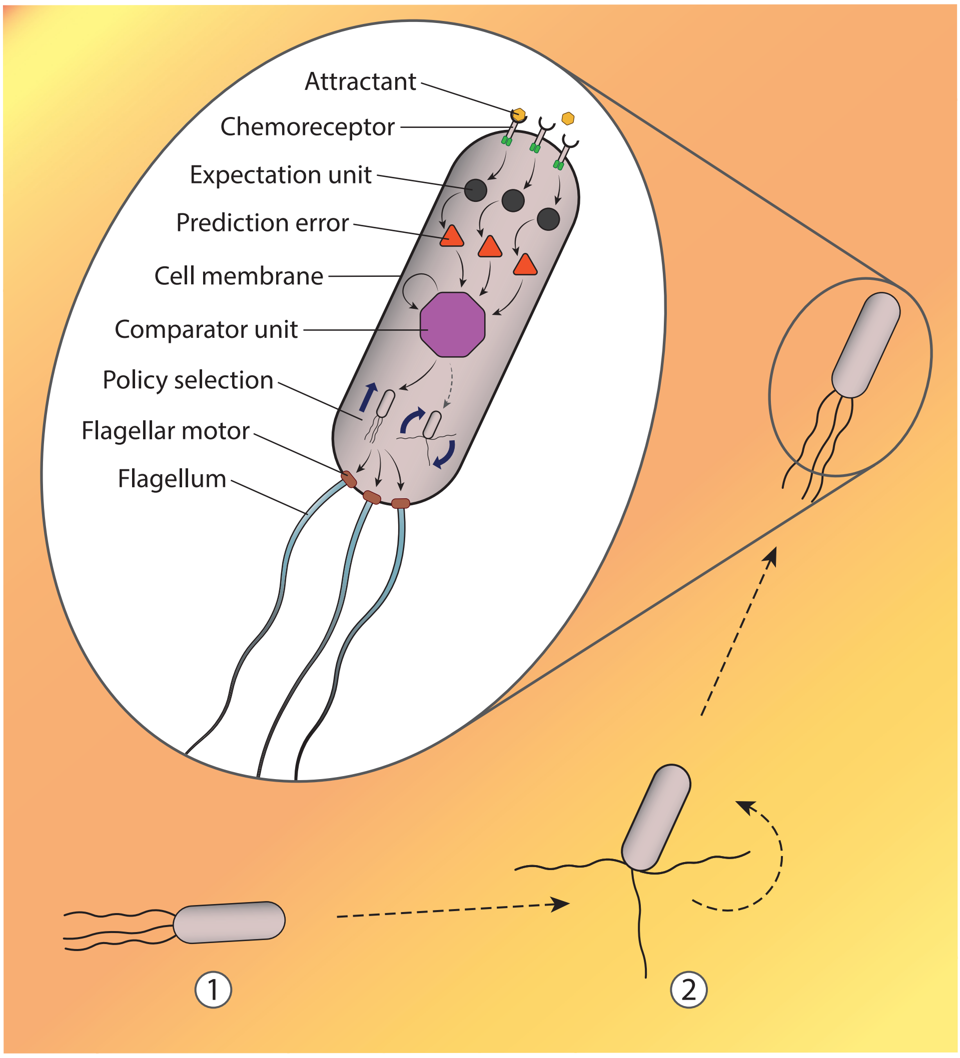

Figure 1: **A simple active inference model of bacterial chemotaxis.** This figure depicts a simple active inference agent that must sample from its sensory states in order to infer the best course of (chemotactic) action. Embedded within its cell membrane are a set of transmembrane proteins that 'activate' when their binding sites encounter attractant molecules. Since the organism expects its chemoreceptors to be occupied by attractant molecules, sensory states indicating a lack of attractant evoke prediction errors. These signals are projected to the agent's motor control network, where they are summed and compared to the summed prediction error accumulated on the previous wave of sensory input. If prediction error is reduced relative to the preceding cycle of perceptual inference, this provides evidence that the agent is ascending the nutrient gradient, and thus evidence favouring the *Run* policy (1). Alternatively, a relative increase in prediction error constitutes evidence of gradient descent, compelling the agent to sample its *Tumble* policy (2). Note that the organism's metabolic system has been omitted from this schematic.

agent cannot discriminate amongst these chemical substances; all it can do is infer the presence (or absence) of 'nutrient' at its various receptor sites. Assuming all forms of chemical attractant are equally nutritious (i.e. equally 'preferable' or 'valuable' given the agent's phenotype), this source of environmental heterogeneity turns out to be entirely irrelevant to the system's ongoing viability. Consequently, the extra structural and functional complexity required to distinguish these substances would afford the organism no adaptive benefit – indeed, the additional metabolic costs incurred by such apparatus might constitute a hindrance.[31]

Our *E. coli*-like creature thus trades in a rather coarse representational currency (i.e. presence/absence of nutrient), enabling it to minimise the costs associated with unwarranted degrees of organisational complexity. This is an example of optimising the trade-off between model accuracy and complexity (FitzGerald et al., 2014; Hobson and Friston, 2012; Moran et al., 2014), where the simplest model to satisfactorily explain observed data (sensory input) defeats more complex alternatives (or where natural selection favours the simplest model that satisfices for survival and reproductive success; Campbell 2016; Friston 2018). This also explains why some creatures might have evolved *simpler* phenotypes from more complicated progenitors – natural selection 'rewards efficiency' over the long-run (McCoy, 1977).

This caveat notwithstanding, however, there are still a great many aspects of the environment that the minimal active inference agent fails to model *despite* their potential bearing on its wellbeing. One such omission is the system's incapacity to represent the evolution of its states over multiple sensory samples. This limitation is significant, since it prevents the organism from discerning patterns of variation over time, which in turn renders it overly sensitive to minor fluctuations in prediction error. For instance, the organism might trigger its *Tumble* policy at the first sign of gradient descent, even though this decrement might stem from a trivial divergence in the quantity of attractant detected across sensory samples. Unable to contextualise incoming sensory information with respect to the broader trajectory of its sensory flows, the agent risks tumbling out of a nutrient-rich stream due to innocuous or transient instability

---

[31] The story changes if the organism's receptors are compatible with molecules it cannot metabolise, or that afford low nutritional value (assuming such molecules are prevalent enough to significantly interfere with chemotaxis). See Sterelny (2003, pp. 20–26) for discussion of the challenges posed by 'informationally translucent environments' that confront organisms with ambiguous (or misleading) cues. Environmental translucence calls for greater model complexity; e.g., the capacity to integrate information harvested across multiple sensory channels (cf. *robust tracking*; Sterelny 2003, pp. 27–29).

of the gradient, or due to the random error introduced by inherently noisy signalling pathways.

Relatedly, the agent's inability to retain and integrate over past experiences precludes the construction of map-like representations of previously-explored territory. The organism thus loses valuable information about the various conditions encountered on previous foraging runs – information that a more sophisticated creature could potentially exploit in order to extrapolate the most promising prospects for future forays. It also lacks the necessary model parameters to track various distal properties that modulate or covary with the distribution of attractant (e.g., weather conditions, conspecifics, etc.). The agent is thus unable to exploit the patterned regularities that obtain between proximal and distal hidden states, and that afford predictive cues about the likely consequences of pursuing a particular policy (cf. fish species whose swim policies are informed by predictions about distal feeding conditions and temperature gradients; Fernö et al. 1998; Neill 1979). Unable to 'see' beyond the present state of its sensory interface with the world, the organism has no option but to tumble randomly towards a new, unknown territory each time prediction error accrues.

In sum, the agent we have described here embodies a very simple active inference scheme; one which supports adaptive responses to an ecologically-relevant dimension of environmental complexity. While the agent does not always succeed in inferring the best chemotactic policy in a given situation, its strategy of alternating between active states in accordance with local nutrient conditions is cheap and efficient, and tends to prevent it from drifting too far beyond its attracting set. But the severe epistemic constraints enforced by the agent's extremely narrow representational repertoire – both in the sense of its highly constricted spatiotemporal horizon, and the poverty of its content – render this organism a creature of hazard. Unable to profit from past experience or future beliefs, it is locked in a perpetual present. This creature is thus thoroughly homeostatic in nature, activating its effector mechanisms whenever error signals indicate deviation beyond setpoint bounds.

The epistemic constraints mentioned above impose substantial limitations on the organism's capacity to resolve uncertainty. Granted, the bacterium's basic perceptual abilities permit the reduction of uncertainty about its current progress along the nutrient gradient, and in this extremely narrow sense it 'discovers' something about its relation to the world each time it samples its sensory states. But the agent never truly learns about the hidden states structuring its sensory experience – no amount of prediction error will compel it to change its fundamental expectations about the nature of its niche. Indeed, even the probabilistic mapping between its actions and their consequences is entirely opaque to the creature, given its incapacity to model statistical regularities over time. This agent is thus incapable of acting in ways that reduce ambiguity over its policies; it cannot acquire information that leads it to update its model parameters in ways that tend to optimise decision-making. Neither can it model volatility – its incapacity to represent contingent relations and temporally-extended dynamics necessarily precluding the representation of time-evolving alterations *amongst* such mappings. Consequently, all sources of uncertainty collapse into risk – the organism can only minimise uncertainty by inferring the least risky policy from its most recent belief update.

Before moving onto our next model, let us briefly consider whether a creature could exist by simply maintaining its homeostatic stability in the absence of exteroceptive modelling and action.[32] When a creature of this sort encounters surprising deviations from its homeostatic expectations it only ever adjusts its internal states, never its active states. It may for instance change its metabolic rate (e.g., slow respiration, inhibit protein synthesis) in response to altered nutrient conditions, rather than acting on the environment in order to reinstate homeostatic equilibrium.[33]

It is difficult to see how such a creature could actually exist in anything but a transitory, serendipitous manner. Changing its internal states in response to interoceptive prediction error is tantamount to yielding entirely to uncertainty. For example, as the nutrient gradient declines the organism's metabolic rate keeps decreasing, until it eventually starves to death – its states disperse throughout all possible states. An organism that fails to act upon its environment is ill-placed to avoid surprise and resist entropy. Only by happening to occupy a perfectly welcoming niche could it survive,

---

[32] See Godfrey-Smith (2016b) for a complementary discussion of this topic in relation to microbial protocognition and metabolic regulation.

[33] One might call this entity a *Spencerian creature*; i.e. an organism that responds to environmental change through "the continuous adjustment of internal relations to external relations" (Spencer 1867, p. 82; see discussion in Godfrey-Smith 1996, pp. 70–71). From an active inference perspective, this creature is the embodiment of pure perception; i.e. an organism that reconfigures its internal states (updates its model) in accordance with external conditions, without ever seeking to alter such conditions (cf. Bruineberg et al. 2018; Corcoran 2019).

but this is just to assume an environment devoid of uncertainty – not our world.[34]

4.3 Model 2: Hierarchical active inference

Next, let us consider a more elaborate version of our creature, now equipped with a more sophisticated, *hierarchical* generative model of its environment. Because higher levels of the generative model subtend increasingly broad temporal scales (Friston et al. 2017d; Kiebel et al. 2008), we shall see that this creature is capable of inferring the causes of slower fluctuations in the nutrient gradient. An implication of this arrangement is the emergence of parameters encoding higher-order expectations about the content and variability of sensory flows over time (cf. the fixed expectation of a high-nutrient state in Model 1).

In the interests of tractability, we limit ourselves to a fairly schematic illustration of hierarchical active inference in the context of circadian regulation. Circadian processes are near ubiquitous features of biological systems (even bacteria like *E. coli* show evidence of circadian rhythmicity; Wen et al. 2015), and provide a useful example of how internal dynamics can be harnessed to anticipate environmental variability.

Circadian clocks are endogenous, self-sustaining timing mechanisms that enable organisms to co-ordinate a host of metabolic processes over an approximately 24 h

---

[34] One might play with the idea of entities that could exist like this quite happily once the ideal, invariant niche is discovered – perhaps deep within rocky crevices or underwater (one is reminded of the sea squirt that consumes its own brain after settling upon a permanent home, but the anecdote turns out to be an exaggeration; see Mackie and Burighel 2005). However, entities of this sort would surely fail to qualify as *adaptive* biological systems – at least insofar as the notion of adaptability implies some capacity to maintain one's viability in the face of time-varying environmental dynamics (cf. 'mere' vs. 'adaptive' active inference; Kirchhoff et al. 2018). Moreover, such entities would also fail to qualify as *agents* in any biologically relevant sense (see for e.g., Moreno and Etxeberria 2005).

Interestingly, this scenario is reminiscent of a common criticism levelled against the free energy principle: the so-called *dark-room problem* (Friston et al., 2012e). The thrust of this argument is that free energy minimisation should compel agents to seek out the least surprising environments possible (e.g., a room devoid of stimulation) and stay there. Various rejoinders to this charge have been made (see for e.g., Clark 2018; Hohwy 2013; Schwartenbeck et al. 2013), including the observation that this strategy will inevitably lead to increasing free energy on account of accumulating interoceptive prediction error (Corcoran, 2019; Pezzulo et al., 2015). More technically, "itinerant dynamics in the environment preclude simple solutions to avoiding surprise" (Friston et al. 2009, p. 2), where the environment referred to here includes the biophysical conditions that obtain *within* the organism, as well as without. This is to say that the attractors around which adaptive biological systems self-organise are inherently unstable – both *autopoietic* ('self-creating') and *autovitiating* ('self-destroying') – thus inducing itinerant trajectories (*heteroclinic cycles*) through state-space (Friston, 2011, 2012b; Friston and Ao, 2012; Friston et al., 2012c).

In other words, dark rooms may very well appeal to creatures like us (e.g., as homeostatic sleep pressure peaks towards the end of the day), but the value such environments afford will inevitably decay as alternative possibilities (e.g., leaving the room to find breakfast after a good night's sleep) become more salient and attractive (cf. *alliesthesia*, the modulation of affective and motivational states according to (time-evolving) physiological conditions; Berridge 2004; Cabanac 1971).

period (Bailey et al., 2014; Dyar et al., 2018). From an allostatic perspective, circadian oscillations furnish a temporal frame of reference enabling the organism to anticipate (and efficiently prepare for) pattered changes in ecologically-relevant variables (e.g., diurnal cycles of light and temperature variation).[35] We can incorporate a molecular clock within our active inference agent by installing oscillatory protein pathways within its metabolic network (Nakajima et al., 2005; Rust et al., 2007; Zwicker et al., 2010). With this timing mechanism in place, our creature may begin to track systematic variations in the temporal dynamics of its internal and sensory states.

Let us imagine our organism exists in a medium that becomes increasingly viscous as temperature declines overnight. The impact of these environmental fluctuations is two-fold: Colder ambient temperatures cool the organism, slowing its metabolic rate; greater viscosity increases the medium's resistance, making chemotaxis more energy-intensive. Initially, the agent might interpret unexpectedly high rates of energy expenditure as indicative of suboptimal chemotaxis, compelling it to sample its *Tumble* policy more frequently in an effort to discover a nutrient-rich patch. Over time, however, the agent may come to associate a particular phase of its circadian cycle with higher average energy expenditure *irrespective* of policy selection. Our creature can capitalise on this information by scheduling its more expensive metabolic operations to coincide with warmer times of day, while restricting its nocturnal activity to a few essential chemical reactions. In other words, the agent can reorganise its behaviour (i.e. develop a rudimentary sleep/wake cycle) in order to improve its fit with its environment.[36]

This scenario is indicative of how a relatively simple hierarchical agent may come to model time-varying hidden states in the distal environment. Like its minimal active inference counterpart, the hierarchical agent registers fluctuations in its sensory and internal states, and responds to them appropriately given its available policies. Unlike the minimal agent, however, these rapid fluctuations are themselves subject to second-order processing in which successive sensory samples are integrated under a probabilistic representation of first-order variation (see Figure 2). The ability

---

[35] Note that the allostatic treatment of circadian regulation may in principle be extended to periodic phenomena spanning shorter or longer timescales; e.g., ultradian and circannual rhythms.

[36] This scenario is not meant to imply that circadian rhythms are actually acquired in this fashion (although they are clearly susceptible to modulation through external cues). Rather, the idea we are trying to illustrate here is the way hierarchical architectures ground adaptive regulation over longer timescales by dint of their capacity to capture recurrent, slowly evolving patterns of environmental variation.

to contextualise such variability in relation to the oscillatory dynamics of the circadian timekeeper enables the agent to infer that it is subject to periodic environmental perturbations, the origin of which can be parsimoniously ascribed to some unitary external process.[37] This example hints at a central tenet of the active inference scheme; namely, that the hierarchical organisation of the generative model implies a hierarchy of temporal scales, where causal dynamics subtending larger timeframes are encoded at higher levels of the model (Friston et al., 2017d; Kiebel et al., 2008).

The hierarchical picture we have sketched here speaks to two complementary aspects of *representational detachment* (cf. Gärdenfors 1995; Pezzulo and Castelfranchi 2007; Pezzulo 2008) engendered by allostatic architectures. First, the separation of processing layers within the model hierarchy gives rise to a kind of temporal decoupling, in which higher layers construct extended representations of low-level sensory states. Although it might be tempting to think of these representations as aggregates of successive sensory samples, this does not do justice to the sophisticated nature of perception under active inference. Rather, higher layers of the hierarchy are perpetually engaged in modelling the evolution of the organism's sensory and internal states, and thus inferring the probable motion of the distal causes of its sensory flows. Consequently, higher-order representations 'reach out' beyond the limits of each sensory moment, extrapolating forwards and backwards in time to synthesise an expanded temporal horizon (Figure 2A).

Second, there is a related sense in which higher-level processing within the hierarchy realises a more negative or reductive kind of detachment from low-level sensory input. Higher-level representations do not merely recapitulate (and predict) the bare contents of sensory experience, but seek instead to extract patterned continuities amidst the flux of sensory stimulation. This is to say that higher levels of the model attempt to carve out biologically-relevant signals within the agent's environment, while dampening or discarding the remaining content of sensory flows. This again speaks to the tension between accuracy and complexity, which can be recapitulated in terms

---

[37] Notice that the agent forms a representation of a hidden cause corresponding to diurnal patterns of temperature variation *despite* its lack of exteroceptive sensitivity to such variables as temperature, viscosity, light, etc. Rather, it detects regular changes in its dynamics that cannot be ascribed to its own actions (which average out across the 24 h period), and infers some hidden external process as being responsible for these changes. It might not be right to say the agent represents ambient temperature *per se*, nor indeed the higher-order causes of its oscillation (sun exposure, planetary rotation, etc.). Our agent lacks sufficient hierarchical depth to arrive at such conclusions, collapsing these fine-grained distinctions into a fairly 'flat', undifferentiated representation of diurnal variation.

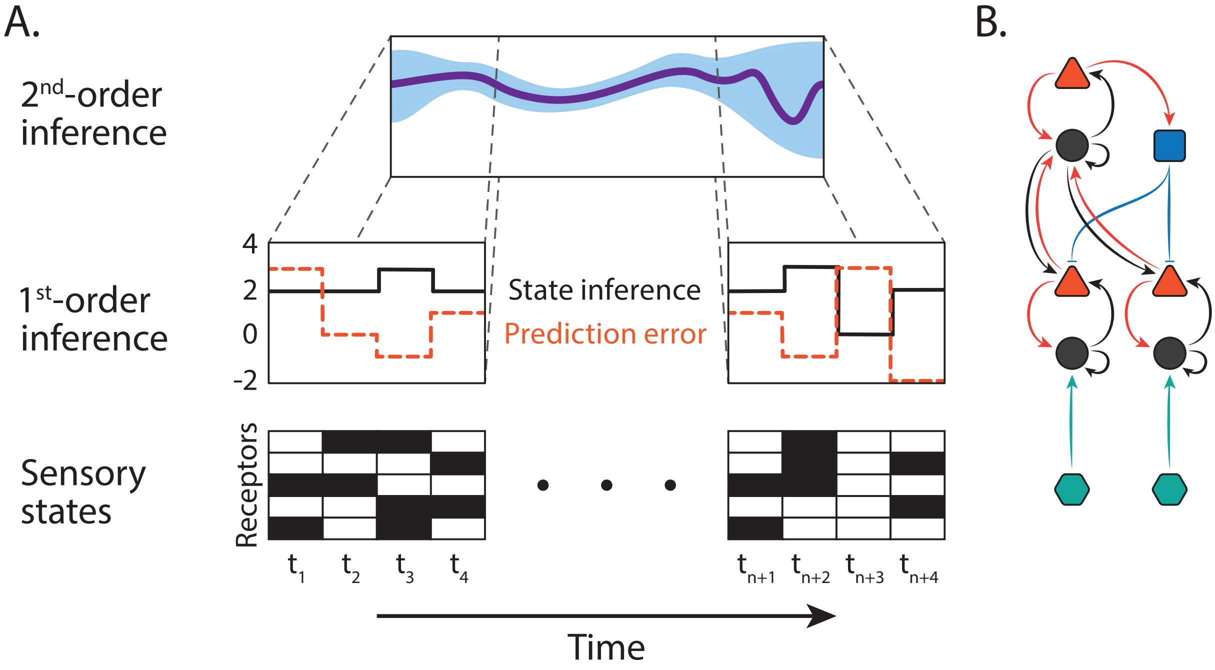

Figure 2: **Perceptual dynamics under hierarchical active inference.** **(A)** In this illustration, the minimal active inference scheme has been augmented with a second-order perceptual inference level that tracks changes in the nutrient gradient over time. The purple function in the top panel indicates the agent's time-evolving estimate of ambient nutrient levels, which is derived from first-order sensory inferences (middle panels) on successive chemosensory receptor states (raster plots, black cells indicate occupied receptor sites at time t). This function oscillates slowly as detected nutrient levels remain more or less stable over time, with each incoming packet of sensory information smoothly integrated within the broader temporal horizon of predicted and postdicted sensory states. The function begins to oscillate more rapidly when the organism experiences marked deviations from its expected states (right panels). This sudden volley of prediction error precipitates an increase in the precision on first-order prediction errors (reflected in the broadening blue ribbons encompassing the function), enhancing the agent's perceptual sensitivity to environmental fluctuations. **(B)** Schematic of a possible implementation of the hierarchical active inference scheme depicted in A. Sensory input from chemoreceptors (green hexagons) is received at the first processing level and compared to sensory expectations (grey circles). Discrepancies between expected and actual input generate prediction errors (red triangles), which are passed up the hierarchy to the second processing level. Crucially, these prediction errors are modulated by precision estimates (blue square), which determine the 'gain' or influence ascribed to error signals (where high gain compels expectation units to conform with prevailing sensory evidence). Expected precision over first-order prediction errors is modulated in turn by second-order prediction error, which increases the gain on first-order errors. See Kanai et al. (2015); Parr and Friston (2018a), and Shipp (2016) for more detailed discussion of how such hierarchical schemes might be implemented in the brain.

of trade-off between risk and ambiguity: Good models reduce ambiguity about the actual state of the world, without being overly sensitive to the data at hand (and thus at risk of accruing prediction error over the long-run; Hohwy 2017b).

If this account is on the right track, the generative model can be construed as a kind of (Bayesian) filter (Friston et al., 2010b) that strips sensory signals of their higher-frequency components as they are passed up the hierarchy. In conjunction with the 'horizontal' temporal processing described above (which can likewise be understood as noncausal filtering or smoothing, where past and future state estimates are updated in light of novel sensory data; Friston et al. 2017a), this 'vertical' filtration scheme enables the organism to form reliable higher-order representations of the slowly-evolving statistical regularities underlying rapid sensory fluctuations. The organism is thus able to model the slow oscillatory dynamics embedded within the distal structure of its eco-niche (e.g., the diurnal temperature cycle), even though the particular sensory states through which these dynamics are accessed may vary considerably over time (e.g., temperature variation may be modulated by multiple interacting factors subtending multiple timescales – momentary occlusion of the sun, daily and seasonal weather cycles, climate change, etc.).

These dual facets of representational detachment help to explain not only how the hierarchical agent learns about invariant properties of an ever-changing environment, but also how it can exploit such regularities to its advantage. Circadian rhythms offer a particularly good example of how abstract representations of oscillatory dynamics foster adaptive behaviour in the context of environmental uncertainty.[38] Given a reliable model of how certain environmental properties are likely to evolve, the agent can form allostatic predictions that enable it to act in preparation for impending conditions, even if such expectations run contrary to current sensory evidence. This observation speaks to the decomposition of (expected) uncertainty into risk and ambiguity: Accurate parameterisation of the processes responsible for generating sensory fluctuations helps resolve ambiguity about the likely consequences of alternative courses of action, thus enabling the agent to confidently sample amongst actions that carry least risk.

An interesting corollary of this view is the role of allostatic representations (e.g., circadian templates or programmes of activity) in compelling the agent to act 'as if'

---

[38] For discussion on the representational status of circadian rhythms, see Bechtel (2011) and Morgan (2018a,b).

particular states of affairs obtain. Under certain conditions, such allostatic predictions amount to a kind of *false inference* about the hidden states that are currently in play. Although such predictions might be expected to engender actions that accumulate prediction error, the agent persists with them on account of their prior precision, which causes conflicting sensory evidence to be downweighted or attenuated (Brown et al., 2013; Wiese, 2017).

Returning to our earlier example, let us imagine that the hierarchical agent leverages its internal representation of diurnal temperature variation to schedule its activities to coincide with favourable environmental conditions. For instance, the organism might preemptively downregulate metabolic activity in preparation for nocturnal quiescence, irrespective of whether the ambient temperature has declined to an extent that would impair its metabolic efficiency. Likewise, the agent might begin to upregulate its activity around its usual time of 'awakening', despite the fact that this routine provokes an elevated rate of energy expenditure on an usually chilly morning.

On first blush, this arrangement might seem suboptimal: Surely the agent would be better off tuning its behaviour to *actual* environmental conditions, rather than relying on error-prone predictions? However, this would simply return us to the kind of closed-loop architecture of the minimal active inference agent; a creature incapable of distinguishing a genuine change in distal conditions from a transient deviation in its sensory states. In this sense our agent's circadian gambit constitutes a more intelligent mode of regulation – armed with knowledge of how state trajectories tend to evolve, the organism acts on the assumption that the future will roughly approximate the past, and treats transient deviations from this prescribed pattern as mere noise (i.e. the inherent risk associated with stochastic processes).

Hence, although circadian rhythms might not guarantee ideal behaviour on shorter timescales, their adaptive value inheres in their ability to approximate the trajectory of homeostatically-relevant states over time. Such allostatic representations provide useful heuristics for guiding action – behaving in accordance with circadian predictions keeps the agent within the vicinity of its attracting set, thus affording a highly efficient means of reducing average uncertainty. Representations of this sort are insensitive to short-term fluctuations precisely because such transient dynamics (e.g., an unseasonably cold morning) are unlikely to afford information that improves its capacity to accurately predict future states. Circadian rhythms are therefore 'robust' to outlying or stochastic fluctuations in sensory data, thus constituting a reliable model of the underlying generative process.[39]

Increasing hierarchical depth affords a dramatic increase in the extent to which sensory information can be integrated and abstracted to form complex representations of hidden dynamics (cf. the transition from sensation to perception in emulation theory; Grush 2004). Such architectures necessarily imply the increasing separation of decision-making from the immediate flux of sensory input: Policy selection is informed by higher-order expectations about the likely evolution of states and outcomes over time, and increasingly oriented towards the resolution of uncertainty across broader temporal scales.

In contrast to the minimal active inference agent, the hierarchical organism can exploit regularities in its environment to predict when and where it will be best placed to act, rather than responding reflexively to online sensory updates. Such inferences are aided by the higher precision with which the organism is able to predict the evolution of its sensory flows under a particular policy. The reliability of such expectations speaks to the agent's capacity to minimise ambiguity through the accurate estimation of ecologically-relevant parameters, which in turn enables the organism to accurately estimate the risk associated with each policy.

While deep hierarchical architectures afford substantial advantages over the minimal scheme of Model 1, their capacity to reduce uncertainty through parameter estimation is most effective in a relatively stable world. Sudden alterations in environmental conditions (e.g., exchanging the European winter for the Australasian summer) require relatively long periods of reparameterisation, and may engender suboptimal, surprise-accruing behaviour in the interim. Volatility poses even more of a challenge for this sort of creature, which lacks any means of evaluating whether its prediction errors emanate from the inherent risk associated with its policies, or a fundamental change in environmental dynamics. We turn to this problem next.

---

[39] The remarkable robustness of circadian oscillations is thrown into relief whenever one traverses several time-zones – a good example of how strongly-held (i.e. high-precision or 'stubborn'; see Yon et al. 2019) allostatic expectations persist in the face of contradictory sensory evidence (i.e. the phase-shifted photoperiod and feeding schedule, to which the system eventually recalibrates; Asher and Sassone-Corsi 2015; Menaker et al. 2013).

4.4 Model 3: Counterfactual active inference

Our final model describes a biological agent equipped to perform *counterfactual* active inference. While this class of inferential architectures is undoubtedly the most complex and sophisticated of our three models, we restrict our treatment to the generic ways in which counterfactual inference supports uncertainty reduction under volatility.[40]

Counterfactual active inference adds to the hierarchical processing of progressively deeper models through subjunctive processing: The agent can evaluate the expected free energy of alternative policies *under a variety of different contexts* before alighting on the best course of action (Friston, 2018; Limanowski and Friston, 2018). Our understanding of subjunctive processing draws on the Stalnaker-Lewis analysis of counterfactual conditionals, where the truth-conditions of a consequent are determined in relation to the *possible world* invoked by its antecedent (Lewis 1973b; Stalnaker 1968, see also Nute 1975; Sprigge 1970; Todd 1964).[41] In the context of active inference, counterfactual processing translates to the simulation of the sensory states the organism *would* observe if it *were* to enact a certain policy *under a particular set of model parameters* (i.e. a possible world).

Our formulation of counterfactual inference implies two complementary forms of uncertainty reduction, which we shall briefly introduce here. The first of these involves counterfactual inference on policies under spatiotemporally distal conditions. For example, the agent could reflect on a previous decision that precipitated a negative outcome, and consider how events might have unfolded differently (for better or worse) had it selected an alternative course of action ('retrospective' inference). Similarly, the agent could envisage a scenario it might encounter in the future, and imagine how various policies might play out under these circumstances ('prospective' inference). This kind of counterfactual processing is useful for reducing ambiguity over the outcomes expected under various policies, and is integral to many sophisticated cognitive processes (e.g., causal induction, mental time travel, mindreading, etc.; Buckner and Carroll 2007; Schacter and Addis 2007; Suddendorf and Corballis 1997, 2007).

The second form of uncertainty reduction mediated by counterfactual processing

---

[40] For further discussion of counterfactual representation under predictive processing, see Clark 2016, Ch. 3; Friston et al. 2012b; Friston 2018; Palmer et al. 2015; Pezzulo et al. 2015; Seth 2014, 2015.

[41] Note that our use of counterfactual semantics here is not intended to imply that cognition bears any necessary resemblance to linguistic processing; it is simply adopted as a convenient way of characterising the logic of model selection under active inference.

pertains to the arbitration of policies when the state of the world is ambiguous. This situation may arise due to uncertainty about the context that currently obtains (or relatedly, uncertainty over the consequences of policies within a particular context), or because the niche it inhabits is inherently volatile (i.e. prone to fluctuate in ways that are relevant for the organism's wellbeing, yet difficult to anticipate). Under such circumstances, counterfactual hypotheses may prove useful in two ways: (1) they may enable the agent to infer the policy that minimises (average) risk across a variety of possible worlds; (2) they may point to 'epistemic' actions that would disambiguate the *actual* state of the world (i.e. disclose which likelihood mapping currently obtains), leading to improved precision over policies.

As a brief illustration of counterfactual inference, let us consider an iteration of our *E. coli*-like creature that can evaluate the outcomes of its policies across several possible worlds. An organism sensitive to incident light could for instance run a counterfactual simulation for a possible world in which there is much scattered sunlight, and compare this to an alternative world featuring relatively little sunlight. If sunlight poses a threat to the bacterium (perhaps sun exposure causes the nutrient patch to dry up), tumbling constitutes a riskier strategy in the sun-dappled world. If it can order these possible worlds on the basis of their similarity to the actual world, then these counterfactual simulations could prove informative about the best action to take in a particular situation.[42] Should the sun-dappled world turn out more similar to the actual world, then the organism would do well to confine its foraging activity to shady regions of the environment. The agent might consequently adapt its policies such that it tolerates gradient descent in the context of low incident light, only risking the *Tumble* policy when the nutrient supply is almost exhausted.

Counterfactual processing enriches the generative model greatly, relative to the hierarchical organisation described in the previous section. Now there is wholly detached generative modelling of fine-grained elements of the prediction error landscape through simulated action; there is (Bayesian) model selection in terms of the best policy (i.e. minimising the free energy between the nutrient gradient simulated under a

---

[42] Interestingly, recent psychological evidence suggests that counterfactual scenarios deemed more similar to previously experienced events are perceived as more plausible and easier to envisage (i.e. simulate) than more distant alternatives (Stanley et al., 2017). This observation lends weight to the idea that humans evaluate competing counterfactual predictions in accordance with their proximity to actual states of affairs, where proximity or similarity might be cashed out in terms of (Bayesian) model evidence (see FitzGerald et al. 2014).

policy and the organism's expected nutrient gradient; cf. FitzGerald et al. 2014; Friston et al. 2016, 2017b; Parr and Friston 2018b); and there is processing that orders possible worlds (i.e. hypotheses entailed under competing model parameterisations) according to their comparative similarity to the actual world (where similarity may be cashed out in terms of representations of law-like relations (e.g., between nutrient gradient and sunlight) and particular matters of fact (e.g., amount of nutrient and sunlight); cf. Lewis 1973a,b, 1979). This contrasts sharply with the hierarchical agent, whose representational states are never completely detached from the content of its sensory flows, and whose active states are modulated gradually in response to reliable patterns of covariation.

With counterfactual inference at its disposal, the organism is potentially even better equipped to meet the demands of a complex and capricious environment.[43] Rather than engaging 'hard-wired' responses to current states (cf. Model 1), or 'soft-wired' responses to anticipated states (cf. Model 2), it can exploit *offline* computation of the likely consequences of different policies under various hypothetical conditions (Gärdenfors, 1995; Grush, 2004; Pezzulo, 2008). This affords the opportunity to generate and test a wide variety of policies in the safety of its imagination, where actions that turn out to be too risky (or downright stupid) can be safely trialed and (hopefully) rejected (cf. Craik 1943, p. 61; Dennett 1995, pp. 375–376; Godfrey-Smith 1996, pp. 105–106). This capacity (or *competence*, see Williams 2018) to disengage from the present and undertake such 'thought experiments' confers a powerful mechanism for innovation, problem-solving, and (vicarious) learning – major advantages in complex environments (Buzsáki et al., 2014; Mugan and MacIver, 2019; Redish, 2016).

The counterfactual active inference scheme described here implies additional degrees of organismic complexity that can be exploited to mitigate the impact of environmental uncertainty. The counterfactual agent is not only capable of 'expecting the unexpected' (inasmuch as it can countenance states of affairs that are unlikely un-

---

[43] One caveat to this claim is that the (neuro)physiological mechanisms and cognitive operations required to enrich and exploit counterfactual predictive models may themselves engender additional costs (e.g., planning a new course of action requires time, energy, and effort; see Zénon et al. 2019). We assume that the costs incurred by such processes 'pay for themselves' over the long-run (or at least tend to on average), insofar as they enable the agent to exploit prior experience in ways that are conducive to adaptive behaviour (see Buzsáki et al. 2014; Pezzulo 2014; Pezzulo et al. 2017; Suddendorf et al. 2018). It is also worth pointing out that some of the costs engendered by counterfactual inference-supporting architectures may be mitigated by a variety of adaptive strategies (e.g., model updating during sleep, habitisation of behaviour under stable and predictable conditions; see FitzGerald et al. 2014; Friston et al. 2017b; Hobson and Friston 2012; Pezzulo et al. 2016).

der its current model of reality), but can prepare for it too – exploiting counterfactual hypotheses to formulate strategies for solving novel problems that might arise in the future (e.g., deciding what one should do in the event of sustaining a puncture while cycling to work). Moreover, the agent may organise its policy sets in ways that are sensitive to outcome contingencies, such that it can choose a backup policy if its initial plan is thwarted (e.g., being prepared to order the apple pie if the tiramisu has sold out). This ability to deftly switch between a subset of low-risk policies may confer a huge advantage under conditions of volatility, where the time and effort required to re-evaluate a large array of policies from scratch could prove extremely costly.

Counterfactual processing is also valuable when the system is confronted with a sudden or sustained volley of prediction error. The counterfactual agent is able to interpret such signals as evidence that the hidden dynamics underwriting its sensory flows may have changed in some significant way (e.g., finding oneself confronted by oncoming traffic), and can draw on alternative possible models to evaluate which parameterisation affords the best explanation for the data at hand (cf. *parameter exploration*; Schwartenbeck et al. 2019). If the contingent relations structuring relevant environmental properties have indeed altered (e.g., realising one is visiting a country where people drive on the opposite side of the road), the agent will need to update its model in order to capture these novel conditions (see Sales et al. 2019). Failure to do so runs the risk of accruing further prediction error, since persisting with policies predicated on inaccurate (i.e. 'out-of-date') likelihood mappings may yield highly surprising outcomes.

One way to assess whether conditions or contexts have indeed changed is to engage in *epistemic action*, the final feature of counterfactual active inference we address here. Epistemic actions are active states that are sampled in order to acquire information about environmental contingencies (Friston et al., 2015b, 2016, 2017a,d).[44] When faced with the problem of identifying which model best captures the causal structure of the world, the agent can run simulations to infer the sensory flows each model predicts under a certain policy. The agent can then put these hypotheses to the test

---

[44] For the purposes of this brief discussion, we limit the scope of epistemic action to instances where the organism actively intervenes on its environment in order to resolve uncertainty. It is worth noting, however, that the concept can also refer to *mental actions* or cognitive operations that reduce uncertainty (see for e.g., Metzinger 2017; Pezzulo et al. 2016; Pezzulo 2017). On this broader understanding, one might construe the different varieties of counterfactual processing described above as covert modes of epistemic action.

by sampling actions designed to arbitrate amongst competing predictions (Seth, 2015). If the agent selects actions that are high in *epistemic value*, it will observe outcomes that afford decisive evidence in favour of the model that best captures the current environmental regime.

The possibility of resolving ambiguity over the parameterisation of state–outcome contingencies through counterfactually-guided epistemic action also extends to ambiguity over policies. Here, the agent may run counterfactual simulations to infer actions that are likely to harvest information that clarifies the best policy to pursue.[45] These epistemic capabilities imply that the policies of the counterfactual agent are not only scored with respect to risk-reduction or expected value (i.e. the extent to which they are expected to realise a *preferred* outcome), but also with respect to ambiguity-reduction or epistemic value (i.e. the extent to which they are expected to produce an *informative* outcome). Such epistemic or uncertainty-reducing actions are unavailable to the (merely) hierarchical agent, who can only reduce ambiguity over model parameters by slowly tuning its estimates to capture stable, enduring patterns of variation.

Finally, it is interesting to remark how epistemic action contributes to the practical utility of cognition as understood under the environmental complexity thesis. Following Dewey (1929), Godfrey-Smith (1996, pp. 116–120) notes that cognition is most likely to be useful in environments that comprise a mixture of regularity and unpredictability. Specifically, distal states should vary in ways that are *a priori* unpredictable (but worth knowing about), while maintaining a stable relationship with proximal states (see also Dunlap and Stephens 2016). The capacity to engage in epistemic action enhances the potential utility of cognition precisely insofar as it helps the agent to reduce uncertainty over this mapping, thus affording more precise knowledge (or novel insight; Friston et al. 2017b) about the state of the world and its possible alternatives.[46]

---

[45] Such activity is sometimes referred to as *epistemic foraging*, where the agent seeks out information about the way state transitions are likely to unfold (Friston et al., 2017d; Mirza et al., 2016; Parr and Friston, 2017). For a nice example of epistemic foraging in wild dolphins, see Arranz et al. (2018).

[46] Note that Godfrey-Smith's exposition of the relation between variability and stability is somewhat ambiguous from an active inference perspective. Unpredictable variation in distal states could be interpreted in terms of volatility (imprecision over state transitions), while reliable correlations between proximal and distal states could be interpreted as predictable state transitions. Epistemic action can be deployed to reduce uncertainty in either case; a suitably sophisticated hierarchical agent could potentially cope with the latter scenario, but not the former.

5 Two options for cognition

We began this paper with the lofty ambition of learning something about the nature and function of cognition, but have for the most part been careful to eschew talk of the cognitive or the mental. In this final section, we sketch out some of the broader implications of our analysis for the concept of biological cognition, and how the latter might be delimited from more general notions of life and adaptive plasticity.

As a precursory step, let us begin by considering how the three schematic models described in Section 4 might relate to real biological agents. One obvious strategy would be to map these architectures onto different taxonomic classes. For instance, one might construe the difference between these models as approximating the difference between relatively primitive organisms (like *E. coli* and other unicellular organisms), creatures with some degree of hierarchical depth (like reptiles or fish), and animals that demonstrate evidence of counterfactual sensitivity (like rodents; e.g., Redish 2016; Steiner and Redish 2014; Sweis et al. 2018; corvids; e.g., Bugnyar et al. 2016; Kabadayi and Osvath 2017; Raby et al. 2007; and primates; e.g., Abe and Lee 2011; Krupenye et al. 2016; Lee et al. 2005).

This approach is immediately undermined however by the remarkable complexity exhibited by (at least some) unicellular organisms. Bacteria like *E. coli* integrate information over a variety of sensory channels, modulate their metabolic and chemotactic activity in response to reliable environmental contingencies, and alternate policy preferences in a context-sensitive fashion (Ben-Jacob 2009; Freddolino and Tavazoie 2012; Hennessey et al. 1979; Mitchell et al. 2009; Salman and Libchaber 2007; Tagkopoulos et al. 2008; Tang and Marshall 2018; see also Van de Cruys 2017, for discussion from a predictive processing perspective). Although this does not rule out the possible existence of minimal active inference agents, it might suggest that *all* extant lifeforms instantiate some form of allostatic architecture. This raises the question of whether meaningful distinctions can be drawn in terms of hierarchical organisation (e.g. shallow vs. deep hierarchies), and whether such distinctions can be systematically mapped to particular functional profiles (e.g., capacities for learning and adaptive flexibility).

It might also be tempting to think of our model organisms as exemplifying creatures that are more or less ‘evolved’ or ‘adapted’ to their environment. Undoubtedly, the counterfactual agent comprises a more complex information-processing architecture than its minimal active inference counterpart, one equipped with a much greater capacity for flexible, selective adaptation to the vicissitudes wrought by uncertainty. However, we must be careful not to conflate adaptation to a specific set of environmental properties with adaptation to environmental complexity *per se*. On both the environmental complexity thesis and the free energy principle, organisms are adapted to their environments to the extent that they successfully track and neutralise *ecologically-relevant* sources of uncertainty (cf. 'frugal' generative models; Baltieri and Buckley 2017; Clark 2015). This means that organisms comprising radically divergent degrees of functional complexity can in principle constitute equally good models of the same environment, assuming they are equally capable of acting in ways that minimise the conditional entropy over their sensory states.

Finally, given that the free energy principle conceives of all biological agents as being engaged in the same essential activity (i.e. the singular project of minimising free energy, maximising self-evidence, and thus conserving self-organisation over time), one might question whether there really are any substantive differences to be found between the levels of our three-tiered scheme. In conjunction with the argument presented in the previous paragraph, it might seem that these architectures differ from one another in a fairly superficial way: They simply illustrate alternative solutions to the fundamental problem of uncertainty reduction over time.

This point notwithstanding, we believe that the distinct functional capacities we have ascribed to these models carry important implications about the origins and limits of cognition. The fact that all three architectures are afforded equal footing by the free energy principle does not speak against this view – despite its neuroscientific origins (Friston, 2002, 2003, 2005), the free energy principle makes no explanatory commitments to cognition *per se*; it simply imposes certain formal constraints on the sort of functional organisation a cognitive system must realise in order to resist entropy. This marks a significant distinction from the environmental complexity thesis, which on Godfrey-Smith's telling limits its explanatory scope to the *subset* of living organisms that count as cognitive agents.

Put differently, the free energy principle is *neutral* on the ontological relation between life and cognition (*pace* Kirchhoff and Froese 2017). The environmental complexity thesis, on the other hand, endorses a *weak continuity* ("Anything that has a mind is alive, although not everything that is alive has a mind"; Godfrey-Smith 1996, p. 72) without specifying a principled way of demarcating the boundary between the cognitive and the non-cognitive.[47] We propose that this boundary can be located at the nexus between hierarchical and counterfactual forms of active inference. This would mean that only those biological systems capable of engaging in fully detached modes of representation, and of exploiting such representations for the purposes of uncertainty reduction, count as cognitive agents.[48]

Associating cognition with counterfactual active inference might strike some as unduly restrictive, limiting category membership to humans and only the most intelligent of mammals and birds (for instance). It is important to bear in mind, however, that our construal of counterfactual processing is a formal one; many kinds of animals are likely to exploit counterfactual inferences in ways that enable them to learn about the world and make sensible (uncertainty-reducing) decisions. Some of these processing architectures might turn out to be highly impoverished compared to the rich counterfactual capacities at our own disposal (cf. Carruthers 2004), but we consider this difference a matter of degree, not kind.

Notably, our counterfactual criterion does not exclude such organisms as bacteria, protists, and plants from the cognitive domain by *fiat*. If clever empirical studies were to reveal that *E. coli* (for example) proactively solicit ambiguity-reducing information to plan their future chemotactic forays, this would afford compelling evidence they constitute cognitive agents. However, as pointed out in recent debates about future-oriented cognition in non-human animals, seemingly complex patterns of behaviour do not always licence the attribution of complex representational or inferential capaci-

---

[47] Godfrey-Smith thus rejects *strong continuity*, the view that "[l]ife and mind have a common abstract pattern or set of basic organizational properties. [...] Mind is literally life-*like*" (1995, p. 320, emphasis in original). Evan Thompson (2007) has defended a position similar to this ('deep continuity'), albeit with the addition of an existential-phenomenological supplement (for discussion, see Wheeler 2011). This view inherits from Maturana's canonical account of autopoiesis, where one finds the strongest expression of life–mind continuity: "Living systems *are* cognitive systems, and living as a process *is* a process of cognition" (Maturana and Varela 1980, p. 13, emphasis added; see also Heschl 1990).

[48] It is perhaps worth noting that other scholars have used the criterion of "detachment" (or "decouplability") to distinguish representational versus non-representational agents, rather than cognitive versus non-cognitive agents (cf. Clark and Grush 1999; Grush 2004). Without digressing into a discussion of the relationship between representational and cognitive systems, we remark that our view conceives of cognition as a computational architecture that engages in a particular subset of representational operations – i.e. the generation, manipulation, and evaluation of counterfactual model predictions. These operations are situated within a broader class of uncertainty-resolving processes, including the homeostatic and allostatic representational schemes outlined in Section 4.

ties (Redshaw and Bulley 2018; Suddendorf and Redshaw 2017; see Mikhalevich et al. 2017, for an environmental complexity-inflected counterargument). If empirical observations can be parsimoniously explained by appeal to such allostatic mechanisms as information integration (Read et al. 2015) and elemental learning (Giurfa 2013; Perry et al. 2013), admittance to the cognitive domain ought to be withheld.

An alternative (and increasingly popular) approach would be to ascribe some form of 'minimal' or 'proto-cognitive' status to bacteria, plants, and other aneural organisms (Ben-Jacob 2009; Calvo Garzón and Keijzer 2011; Gagliano 2015; Godfrey-Smith 2016a,b; Lyon 2015; Segundo-Ortin and Calvo 2019; Smith-Ferguson and Beekman 2019; van Duijn et al. 2006; for a dissenting view, see Adams 2018). Such terms might seem appealing in light of the mounting body of research claiming that many 'simple' organisms engage in primitive or precursory forms of cognitive activity (Baluška and Levin, 2016; Levin et al., 2017; Tang and Marshall, 2018). Granting such cases do indeed demonstrate genuine instances of learning, memory, decision-making, and so on, it seems only the staunchest of neuro-chauvinists would persist in denying the cognitive status of such organisms.

While we cannot do justice to this complex topic here, a few remarks are in order. First, we should acknowledge that there may be few substantive differences between the kinds of organisms we designate as hierarchical or allostatic agents, and the biological systems Godfrey-Smith and others would identify as exhibiting 'minimal' or 'proto-cognitive' capacities (e.g., Godfrey-Smith 2002, 2016b).[49] Both categories imply systems that track relevant states in their (internal and external) environments, and exploit this information to adaptively regulate their activity. Both categories also imply some form of evolutionary precedence over 'fully' cognitive agents – cognition 'proper' builds on the foundations laid by allostatic/proto-cognitive architectures.

Our concern with such terminology is that it implies the ascription of some form of cognitive capacity, while remaining opaque as to its precise relation to 'full-blown' cognition – including the reason for its demarcation from the latter. Is there some fundamental cognitive ingredient that proto-cognition lacks, or is it simply a scaled-down, severely degraded version of (say) animal cognition? If the latter, is the distinc-

---

[49] 'Minimal cognition' is perhaps more closely associated with a rather different set of philosophical views than those espoused by Godfrey-Smith (e.g., anti-representationalism, situated and embodied cognition; Barandiaran and Moreno 2006; Beer 2003; van Duijn et al. 2006). We take the main thrust of our argument to be equally applicable to these positions.

tion between proto- and 'genuine' cognition marked by a critical boundary, or is the difference gradual and indeterminate? Godfrey-Smith explicitly endorses some variety of the latter view, frequently remarking that cognition 'shades-off' into other biological processes. But if proto-cognitive organisms ultimately fail to qualify as cognitive agents,[50] such talk obscures a fundamental *dis*continuity.

We take it that the distinction between relatively sophisticated, 'intelligent' forms of allostatic regulation on the one hand, and rather simple forms of counterfactual processing on the other, marks the subtle but significant functional boundary hinted at in Godfrey-Smith's analysis. Defining cognition in this way implies a clear functional demarcation between organisms that exhibit or lack the capacity for counterfactual processing. This proposed requirement is – in most cases – stricter than other criteria often mentioned in the debate about minimal cognition: it implies that organisms that only engage in allostatic regulation (sometimes requiring forms of learning, memory or decision-making) would not be ascribed fully-fledged cognition. Of course, testing which organisms meet this counterfactual criterion remains an important conceptual and empirical challenge.

In this respect, our proposed definition is not neuro-chauvinistic, but is focused rather on a functional (computationally-grounded) definition of cognition that can be met – at least in principle – by many different kinds of organisms. On this view, a minimally cognitive agent is a minimally counterfactual agent – an organism that not only learns about itself and its environment, but is capable of imagining them anew. If we are wrong, and sophisticated forms of cognitive activity simply emerge as allostatic processing schemes become increasingly more powerful and hierarchically elaborate, then a single dimension along which cognition 'shades off' into primitive forms of sensorimotor control and metabolic regulation would seem the better option.

---

[50] When pressed, Godfrey-Smith seems to hold this view: "I do *not* claim that bacteria exhibit cognition; this is *at most* a case of proto-cognition" (2002, p. 223, emphasis added).

Acknowledgements

AWC is supported by an Australian Government Research Training Program (RTP) scholarship. JH is supported by the Australian Research Council (DP160102770, DP190101805). This research has received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 785907 (Human Brain Project SGA2 to GP). We would like to thank participants at the *Science of the Self* research forum and the *22nd Annual Meeting of the Association for the Scientific Study of Consciousness* for feedback on earlier presentations of this work. We also wish to thank Louise Kyriaki, Dan Williams, and members of the Cognition & Philosophy Lab – especially Stephen Gadsby, Andy Mckilliam, Kelsey Perrykkad, and Iwan Williams – for insightful comments on earlier versions of this manuscript.

References

Abe, H. and Lee, D. (2011). Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. *Neuron*, 70(4):731–741.

Adams, F. (2018). Cognition wars. *Studies in History & Philosophy of Science*, 68:20–30.

Adams, R. A., Shipp, S., and Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure & Function*, 218(3):611–643.

Ainley, V., Apps, M. A. J., Fotopoulou, A., and Tsakiris, M. (2016). 'bodily precision': A predictive coding account of individual differences in interoceptive accuracy. *Philosophical Transactions of the Royal Society B*, 371(20160003):1–9.

Allen, M. and Friston, K. J. (2018). From cognitivism to autopoiesis: Towards a computational framework for the embodied mind. *Synthese*, 195(6):2459–2482.

Allen, M. and Tsakiris, M. (2018). The body as first prior: Interoceptive predictive processing and the primacy of self-models. In Tsakiris, M. and De Preester, H., editors, *The interoceptive mind: From homeostasis to awareness*, pages 27–45. Oxford: Oxford University Press.

Andrews, B. W., Yi, T.-M., and Iglesias, P. A. (2006). Optimal noise filtering in the chemotactic response of escherichia coli. *PLoS Computational Biology*, 2(11):e154.

Arranz, P., Benoit-Bird, K. J., Southall, B. L., Calambokidis, J., Friedlaender, A. S., and Tyack, P. L. (2018). Risso's dolphins plan foraging dives. *Journal of Experimental Biology*, 221(4):jeb165209.

Ashby, W. R. (1940). Adaptiveness and equilibrium. *British Journal of Psychiatry*, 86(362):478–483.

Ashby, W. R. (1956). *An introduction to cybernetics*. London: Chapman & Hall Ltd.

Ashby, W. R. (1958). Requisite variety and its implications for the control of complex systems. *Cybernetica*, 1(2):83–99.

Asher, G. and Sassone-Corsi, P. (2015). Time for food: The intimate interplay between nutrition, metabolism, and the circadian clock. *Cell*, 161(1):84–92.

Attias, H. (2003). Planning by probabilistic inference. In Bishop, C. M. and Frey, B. J., editors, *Proceedings of the Ninth International Conference on Artificial Intelligence and Statistics*. Society for Artificial Intelligence and Statistics.

Bach, D. R. and Dolan, R. J. (2012). Knowing how much you don't know: A neural organization of uncertainty estimates neural organization of uncertainty estimates. *Nature Reviews Neuroscience*, 13:572–586.

Badcock, P. B., Davey, C. G., Whittle, S., Allen, N. B., and Friston, K. J. (2017). The depressed brain: An evolutionary systems theory. *Trends in Cognitive Sciences*, 21(3):182–194.

Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences*, 12(5):193–200.

Bailey, S. M., Udoh, U. S., and Young, M. E. (2014). Circadian regulation of metabolism. *Journal of Endocrinology*, 222(2):R75–R96.

Baltieri, M. and Buckley, C. L. (2017). An active inference implementation of phototaxis. In Knibbe, C., Beslon, G., Parsons, D., Misevic, J. R.-C., Bredèche, N., Hassas, S., Simonin, O., and Soula, H., editors, *Proceedings of ECAL 2017: The 14th European Conference on Artificial Life*, pages 36–43. Cambridge, MA: MIT Press.

Baluška, F. and Levin, M. (2016). On having no head: Cognition throughout biological systems. *Frontiers in Psychology*, 7:902.

Barandiaran, X. E. and Moreno, A. (2006). On what makes certain dynamical systems cognitive: A minimally cognitive organization program. *Adaptive Behavior*, 14(2):171–185.

Barrett, L. F., Quigley, K. S., and Hamilton, P. (2016). An active inference theory of allostasis and interoception in depression. *Philosophical Transactions of the Royal Society B*, 371(20160011):1–17.

Barrett, L. F. and Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature Reviews Neuroscience*, 16(7):419–429.

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711.

Bauman, D. E. (2000). Regulation of nutrient partitioning during lactation: Homeostasis and homeorhesis revisited. In Cronjé, P. B., editor, *Ruminant physiology: Digestion, metabolism, growth and reproduction*, chapter 18, pages 311–328. New York: CABI Publishing.

Bechtel, W. (2011). Representing time of day in circadian clocks. In Newen, A., Bartels, A., and Jung, E.-M., editors, *Knowledge and representation*, chapter 7, pages 129–162. Stanford, C.A.: CSLI Publications.

Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–243.

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., and Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9):1214–1221.

Ben-Jacob, E. (2009). Learning from bacteria about natural information processing. *Annals of the New York Academy of Sciences*, 1178:78–90.

Berg, H. C. and Brown, D. A. (1972). Chemotaxis in escherichia coli analysed by three-dimensional tracking. *Nature*, 239(5374):500–504.

Bernard, C. (1974). *Lectures on the phenomena of life common to animals and plants*. American Lecture Series. Springfield, IL: Charles C. Thomas Pub Ltd.

Berntson, G. G. and Cacioppo, J. T. (2000). From homeostasis to allodynamic regulation. In Cacioppo, J. T., Tassinary, L. G., and Berntson, G. G., editors, *Handbook of psychophysiology*, chapter 17, pages 459–481. Cambridge: Cambridge University Press, 2nd edition.

Berridge, K. C. (2004). Motivation concepts in behavioral neuroscience. *Physiology & Behavior*, 81(2):179–209.

Bich, L., Mossio, M., Ruiz-Mirazo, K., and Moreno, A. (2016). Biological regulation: Controlling the system from within. *Biology & Philosophy*, 31(2):237–265.

Birkhoff, G. D. (1931). Proof of the ergodic theorem. *Proceedings of the National Academy of Sciences*, 17(12):656–660.

Bland, A. R. and Schaefer, A. (2012). Different varieties of uncertainty in human decision-making. *Frontiers in Neuroscience*, 6:85.

Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877.

Bogacz, R. (2017). A tutorial on the free-energy framework for modelling perception and learning. *Journal of Mathematical Psychology*, 76:198–211.

Botvinick, M. and Toussaint, M. (2012). Planning as inference. *Trends in Cognitive Sciences*, 16(10):485–488.

Bradley, R. and Drechsler, M. (2014). Types of uncertainty. *Erkenntnis*, 79(6):1225–1248.

Brown, H., Adams, R. A., Parees, I., Edwards, M., and Friston, K. J. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, 14(4):411–427.

Bruineberg, J., Rietveld, E., Parr, T., van Maanen, L., and Friston, K. J. (2018). Free-energy minimization in joint agent-environment systems: A niche construction perspective. *Journal of Theoretical Biology*, 455:161–178.

Buckley, C. L., Chang, S. K., McGregor, S., and Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81:55–79.

Buckner, R. L. and Carroll, D. C. (2007). Self-projection and the brain. *Trends in Cognitive Sciences*, 11(2):49–57.

Bugnyar, T., Reber, S. A., and Buckner, C. (2016). Ravens attribute visual access to unseen competitors. *Nature Communications*, 7:10506.

Burdakov, D. (2019). Reactive and predictive homeostasis: Roles of orexin/hypocretin neurons. *Neuropharmacology*, 154:61–67.

Buzsáki, G., Peyrache, A., and Kubie, J. (2014). Emergence of cognition from action. *Cold Spring Harbor Symposia on Quantitative Biology*, 79:41–50.

Cabanac, M. (1971). Physiological role of pleasure. *Science*, 173(4002):1103–1107.

Cabanac, M. (2006). Adjustable set point: To honor harold t. hammel. *Journal of Applied Physiology*, 100(4):1338–1346.

Calvo, P. and Friston, K. J. (2017). Predicting green: Really radical (plant) predictive processing. *Journal of the Royal Society Interface*, 14(20170096):1–11.

Calvo Garzón, P. and Keijzer, F. (2011). Plants: Adaptive behavior, root-brains, and minimal cognition. *Adaptive Behavior*, 19(3):155–171.

Camerer, C. and Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk & Uncertainty*, 5:325–370.

Campbell, J. O. (2016). Universal darwinism as a process of bayesian inference. *Frontiers in Systems Neuroscience*, 10(49).

Cannon, W. B. (1914). The emergency function of the adrenal medulla in pain and the major emotions. *American Journal of Physiology*, 33(2):356–372.

Cannon, W. B. (1915). *Bodily changes in pain, hunger, fear and rage: An account of recent researches into the function of emotional excitement*. New York, NY: D. Appleton and Company.

Cannon, W. B. (1929). Organization for physiological homeostasis. *Physiological Reviews*, 9(3):399–431.

Cannon, W. B. (1939). *The wisdom of the body: Revised and enlarged edition*. New York: W. W. Norton & Company, Inc.

Carruthers, P. (2004). On being simple minded. *American Philosophical Quarterly*, 41(3):205–220.

Clark, A. (2015). Radical predictive processing. *Southern Journal of Philosophy*, 53:3—27.

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

Clark, A. (2017). How to knit your own markov blanket: Resisting the second law with metamorphic minds. In Metzinger, T. and Wiese, W., editors, *Philosophy and Predictive Processing*, chapter 3, pages 1–19. Frankfurt am Main: MIND Group.

Clark, A. (2018). A nice surprise? predictive processing and the active pursuit of novelty. *Phenomenology & the Cognitive Sciences*, 17(3):521–534.

Clark, A. and Grush, R. (1999). Toward a cognitive robotics. *Adaptive Behavior*, 7(1):5–16.

Conant, R. C. and Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2):89–97.

Corcoran, A. W. (2019). Cephalopod molluscs, causal models, and curious minds. *Animal Sentience*, 4(26):13.

Corcoran, A. W. and Hohwy, J. (2018). Allostasis, interoception, and the free energy principle: Feeling our way forward. In Tsakiris, M. and De Preester, H., editors, *The interoceptive mind: From homeostasis to awareness*, chapter 15, pages 272–292. Oxford: Oxford University Press.

Corcoran, A. W., Pezzulo, G., and Hohwy, J. (2018). Commentary: Respiration-entrained brain rhythms are global but often overlooked. *Frontiers in Systems Neuroscience*, 12:25.

Craig, A. D. (2009). How do you feel – now? the anterior insula and human awareness. *Nature Reviews Neuroscience*, 10(1):59–70.

Craik, K. (1943). *The nature of explanation*. Cambridge: Cambridge University Press.

Crauel, H. and Flandoli, F. (1994). Attractors for random dynamical systems. *Probability Theory and Related Fields*, 100:365–393.

Critchley, H. D. and Harrison, N. A. (2013). Visceral influences on brain and behavior. *Neuron*, 77(4):624–638.

Dampney, R. A. L. (2016). Central neural control of the cardiovascular system: Current perspectives. *Advances in Physiology Education*, 40(3):283–296.

Daw, N. D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12):1704–1711.

de Berker, A. O., Rutledge, R. B., Mathys, C., Marshall, L., Cross, G. F., Dolan, R. J., and Bestmann, S. (2016). Computations of uncertainty mediate acute stress responses in humans. *Nature Communications*, 7:10996.

de Vries, B. and Friston, K. J. (2017). A factor graph description of deep temporal active inference. *Frontiers in Computational Neuroscience*, 11(95).

Degaute, J. P., van de Borne, P., Linkowski, P., and Van Cauter, E. (1991). Quantitative analysis of the 24-hour blood pressure and heart rate patterns in young men. *Hypertension*, 18(2):199–210.

Dennett, D. C. (1987). *The intentional stance.* MIT Press.

Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life.* London: Penguin Books Ltd.

Dewey, J. (1929). *Experience and nature.* London: George Allen & Unwin, Ltd.

Dolan, R. J. and Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2):312–325.

Dunlap, A. S. and Stephens, D. W. (2016). Reliability, uncertainty, and costs in the evolution of animal learning. *Current Opinion in Behavioral Sciences*, 12:73–79.

Dworkin, B. R. (1993). *Learning and physiological regulation.* Chicago, IL: University of Chicago Press.

Dyar, K. A., Lutter, D., Artati, A., Ceglia, N. J., Liu, Y., Armenta, D., Jastroch, M., Schneider, S., de Mateo, S., Cervantes, M., Abbondante, S., Tognini, P., Orozco-Solis, R., Kinouchi, K., Wang, C., Swerdloff, R., Nadeef, S., Masri, S., Magistretti, P., Orlando, V., Borrelli, E., Uhlenhaut, N. H., Baldi, P., Adamski, J., Tschöp, M. H., Eckel-Mahan, K., and Sassone-Corsi, P. (2018). Atlas of circadian metabolism reveals system-wide coordination and communication between clocks. *Cell*, 174(6):1571–1585.

Elias, P. (1955). Predictive coding – part i. *IRE Transactions – Information Theory*, 1(1):16–24.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75(4):643–669.

Evans, D. J. and Searles, D. J. (1994). Equilibrium microstates which generate second law violating steady states. *Physical review. E*, 50(2):1645–1648.

Evans, D. J. and Searles, D. J. (2002). The fluctuation theorem. *Advances in Physics*, 51(7):1529–1585.

Falke, J. J., Bass, R. B., Butler, S. L., Chervitz, S. A., and Danielson, M. A. (1997). The two-component signaling pathway of bacterial chemotaxis: A molecular view of signal transduction by receptors, kinases, and adaptation enzymes. *Annual Review of Cell & Developmental Biology*, 13:457–512.

Fernö, A., Pitcher, T. J., Melle, W., Nøttestad, L., Mackinson, S., Hollingworth, C., and Misund, O. A. (1998). The challenge of the herring in the norwegian sea: Making optimal collective spatial decisions. *Sarsia*, 83(2):149–167.

Feynman, R. P. (1972). *Statistical mechanics: A set of lectures*. Reading, M.A. : W. A. Benjamin, Inc.

FitzGerald, T. H. B., Dolan, R. J., and Friston, K. J. (2014). Model averaging, optimal inference, and habit formation. *Frontiers in Human Neuroscience*, 8(457):1–11.

FitzGerald, T. H. B., Dolan, R. J., and Friston, K. J. (2015). Dopamine, reward learning, and active inference. *Frontiers in Computational Neuroscience*, 9(136):1–16.

Fotopoulou, A. and Tsakiris, M. (2017). Mentalizing homeostasis: The social origins of interoceptive inference. *Neuropsychoanalysis*, pages 1–26.

Freddolino, P. L. and Tavazoie, S. (2012). Beyond homeostasis: A predictive-dynamic framework for understanding cellular behavior. *Annual Review of Cell & Development*, 28:363–384.

Friston, K. J. (2002). Functional integration and inference in the brain. *Progress in Neurobiology*, 68(2):113–143.

Friston, K. J. (2003). Learning and inference in the brain. *Neural Networks*, 16(9):1325–1352.

Friston, K. J. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B*, 360(1456):815–836.

Friston, K. J. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4(11):e1000211.

Friston, K. J. (2009). The free-energy principle: A rough guide to the brain? *Trends in Cognitive Sciences*, 13(7):293–301.

Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138.

Friston, K. J. (2011). Embodied inference: Or ”i think therefore i am, if i am what i think”. In Tschacher, W. and Bergomi, C., editors, *The implications of embodiment: Cognition and communication*, pages 89–125. Exeter: Imprint Academic.

Friston, K. J. (2012a). A free energy principle for biological systems. *Entropy*, 14(11):2100–2121.

Friston, K. J. (2012b). Policies and priors. In Gutkin, B. and Ahmed, S. H., editors, *Computational neuroscience of drug addiction*, Springer series in computational neuroscience 10, pages 237—283. New York, NY: Springer Science + Business Media.

Friston, K. J. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86):20130475.

Friston, K. J. (2017). Self-evidencing babies: Commentary on “mentalizing homeostasis: The social origins of interoceptive inference” by fotopoulou & tsakiris. *Neuropsychoanalysis*, 19(1):43—47.

Friston, K. J. (2018). Am i self-conscious? (or does self-organisation entail self-consciousness?). *Frontiers in Psychology*, 9(579).

Friston, K. J., Adams, R. A., and Montague, R. (2012a). What is value – accumulated reward or evidence? *Frontiers in Neurorobotics*, 6:11.

Friston, K. J., Adams, R. A., Perrinet, L., and Breakspear, M. (2012b). Perceptions as hypotheses: Saccades as experiments. *Frontiers in Psychology*, 3(151):1–20.

Friston, K. J. and Ao, P. (2012). Free energy, value, and attractors. *Computational & Mathematical Methods in Medicine*, page 937860.

Friston, K. J., Breakspear, M., and Deco, G. (2012c). Perception and self-organized instability. *Frontiers in Computational Neuroscience*, 6(44):1–19.

Friston, K. J., Daunizeau, J., and Kiebel, S. J. (2009). Reinforcement learning or active inference? *PLoS One*, 4(7):e6421.

Friston, K. J., Daunizeau, J., Kilner, J., and Kiebel, S. J. (2010a). Action and behavior: A free-energy formulation. *Biological Cybernetics*, 102(3):227–260.

Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J., and Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68:862–879.

Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017a). Active inference: A process theory. *Neural Computation*, 29(1):1–49.

Friston, K. J. and Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B*, 364(1521):1211–1221.

Friston, K. J., Kilner, J., and Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology-Paris*, 100(1-3):70—87.

Friston, K. J., Levin, M., Sengupta, B., and Pezzulo, G. (2015a). Knowing one's place: A free-energy approach to pattern regulation. *Journal of the Royal Society Interface*, 12(20141383):1–12.

Friston, K. J., Lin, M., Frith, C. D., Pezzulo, G., Hobson, J. A., and Ondobaka, S. (2017b). Active inference, curiosity and insight. *Neural Computation*, 29(10):2633–2683.

Friston, K. J., Mattout, J., Trujillo-Barreto, N., Ashburner, J., and Penny, W. D. (2007). Variational free energy and the laplace approximation. *NeuroImage*, 34(1):220–234.

Friston, K. J., Parr, T., and de Vries, B. (2017c). The graphical brain: Belief propagation and active inference. *Network Neuroscience*, 1(4):381–414.

Friston, K. J., Parr, T., and Zeidman, P. (2018). Bayesian model reduction. *arXiv*, page 1805.07092.

Friston, K. J., Rigoli, F., Ognibene, D., Mathys, C. D., Fitzgerald, T., and Pezzulo, G. (2015b). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4):187–224.

Friston, K. J., Rosch, R., Parr, T., Price, C., and Bowman, H. (2017d). Deep temporal models and active inference. *Neuroscience & Biobehavioral Reviews*, 77:388–402.

Friston, K. J., Samothrakis, S., and Montague, R. (2012d). Active inference and agency: optimal control without cost functions. *Biological Cybernetics*, 106(8-9):523–541.

Friston, K. J., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., and Dolan, R. J. (2013). The anatomy of choice: Active inference and agency. *Frontiers in Human Neuroscience*, 7(598):1–18.

Friston, K. J. and Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, 159(3):417–458.

Friston, K. J., Stephan, K. E., Li, B., and Daunizeau, J. (2010b). Generalised filtering. *Mathematical Problems in Engineering*, 3:621670.

Friston, K. J., Thornton, C., and Clark, A. (2012e). Free-energy minimization and the dark-room problem. *Frontiers in Psychology*, 3:130.

Fuster, J. M. (2001). The prefrontal cortex – an update: Time is of the essence. *Neuron*, 30(2):319–333.

Fuster, J. M. (2004). Upper processing stages of the perception-action cycle. *Trends in Cognitive Sciences*, 8(4):143–145.

Gagliano, M. (2015). In a green frame of mind: Perspectives on the behavioural ecology and cognitive nature of plants. *AoB Plants*, 7:plu075.

Gärdenfors, P. (1995). Cued and detached representations in animal cognition. *Behavioural Processes*, 35:263–273.

Ginty, A. T., Kraynak, T. E., Fisher, J. P., and Gianaros, P. J. (2017). Cardiovascular and autonomic reactivity to psychological stress: Neurophysiological substrates and links to cardiovascular disease. *Autonomic Neuroscience: Basic & Clinical*, 207:2–9.

Giurfa, M. (2013). Cognition with few neurons: Higher-order learning in insects. *Trends in Neurosciences*, 36(5):285–294.

Godfrey-Smith, P. (1995). Spencer and dewey on life and mind. In Boden, M. A., editor, *The philosophy of artificial life*, Oxford Readings in Philosophy, chapter 12, pages 314–331. Oxford: Oxford University Press.

Godfrey-Smith, P. (1996). *Complexity and the function of mind in nature*. Cambridge Studies in Philosophy and Biology. Cambridge: Cambridge University Press.

Godfrey-Smith, P. (2002). Environmental complexity and the evolution of cognition. In Sternberg, R. J. and Kaufman, J. C., editors, *The evolution of intelligence*, chapter 10, pages 223–250. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Godfrey-Smith, P. (2016a). Individuality, subjectivity, and minimal cognition. *Biology & Philosophy*, 31(6):775–796.

Godfrey-Smith, P. (2016b). Mind, matter, and metabolism. *Journal of Philosophy*, 113(10):481–506.

Goodwin, G. M., McCloskey, D. I., and Mitchell, J. H. (1972). Cardiovascular and respiratory responses to changes in central command during isometric exercise at constant muscle tension. *Journal of Physiology*, 226(1):173–190.

Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral & Brain Sciences*, 27(3):377–442.

Gu, X., Hof, P. R., Friston, K. J., and Fan, J. (2013). Anterior insular cortex and emotional awareness. *Journal of Comparative Neurology*, 521(15):3371–3388.

Hennessey, T. M., Rucker, W. B., and McDiarmid, C. G. (1979). Classical conditioning in paramecia. *Animal Learning & Behavior*, 7(4):417–423.

Heschl, A. (1990). L = c: A simple equation with astonishing consequences. *Journal of Theoretical Biology*, 145:13–40.

Hobson, J. A. and Friston, K. J. (2012). Waking and dreaming consciousness: Neurobiological and functional considerations. *Progress in Neurobiology*, 98(1):82–98.

Hohwy, J. (2013). *The predictive mind.* Oxford: Oxford University Press.

Hohwy, J. (2016). The self-evidencing brain. *Noûs*, 50(2):259—285.

Hohwy, J. (2017a). How to entrain your evil demon. In Metzinger, T. and Wiese, W., editors, *Philosophy and Predictive Processing*, chapter 2, pages 1–15. Frankfurt am Main: MIND Group.

Hohwy, J. (2017b). Priors in perception: Top-down modulation, bayesian perceptual learning rate, and prediction error minimization. *Consciousness & Cognition*, 47:75–85.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., and Camerer, C. F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science*, 310(5754):1680–1683.

Huang, Y. and Rao, R. P. N. (2011). Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5):580–593.

Huettel, S. A., Stowe, C. J., Gordon, E. M., Warner, B. T., and Platt, M. L. (2006). Neural signatures of economic preferences for risk and ambiguity. *Neuron*, 49(5):765–775.

Kabadayi, C. and Osvath, M. (2017). Ravens parallel great apes in flexible planning for tool-use and bartering. *Science*, 357(6347):202–204.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45.

Kalman, R. E. and Bucy, R. S. (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83(1):95–108.

Kanai, R., Komura, Y., Shipp, S., and Friston, K. J. (2015). Cerebral hierarchies: Predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B*, 370(1668):69–81.

Kaplan, R. and Friston, K. J. (2018). Planning and navigation as active inference. *Biological Cybernetics*, 112(4):323–343.

Keramati, M. and Gutkin, B. (2014). Homeostatic reinforcement learning for integrating reward collection and physiological stability. *eLife*, 3(e04811):1–26.

Khalsa, S. S., Adolphs, R., Cameron, O. G., Critchley, H. D., Davenport, J. S., Feinstein, J. S., Feusner, J. D., Garfinkel, S. N., Lane, R. D., Mehling, W. E., Meuret, A. E., Nemeroff, C. B., Oppenheimer, S., Petzschner, F. H., Pollatos, O., Rhudy, J. L., Schramm, L. P., Simmons, W. K., Stein, M. B., Stephan, K. E., Van Den Bergh, O., Van Diest, I., von Leupoldt, A., and Paulus, M. P. (2018). Interoception and mental health: A roadmap. *Biological Psychiatry: Cognitive Neuroscience & Neuroimaging*, 3:501–513.

Kiebel, S. J., Daunizeau, J., and Friston, K. J. (2008). A hierarchy of time-scales and the brain. *PLoS Computational Biology*, 4(11):e1000209.

Kirchhoff, M., Parr, T., Palacios, E., Friston, K. J., and Kiverstein, J. (2018). The markov blankets of life: Autonomy, active inference and the free energy principle. *Journal of The Royal Society Interface*, 15(138):20170792.

Kirchhoff, M. D. and Froese, T. (2017). Where there is life there is mind: In support of a strong life-mind continuity thesis. *Entropy*, 19(4):169.

Knight, F. H. (1921). *Risk, uncertainty, and profit*. New York, NY: Sentry Press.

Kozyreva, A. and Hertwig, R. (2019). The interpretation of uncertainty in ecological rationality. *Synthese*, pages 1–31.

Kräuchi, K. and Wirz-Justice, A. (1994). Circadian rhythm of heat production, heart rate, and skin and core temperature under unmasking conditions in men. *American Journal of Physiology*, 267(3 Pt 2):R819–R829.

Krogh, A. and Lindhard, J. (1913). The regulation of respiration and circulation during the initial stages of muscular work. *Journal of Physiology*, 47:112–136.

Krupenye, C., Kano, F., Hirata, S., Call, J., and Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354(6308):110–114.

Lee, D., McGreevy, B. P., and Barraclough, D. J. (2005). Learning and decision making in monkeys during a rock-paper-scissors game. *Cognitive Brain Research*, 25(2):416–430.

Lee, T. S. and Mumford, D. (2003). Hierarchical bayesian inference in the visual cortex. *Journal of the Optical Society of America A*, 20(7):1434–1448.

Levin, M., Pezzulo, G., and Finkelstein, J. M. (2017). Endogenous bioelectric signaling networks: Exploiting voltage gradients for control of growth and form. *Annual Review of Biomedical Engineering*, 19:353–387.

Levy, I., Snell, J., Nelson, A. J., Rustichini, A., and Glimcher, P. W. (2010). Neural representation of subjective value under risk and ambiguity. *Journal of Neurophysiology*, 103(2):1036–1047.

Lewis, D. (1973a). Causation. *Journal of Philosophy*, 70(17):556–567.

Lewis, D. (1973b). *Counterfactuals*. Oxford: Basil Blackwell Ltd.

Lewis, D. (1979). Counterfactual dependence and time's arrow. *Noûs*, 13:455–476.

Limanowski, J. and Friston, K. J. (2018). 'seeing the dark': Grounding phenomenal transparency and opacity in precision estimation for active inference. *Frontiers in Psychology*, 9:643.

Linson, A., Clark, A., Ramamoorthy, S., and Friston, K. J. (2018). The active inference approach to ecological perception: General information dynamics for natural and artificial embodied cognition. *Frontiers in Robotics & AI*, 5(21).

Lyon, P. (2015). The cognitive cell: Bacterial behavior reconsidered. *Frontiers in Microbiology*, 6:264.

Mackie, G. O. and Burighel, P. (2005). The nervous system in adult tunicates: Current research directions. *Canadian Journal of Zoology*, 83:151–183.

Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., and Stephan, K. E. (2014). Uncertainty in perception and the hierarchical gaussian filter. *Frontiers in Human Neuroscience*, 8(825):1–24.

Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. Dordrecht: D. Reidel Publishing Company.

McCoy, J. W. (1977). Complexity in organic evolution. *Journal of Theoretical Biology*, 68(3):457–488.

McEwen, B. S. and Stellar, E. (1993). Stress and the individual: Mechanisms leading to disease. *Archives of Internal Medicine*, 153(18):2093–2101.

McGregor, S., Baltieri, M., and Buckley, C. L. (2015). A minimal active inference agent. *arXiv*, 1503.04187.

Menaker, M., Murphy, Z. C., and Sellix, M. T. (2013). Central control of peripheral circadian oscillators. *Current Opinion in Neurobiology*, 23(5):741–746.

Menon, V. and Uddin, L. Q. (2010). Saliency, switching, attention and control: A network model of insula function. *Brain Structure & Function*, 214(5-6):655–667.

Metzinger, T. (2017). The problem of mental action: Predictive control without sensory sheets. In Metzinger, T. and Wiese, W., editors, *Philosophy and Predictive Processing*, chapter 19, pages 1–26. Frankfurt am Main: MIND Group.

Mikhalevich, I., Powell, R., and Logan, C. (2017). Is behavioural flexibility evidence of cognitive complexity? how evolution can inform comparative cognition. *Interface Focus*, 7(3):20160121.

Miracchi, L. (2019). A competence framework for artificial intelligence research. *Philosophical Psychology*, 32(5):588–633.

Mirza, M. B., Adams, R. A., Mathys, C. D., and Friston, K. J. (2016). Scene construction, visual foraging, and active inference. *Frontiers in Computational Neuroscience*, 10(56):1–16.

Mitchell, A., Romano, G. H., Groisman, B., Yona, A., Dekel, E., Kupiec, M., Dahan, O., and Pilpel, Y. (2009). Adaptive prediction of environmental changes by microorganisms. *Nature*, 460(7252):220–224.

Moore, B. R. (2004). The evolution of learning. *Biological Reviews*, 79(2):301–335.

Moran, R. J., Symmonds, M., Dolan, R. J., and Friston, K. J. (2014). The brain ages optimally to model its environment: Evidence from sensory learning over the adult lifespan. *PLoS Computational Biology*, 10(1):e1003422.

Moreno, A. and Etxeberria, A. (2005). Agency in natural and artificial systems. *Artificial Life*, 11:161–175.

Morgan, A. (2018a). Mindless accuracy: On the ubiquity of content in nature. *Synthese*, 195(12):5403–5429.

Morgan, A. (2018b). Pictures, plants, and propositions. *Minds & Machines*.

Morville, T., Friston, K. J., Burdakov, D., Siebner, H. R., and Hulme, O. J. (2018). The homeostatic logic of reward. *bioRxiv*.

Mugan, U. and MacIver, M. A. (2019). The shift from life in water to life on land advantaged planning in visually-guided behavior. *bioRxiv*.

Nakajima, M., Imai, K., Ito, H., Nishiwaki, T., Murayama, Y., Iwasaki, H., Oyama, T., and Kondo, T. (2005). Reconstitution of circadian oscillation of cyanobacterial kaic phosphorylation in vitro. *Science*, 308(5720):414–415.

Neill, W. H. (1979). Mechanisms of fish distribution in heterothermal environments. *American Zoologist*, 19(1):305–317.

Nicolis, G. and Prigogine, I. (1977). *Self-organization in nonequilibrium systems: From dissipative structures to order through fluctuations*. New York, NY: John Wiley & Sons.

Nute, D. (1975). Counterfactuals. *Notre Dame Journal of Formal Logic*, 16(4):476–482.

Owens, A. P., Allen, M., Ondobaka, S., and Friston, K. J. (2018). Interoceptive inference: From computational neuroscience to clinic. *Neuroscience & Biobehavioral Reviews*, 90:174–183.

Palacios, E. R., Razi, A., Parr, T., Kirchhoff, M. D., and Friston, K. J. (2017). Biological self-organisation and markov blankets. *bioRxiv*.

Palmer, C. J., Seth, A. K., and Hohwy, J. (2015). The felt presence of other minds: Predictive processing, counterfactual predictions, and mentalising in autism. *Consciousness & Cognition*, 36:376–389.

Parr, T., Corcoran, A. W., Friston, K. J., and Hohwy, J. (2019). Perceptual awareness and active inference. *Neuroscience of Consciousness*, 5(1):niz012.

Parr, T. and Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of the Royal Society Interface*, 14(20170376):1–10.

Parr, T. and Friston, K. J. (2018a). The anatomy of inference: Generative models and brain structure. *Frontiers in Computational Neuroscience*, 12:90.

Parr, T. and Friston, K. J. (2018b). The discrete and continuous brain: From decisions to movement–and back again. *Neural Computation*, 30:1–29.

Paulus, M. P. and Stein, M. B. (2006). An insular view of anxiety. *Biological Psychiatry*, 60(4):383–387.

Pavlov, I. P. (1902). *The work of the digestive glands*. London: Charles Griffin & Co., Ltd.

Payzan-LeNestour, E. and Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Computational Biology*, 7(1):e1001048.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo: Morgan Kaufmann Publishers.

Penny, W. and Stephan, K. (2014). A dynamic bayesian model of homeostatic control. *Lecture Notes in Computer Science*, 8779:60–69.

Penny, W. D., Zeidman, P., and Burgess, N. (2013). Forward and backward inference in spatial cognition. *PLoS Computational Biology*, 9(12):e1003383.

Perry, C. J., Barron, A. B., and Cheng, K. (2013). Invertebrate learning and cognition: Relating phenomena to neural substrate. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(5):561–582.

Peters, A., McEwen, B. S., and Friston, K. J. (2017). Uncertainty and stress: Why it causes diseases and how it is mastered by the brain. *Progress in Neurobiology*, 156:164–188.

Petzschner, F. H., Weber, L. A. E., Gard, T., and Stephan, K. E. (2017). Computational psychosomatics and computational psychiatry: Toward a joint framework for differential diagnosis. *Biological Psychiatry*, 82:421–430.

Pezzulo, G. (2008). Coordinating with the future: The anticipatory nature of representation. *Minds & Machines*, 18(2):179–225.

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.

Pezzulo, G. (2017). Tracing the roots of cognition in predictive processing. In Metzinger, T. and Wiese, W., editors, *Philosophy and Predictive Processing*, chapter 20, pages 1–20. Frankfurt am Main: MIND Group.

Pezzulo, G., Cartoni, E., Rigoli, F., Pio-Lopez, L., and Friston, K. J. (2016). Active inference, epistemic value, and vicarious trial and error. *Learning & Memory*, 23(7):322–338.

Pezzulo, G. and Castelfranchi, C. (2007). The symbol detachment problem. *Cognitive Processing*, 8(2):115–131.

Pezzulo, G. and Castelfranchi, C. (2009). Thinking as the control of imagination: A conceptual framework for goal-directed systems. *Psychological Research*, 73(4):559–577.

Pezzulo, G. and 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.

Pezzulo, G., Kemere, C., and van der Meer, M. A. A. (2017). Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition. *Annals of the New York Academy of Sciences*, 1396(1):144–165.

Pezzulo, G., Rigoli, F., and Friston, K. J. (2015). Active inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, 134:17–35.

Pezzulo, G., Rigoli, F., and Friston, K. J. (2018). Hierarchical active inference: A theory of motivated control. *Trends in Cognitive Sciences*, 22(4):294–306.

Powers, W. T. (1973). Feedback: Beyond behaviorism. *Science*, 179(4071):351–356.

Preuschoff, K., Quartz, S. R., and Bossaerts, P. (2008). Human insula activation reflects risk prediction errors as well as risk. *Journal of Neuroscience*, 28(11):2745–2752.

Quadt, L., Critchley, H. D., and Garfinkel, S. N. (2018). The neurobiology of interoception in health and disease. *Annals of the New York Academy of Sciences*, 1428(1):112–128.

Raby, C. R., Alexis, D. M., Dickinson, A., and Clayton, N. S. (2007). Planning for the future by western scrub-jays. *Nature*, 445(7130):919–921.

Ramsay, D. S. and Woods, S. C. (2014). Clarifying the roles of homeostasis and allostasis in physiological regulation. *Psychological Review*, 121(2):225–247.

Ramsay, D. S. and Woods, S. C. (2016). Physiological regulation: How it really works. *Cell Metabolism*, 24(3):361–364.

Ramstead, M. J. D., Badcock, P. B., and Friston, K. J. (2018). Answering schrödinger's question: A free-energy formulation. *Physics of Life Reviews*, 24:1–16.

Rao, R. P. N. and Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87.

Read, C. R., Garnier, S., Beekman, M., and Latty, T. (2015). Information integration and multiattribute decision making in non-neuronal organisms. *Animal Behaviour*, 100:44–50.

Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, 17(3):147–159.

Redshaw, J. and Bulley, A. (2018). Future-thinking in animals: Capacities and limits. In Oettingen, G., Sevincer, A. T., and Gollwitzer, P. M., editors, *The psychology of thinking about the future*, chapter 2, pages 31–51. New York, NY: The Guilford Press.

Requin, J., Brener, J., and Ring, C. (1991). Preparation for action. In Jennings, J. R. and Coles, M. G. H., editors, *Handbook of cognitive psychophysiology: Central and* *autonomic nervous system approaches*, chapter 4, pages 357–448. John Wiley & Sons Ltd.

Rust, M. J., Markson, J. S., Lane, W. S., Fisher, D. S., and O'Shea, E. K. (2007). Ordered phosphorylation governs oscillation of a three-protein circadian clock. *Science*, 318(5851):809–812.

Sales, A. C., Friston, K. J., Jones, M. W., Pickering, A. E., and Moran, R. J. (2019). Locus coeruleus tracking of prediction errors optimises cognitive flexibility: An active inference model. *PLoS Computational Biology*, 15(1):e1006267.

Salman, H. and Libchaber, A. (2007). A concentration-dependent switch in the bacterial response to temperature. *Nature Cell Biology*, 9(9):1098–1100.

Sanchez-Fibla, M., Bernardet, U., Wasserman, E., Pelc, T., Mintz, M., Jackson, J. C., Pennartz, C. M. A., and Verschure, P. F. M. J. (2010). Allostatic control for robot behavior regulation: A comparative rodent-robot study. *Advances in Complex Systems*, 13(3):377–403.

Schacter, D. L. and Addis, D. R. (2007). The cognitive neuroscience of constructive memory: Remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481):773–786.

Schrödinger, E. (1992). *What is life? with "Mind and matter" and "Autobiographical sketches"*. Cambridge: Cambridge University Press.

Schulkin, J. and Sterling, P. (2019). Allostasis: A brain-centered, predictive mode of physiological regulation. *Trends in Neurosciences*, 42(10):740–752.

Schwartenbeck, P., FitzGerald, T., Dolan, R. J., and Friston, K. J. (2013). Exploration, novelty, surprise, and free energy minimization. *Frontiers in Psychology*, 4(710):1–5.

Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C. D., Dolan, R., Kronbichler, M., and Friston, K. J. (2015). Evidence for surprise minimization over value maximization in choice behavior. *Scientific Reports*, 5(16575):1–14.

Schwartenbeck, P., Passecker, J., Hauser, T. U., FitzGerald, T. H. B., Kronbichler, M., and Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration. *eLife*, 8.

Segundo-Ortin, M. and Calvo, P. (2019). Are plants cognitive? a reply to adams. *Studies in History & Philosophy of Science*, 73:64–71.

Seifert, U. (2012). Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75(12):126001.

Sengupta, B., Stemmler, M. B., and Friston, K. J. (2013). Information and efficiency in the nervous system – a synthesis. *PLoS Computational Biology*, 9(7):e1003157.

Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, 17(11):565–573.

Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5(2):97–118.

Seth, A. K. (2015). The cybernetic bayesian brain: From interoceptive inference to sensorimotor contingencies. In Metzinger, T. and Windt, J. M., editors, *Open MIND*, pages 1–24. Frankfurt am Main: MIND Group.

Seth, A. K. and Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Philosophical Transactions of the Royal Society B*, 371(1708):1–10.

Seth, A. K., Suzuki, K., and Critchley, H. D. (2012). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, 2(395):1–16.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, 27(3):379–423.

Shipp, S. (2016). Neural elements for predictive coding. *Frontiers in Psychology*, 7(1792):1–21.

Shipp, S., Adams, R. A., and Friston, K. J. (2013). Reflections on agranular architecture: Predictive coding in the motor cortex. *Trends in Cognitive Sciences*, 36(12):706–716.

Smith, G. P. (2000). Pavlov and integrative physiology. *American Journal of Physiology - Regulatory, Integrative and Comparative Physiology*, 279(3):R743–R755.

Smith, R., Thayer, J. F., Khalsa, S. S., and Lane, R. D. (2017). The hierarchical basis of neurovisceral integration. *Neuroscience & Biobehavioral Reviews*, 75:274–296.

Smith-Ferguson, J. and Beekman, M. (2019). Who needs a brain? slime moulds, behavioural ecology and minimal cognition. *Adaptive Behavior*.

Solway, A. and Botvinick, M. M. (2012). Goal-directed decision making as probabilistic inference: A computational framework and potential neural correlates. *Psychological Review*, 119(1):120–154.

Spencer, H. (1867). *First principles*. London: Williams & Norgate, 2nd edition.

Spratling, M. W. (2017). A review of predictive coding algorithms. *Brain & Cognition*, 112:92–97.

Sprigge, T. L. S. (1970). *Facts, words and beliefs*. London: Routledge & Keegan Paul.

Srinivasan, M. V., Laughlin, S. B., and Dubs, A. (1982). Predictive coding: A fresh view of inhibition in the retina. *Proceedings of the Royal Society B*, 216(1205):427–459.

Stalnaker, R. C. (1968). A theory of conditionals. In Rescher, N., editor, *Studies in logical theory*, American Philosophical Quarterly supplementary monograph series, pages 98–112. Oxford: Basil Blackwell Ltd.

Stanley, M. L., Stewart, G. W., and De Brigard, F. (2017). Counterfactual plausibility and comparative similarity. *Cognitive Science*, 41(Suppl 5):1216–1228.

Steiner, A. P. and Redish, A. D. (2014). Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nature Neuroscience*, 17(7):995–1002.

Stephan, K. E., Manjaly, Z. M., Mathys, C. D., Weber, L. A. E., Paliwal, S., Gard, T., Tittgemeyer, M., Fleming, S. M., Haker, H., Seth, A. K., and Petzschner, F. H. (2016). Allostatic self-efficacy: A metacognitive theory of dyshomeostasis-induced fatigue and depression. *Frontiers in Human Neuroscience*, 10(550):1–27.

Sterelny, K. (2003). *Thought in a hostile world: The evolution of human cognition*. Malden, MA: Blackwell Publishing.

Sterling, P. (2004). Principles of allostasis: Optimal design, predictive regulation, pathophysiology and rational therapeutics. In Schulkin, J., editor, *Allostasis, homeostasis, and the costs of physiological adaptation*, chapter 1, pages 17–64. Cambridge: Cambridge University Press.

Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology & Behavior*, 106(1):5–15.

Sterling, P. and Eyer, J. (1988). Allostasis: A new paradigm to explain arousal pathology. In Fisher, S. and Reason, J., editors, *Handbook of life stress, cognition and health*, chapter 34, pages 629–649. John Wiley & Sons Ltd.

Suddendorf, T., Bulley, A., and Miloyan, B. (2018). Prospection and natural selection. *Current Opinion in Behavioral Sciences*, 24:26–31.

Suddendorf, T. and Corballis, M. C. (1997). Mental time travel and the evolution of the human mind. *Genetic Social & General Psychology Monographs*, 123(2):133–167.

Suddendorf, T. and Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral & Brain Sciences*, 30(3):299–313.

Suddendorf, T. and Redshaw, J. (2017). Anticipation of future events. In Vonk, J. and Shackelford, T. K., editors, *Encyclopedia of animal cognition and behavior*. Springer International Publishing.

Sweis, B. M., Thomas, M. J., and Redish, A. D. (2018). Mice learn to avoid regret. *PLoS Biology*, 16(6):e2005853.

Tagkopoulos, I., Liu, Y.-C., and Tavazoie, S. (2008). Predictive behavior within microbial genetic networks. *Science*, 320(5881):1313–1317.

Tang, S. K. Y. and Marshall, W. F. (2018). Cell learning. *Current Biology*, 28(20):R1180–R1184.

Teff, K. L. (2011). How neural mediation of anticipatory and compensatory insulin release helps us tolerate food. *Physiology & Behavior*, 103(1):44–50.

Thompson, E. (2007). *Mind in life: Biology, phenomenology and the sciences of mind*. Cambridge, MA: Harvard University Press.

Todd, W. (1964). Counterfactual conditionals and the presuppositions of induction. *Philosophy of Science*, 31(2):101–110.

Van de Cruys, S. (2017). Affective value in the predictive mind. In Metzinger, T. and Wiese, W., editors, *Philosophy and Predictive Processing*, chapter 24, pages 1–21. Frankfurt am Main: MIND Group.

van Duijn, M., Keijzer, F., and Franken, D. (2006). Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior*, 14(2):157–170.

Verschure, P. F. M. J., Pennartz, C. M. A., and Pezzulo, G. (2014). The why, what, where, when and how of goal-directed choice: Neuronal and computational principles. *Philosophical Transactions of the Royal Society B*, 369(1655):20130483.

Vincent, P., Parr, T., Benrimoh, D., and Friston, K. J. (2019). With an eye on uncertainty: Modelling pupillary responses to environmental volatility. *PLoS Computational Biology*, 15(7):e1007126.

Wen, Y., Zhou, W., Zhu, X., Cheng, S., Xiao, G., Li, Y., Zhu, Y., Wang, Z., and Wan, C. (2015). An investigation of circadian rhythm in escherichia coli. *Biological Rhythm Research*, 46(5):753–762.

Wheeler, M. (2011). Mind in life or life in mind? making sense of deep continuity. *Journal of Consciousness Studies*, 18(5):148–168.

Wiener, N. (1961). *Cybernetics: Or control and communication in the animal and the machine*. Cambridge, MA: MIT Press, 2nd edition.

Wiese, W. (2017). Action is enabled by systematic misrepresentations. *Erkenntnis*, 82(6):1233–1252.

Wiese, W. and Metzinger, T. (2017). Vanilla pp for philosophers: A primer on predictive processing. In Metzinger, T. and Wiese, W., editors, *Philosophy and Predictive Processing*, chapter 1, pages 1–18. Frankfurt am Main: MIND Group.

Williams, D. (2018). Predictive minds and small-scale models: Kenneth craik's contribution to cognitive science. *Philosophical Explorations*, 21(2):245–263.

Williams, D. and Colling, L. (2018). From symbols to icons: The return of resemblance in the cognitive neuroscience revolution. *Synthese*, 195(5):1941–1967.

Yon, D., de Lange, F. P., and Press, C. (2019). The predictive brain as a stubborn scientist. *Trends in Cognitive Sciences*, 23(1):6–8.

Yu, A. J. and Dayan, P. (2003). Expected and unexpected uncertainty: Ach and ne in the neocortex. *Advances in Neural Information Processing Systems*, 15:157–164.

Yu, A. J. and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, 46(4):681–692.

Zénon, A., Solopchuk, O., and Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, 123:5–18.

Zwicker, D., Lubensky, D. K., and ten Wolde, P. R. (2010). Robust circadian clocks from coupled protein-modification and transcription-translation cycles. *Proceedings of the National Academy of Sciences*, 107(52):22540–22545.