Title: The Predictive Dynamics of Happiness and Well-Being

Authors:
Mark Miller [1] (markmiller@chain.hokudai.ac.jp)
Erik Rietveld [2, 3] (d.w.rietveld@amsterdamumc.nl)
Julian Kiverstein [2] (j.d.kiverstein@amsterdamumc.nl)

Affiliations:
1. Center for Human Nature, Artificial Intelligence and Neuroscience, Hokkaido University, Japan.
2. Department of Psychiatry, Amsterdam UMC, University of Amsterdam, Meibergdreef 9, Amsterdam, Netherlands.
3. Department of Philosophy, University of Twente, Ravelijn 4232, P.O. Box 217, 7500 AE Enschede, Netherlands.

Acknowledgements:
We are grateful to Abby Tabor, Mahault Albarracin, Brett Anderson, Ines Hipolito and Andy Clark for helpful comments and discussions along the way. Mark Miller carried out this work with the support of Horizon 2020 European Union ERC Advanced Grant XSPECT - DLV-692739. Julian Kiverstein and Erik Rietveld are supported by the European Research Council in the form of ERC Starting Grant 679190 (EU Horizon 2020) for the project AFFORDS-HIGHER awarded to Erik Rietveld, and by a project grant from the Amsterdam Brain and Cognition research group at the University of Amsterdam.

Abstract:
We offer an account of mental health and well-being using the Predictive Processing Framework (PPF). According to this framework, the difference between mental health and psychopathology can be located in the goodness of the predictive model as a regulator of action. What is crucial for avoiding the rigid patterns of thinking, feeling and acting associated with psychopathology is the regulation of action based on the valence of affective states. In PPF valence is modelled as error dynamics - the change in prediction errors over time. Our aim in this paper is to show how error dynamics can account for both momentary happiness and longer-term well-being. What will emerge is a new neurocomputational framework for making sense of human flourishing.

Key words: predictive processing, error dynamics, valence, happiness, reward, well-being

"*Whatever Is Flexible and Flowing Will Tend to Grow, Whatever Is Rigid and Blocked Will Wither and Die*". Tao Te Ching

Introduction

The predictive processing framework (henceforth "PPF") has recently been proposed as a unifying theory of the embodied brain and its cognitive functions (Friston 2010; Hohwy 2013; Clark 2013, 2016). The central idea behind PPF is that the embodied brain is a predictive model of the body in the world. This model is used to generate predictions of, among other things, the sensory outcomes of the organism's actions in the world. These predictions can be compared with the sensory states the organism actually visits when it acts. So long as the organism keeps the error in its predictions to a minimum over time, the organism will typically succeed in achieving the valued outcomes it aims for in acting. PPF is an increasingly influential theoretical framework for studying mental illness in computational psychiatry.[1] However, so far little attention has been given to the question of mental health and well-being from the perspective of the PPF. We take first steps in this direction in our paper.

We take as our starting point the proposal that to be mentally healthy an organism must be a good predictor of the hidden causes (environmental and bodily) of its sensory states. Such an organism will tend to behave in ways that maintain homeostasis at each moment in time. While we take this to be an important part of the story, we use the example of substance addiction to explain why moment by moment prediction-error minimisation is probably not sufficient for mental well-being (Miller, Kiverstein & Rietveld 2020). What goes wrong in addiction offers clues about what it means for a human being to be well. Addiction is an example of a sub-optimal strategy that an agent can pursue for reducing prediction error (i.e., bring about a certain goal) by relying on an over-learned, habitual form of behaviour. What is harmful to an agent here is the ways in which drugs of addiction engender a rigidity in thinking and acting (see Miller et al. 2020, Barrett & Simmons 2017). Drugs of addiction do this by acting on dopaminergic systems that strengthen drug seeking and using policies at the expense of alternatives.

We will argue that what is crucial for avoiding pathological forms of rigidity as seen in addiction is *managing* error based on sensitivity to *error dynamics* - the change in the rate of error reduction (Joffily & Corricelli 2013; Van de Cruys 2017; Kiverstein et al. 2017). Change in rate means that error reduction has either unfolded worse or better than expected. When agents do

[1] See e.g. Fletcher & Frith 2009; Seth, Suzuki & Critchley 2012; Edwards et al 2014; Friston et al 2014; Corlett & Fletcher 2014; Seth & Friston 2016; Barrett et al 2017; Badcock et al 2017; Smith et al. 2020; Linson & Friston 2019; Paulus et al. 2019; Barca & Pezzulo 2020; Ciaunica et al. 2020; Gerrans & Gadsby 2020; Prosser 2018; Hipolito et al. forthcoming;; Fotopolou et al. 201x; Linson & Friston 2019; Wilkinson (delusions); Linson, Parr & Friston 2020; Parr, Rees & Friston 2018; Adams et al. 2013; Schwartenbeck & Friston 2016; Smith, Badcock & Friston 2020; Ramsteade et al. forthcoming; Miller et al. 2020; Kiverstein et al. 2020; Dean, Miller & Wilkinson 2020).

better or worse than expected this is registered in the body as positively or negatively valenced affect. Agents that use these affective states to regulate their behaviour will be driven to continuously make progress in error reduction. As we will show, this will require them to sometimes disrupt their own habits of thinking and acting in ways that temporarily lead to increases in error and uncertainty but that in the long-run allow them to make progress in learning. That is, they will sometimes perform actions that temporarily lead to an increase in uncertainty if doing so will help them to do better at reducing error in the long-run. This is precisely what does not happen in long term addicts. What turns out to be important for optimality is being attuned to opportunities for making progress in error reduction. We will characterise this attunement in terms of metastable dynamics[2] or what we will call *metastable attunement*. We will argue that metastable attunement is conducive to well-being because it allows an agent to remain in touch with and integrate their various cares and concerns over a life-time.

1. A Predictive Processing Account of Mental Health

The definition of well-being has proven to be controversial among psychologists. The debate has turned upon different conceptions of "the good life", with some psychologists favouring a hedonic view that understands well-being to consist of a life of positive experiences such as pleasure and happiness (Kahneman et al. 1999). The other tradition has drawn upon ancient ideas of *eudaimonia* understanding well-being to consist of fulfilling or realising one's potential as a human being (Ryff 1995). What makes a person better-off according to the eudaimonia tradition need not include pleasure or the satisfaction of their desires. In what follows we will have something to say about the computational differences between momentary subjective happiness and overall well-being. However, for the most part we set aside the debate concerning the nature of psychological well-being. Instead, we take as our starting point the well-established conceptual connection between mental health and well-being.[3]

There is currently a growing literature within the field of computational psychiatry that applies the predictive processing framework (PPF) to model various psychopathologies including schizophrenia, depersonalisation, autism spectrum disorder, obsessive compulsive disorder, major depression, eating disorders, post-traumatic stress disorder, among others.[4] These psychopathologies are characterised by diverse behavioural, cognitive and emotional symptoms that manifest differently across individuals. Computational psychiatry provides a formal

[2] Interestingly, metastability has been shown to be related to well-being elsewhere in the literature on the neurobiology of eudaimonia (Kringelbach & Berridge 2017).
[3] This connection is enshrined for instance in the 1948 Constitution of the World Health Organisation health is defined as "a state of complete physical, mental and social well-being." We will assume the WHO definition of health is roughly along the right lines and health does indeed consist in a state of complete well-being. Our overall argumentative strategy will be to consider whether this account of mental health can also help us to understand the state of "complete well-being" in naturalistic terms as a biological condition of persons.
[4] For a non-exhaustive sampling of this literature see the references in footnote 1.

framework for relating symptom expression to neurocomputational mechanisms based upon a general theory of inference and control in biological systems (e.g., Montague et al. 2012). What seems to be common to psychopathologies are abnormal beliefs of various kinds and their behavioural consequences (Friston et al. 2014). Thus, it makes sense to seek an explanation of the expression of complex symptoms characteristic of a given psychopathology in terms of the inferential mechanisms that lead to the formation of these abnormal beliefs, and the control processes that result in pathological behaviours.

According to the PPF, agents infer beliefs and control their actions by maintaining and continuously updating a hierarchically structured generative model of their environment. The generative model is used to approximate Bayesian probabilistic inferences under conditions of uncertainty. When the agent gathers new sensory evidence, it must combine a likelihood function (a probabilistic mapping from hidden states of the world and their dynamics x to sensory inputs y) with its prior beliefs (a probability distribution that predicts possible states of the world over time x). These two probability distributions (the prior beliefs and the likelihood function) are referred to as the "generative model" (Ramstead et al. 2020). The likelihood and prior beliefs are described as a generative model because they can be interpreted as mapping how sensory inputs y are believed to be generated by states x of the environment. Given some sensory observations the generative model is used to compute the posterior probability of a possible state of the world that is the cause of those observations.

The generative model, as instantiated in humans, has a deep temporal structure that tracks the sensory consequences of actions over multiple timescales. Higher cortical layers of the model track regularities that unfold over longer time scales while lower layers of the generative model track faster changing events such as the sensory consequences of motor movements or the control of homeostatic setpoints (Kiebel et al. 2008). Minimising prediction errors in the long-run requires predicting the sensory outcomes of sequences of actions, sometimes reaching far into the future. Think for instance of organising your summer vacation. This calls for inferring plans that, when acted on in the future, bring about the preferred outcomes that are predicted - you're visiting a holiday destination in Greece. Prediction error here signals a mismatch between the predicted outcomes that one is aiming at in the future and the actual outcomes of one's actions. The generative model has temporal depth insofar as it aims to control the future outcomes of actions or expected prediction error.

The goodness of a model m can be measured by the model-evidence E, where E is the probability of sensory observations the agent samples when it acts, given the model m it uses to control its actions. Model-evidence is identical to negative surprise of a sensory observation y under a model m. A good model is one that minimises "surprise" understood in the technical information-theoretic sense of the negative log probability of a sensory observation given a model. Given some new sensory observations (prediction errors) that do not fit with the model's prior predictions, the agent should update its model so as to infer a new posterior prediction that better explains its sensory observations. A model that minimises surprise in the long-run (understood as long-term prediction error) qualifies as a good model. Our initial proposal is therefore to understand mental health very broadly, in terms of the "goodness" of a generative model the agent uses to form beliefs about the world and to control its actions.

According to the PPF, the abnormal beliefs that arise in various psychopathologies are hypothesised to be the consequence of an agent making use of a generative model whose prior predictions persistently fail to match with its sensory observations (Friston et al. 2014, for additional references see footnote 1). In order to avoid persistent surprise an agent will need some means of assessing the uncertainty of its prior predictions and of the likelihood function in a given context. These estimations of uncertainty can then be used to modulate the influence of new sensory evidence (prediction errors) on the model's subsequent predictions. The agent's uncertainty is referred to as *precision*, a measure of the reliability of information. The weighting that is given to the likelihood relative to past learning is referred to as the precision of the prediction error where precision refers to the inverse of the variance of a probability distribution. We can think of the precision of the prediction error as equivalent to the learning rate. Thus, precision of the prediction error is high when the likelihood is estimated to be precise, but decreases with precision of the prior predictions. The result of this kind of precision weighting is that inferential processes rely on past learning when new sensory information is weighed as imprecise and unreliable.[5]

Aberrant precision estimation is what leads to abnormal beliefs of the kind seen in psychopathology. The PPF claims that this failure to find the right balance between the precision of prior beliefs and current sensory evidence may be common to many different psychopathologies.[6] When too much, or too little, precision is given to prediction error signals, the agent will operate with a model whose predictions will come to diverge substantially from the sensory states it samples. Decreasing precision, for example, can lead to prior predictions dominating, as happens in schizophrenic delusion (Fletcher & Frith 2008; Corlett et al. 2010; Corlett & Fletcher 2014). By contrast, in Autism Spectrum Disorder too much precision is given to prediction errors relative to prior predictions (Pellicano & Burr 2012; Lawson et al. 2014; Palmer, Lawson & Hohwy 2017, Karvelis et al. 2018). People with autism are hypothesised to

[5] Three different sources of uncertainty can be distinguished (Parr & Friston 2017). The first source of uncertainty is the likelihood that may be excessively precise or imprecise. The result of estimating the precision of the likelihood function is a weighing of the reliability of a prediction error signal. A second source of uncertainty are the priors that map the dynamics of environmental causes. Volatility and noise may make for a high degree of unreliability in such mappings. A third source of uncertainty concerns the sensory states the agent has control over through their actions. The agent can be more or less confident that an action policy (a sequence of actions) will lead to the sensory states it predicts.

[6] See Hipolito et a. (forthcoming) for a computationally rich account of how "insulated" internal states, states that fail to be updated relative to incoming information, show psychotic (maladaptive) behaviours due to inevitable increases in deluded beliefs.

rely too much on current sensory information and only weakly on prior beliefs in making inferences about the state of the world over time.

Interestingly, the pathological behaviours that ensue are the result of processes that *approximate Bayesian inference* (Schwartenbeck et al. 2015). What makes the agent's behaviour pathological and sub-optimal is the generative model, and the prediction errors the agent repeatedly encounters, when they use the predictions of this model to control their actions. In the next section, we will show how health more generally can be tied to processes of allostatic control that ensure the body has the necessary metabolic resources available to meet the challenges of its environment. We will show how in the PPF, allostasis can be modelled as a process of prediction error minimisation.

2. Health and Allostatic Control

In this section, we expand on the claim made above, in part based on this literature in computational psychiatry, that the difference between mental health and psychopathology can be located in the goodness of the generative model as the regulator of the agent's behaviour. This proposal is related to what Conant and Ashby (1970) called the *good regulator theorem*, which states that an agent is only able to effectively regulate or control the states of its environment if it is a "good" model of its environment. We have seen above that the goodness of a generative model derives from its model-evidence. A good model is a model that maximises model-evidence or minimises surprise. Recall that surprise is to be understood as a mathematical measure of the unexpectedness of sampling a sensory state given a model. Maximising model evidence is identical to minimising surprise because the evidence the agent gathers for a model comes in the form of the sensory states the agent samples when they act to test the model's predictions. So long as the agent keeps the surprise of its sensory states to a minimum, they will succeed in maximising the evidence for the predictions of their model. In order for a generative model to minimise surprise, the sensory states the generative model predicts must be equal to the number of ways the environment can influence the agent, combined with the number of ways the agent can influence the environment through its actions. A generative model whose predictions systematically diverge from the states of the world, and their dynamics will fail to function as a good regulator of the agent's behaviour.

Of all the possible sensory states the organism can find itself in, a small subset will prove to be consistent with the organism remaining well-adapted to its changing environment. Sensory states belonging to this subset will include internal states of the body sensed through interoception and vital to the organism's continued existence (e.g., respiratory rate, blood acidity, glucose levels, bodily temperature, and plasma osmolality). These states of the body are maintained within a tight range of values compatible with the organism's viability through feedback control. Whenever the organism senses a deviation from these set-points, processes of physiological, hormonal and immunological regulation ensure that internal bodily states swiftly return to the set-points consistent with the organism's continued existence.

However, many of the brain and body's regulatory responses are not reactive but anticipatory. A predicted deviation from so-called "homeostatic setpoints" is avoided by taking preparatory action in advance of the deviation's occurrence. This process is referred to as "allostasis", meaning the stability of the internal conditions of the body through change (Sterling & Eyer 1988; McEwen & Stellar 1993; McEwen 2000).[7] Allostasis is a form of predictive regulation where anticipatory actions are selected that ensure that the organism's needs are prioritised and opportunities are weighed against dangers.[8] Examples of allostatic systems include hormonal, autonomic, and immune systems. The responsiveness of these systems is optimal when the brain is able to predict and accommodate the demands on the body before they arise. For instance, blood pressure varies continuously throughout the day. When the individual can let their guard down (e.g., during sleep) blood pressure drops sharply (Young et al. 2004; Lightman et al. 2020). When we wake in the morning, blood pressure ramps up in anticipation of stress and the need to remain vigilant. A comparable increase in blood pressure occurs during sexual intercourse. Fluctuations above and below an average state occur throughout the day depending on the need for the organism to maintain a state of vigilant arousal. Blood pressure is thus regulated to match the demands of a dynamically changing environment. The result of this matching of the body's resources to the predicted demands on its physiology and metabolism is the efficient regulation of the body's responsiveness to its environment.

If the body encounters constant high demand, this can result in the body adapting its predictions and remaining in a state of high arousal. Chronic stress arising from poverty, physical and emotional abuse, or loneliness leads the body to predict constant environmental challenges (McEwen 1998, 2000; Seeman & McEwen 1996; Adler & Ostrove 1999; Sterling 2011; Seeman et al. 2010; Cacciopo et al. 2015; Wilkinson & Pickett 2010). Just as muscles can learn to anticipate exercise, so also the body can learn to anticipate stress. This regulatory circuit can eventually enter into a pathological feedback loop. Arteries thicken and harden, consequently requiring higher pressure which further reinforces their stiffness (Sterling 2018: p.9). Chronically high blood pressure leads to inflammation of the nervous system and eventually to heart disease or stroke. So long as the body continues to predict the need for high blood pressure (an example of high "allostatic load"), the cycle will be very difficult to break.

[7] Cf. Sterling 2011; Power & Shulkin (2012). The latter defines allostasis as the means by which the body reestablishes homeostasis in the face of a challenge (p.25, cited by Corcoran & Hohwy (2018: p.4)). McEwen conceives of allostasis as anticipatory physiological responses aimed at restoring homeostatic variables to the range of values that allow for the maintaining of the organism's biological viability

[8] Sterling distinguishes allostasis from homeostasis on the grounds that the latter is reactive relying on negative feedback however the distinction cannot be drawn in this way if one thinks of homeostasis as working through active inference. Both processes are equally anticipatory, proactive and predictive. For a discussion of the relation between homeostasis and allostasis in the PPF see Stephan et al. 2017; Corcoran & Hohwy 2018.

We propose that health can be understood in terms of processes that forecast the likely demands on the organism's body by maintaining a generative model. This model is used to predict how signals arising internally and externally to the body are likely to evolve over time. Predictions track the likelihood that actions will maintain the body within the range of physiological, hormonal and immunological values consistent with its remaining well-adapted to the challenges of its environment. The proposal we explore in the remainder of this paper accounts for mental health and well-being in terms of a generative model that works in the service of allostatic control. In the next section we will take up the idea introduced above that maintaining a good model depends on the agent being able to estimate their own uncertainty in relation to who they are, what they are doing, and the world around them. Failure to accurately estimate uncertainty is thought to underlie various pathologies including addiction, a point we will return to in section four.

3. Reward, Error Dynamics and Momentary Happiness

We have seen in the previous section how mental health depends on estimations of uncertainty. Assigning too much precision, or too little precision, to prediction errors can result in abnormal beliefs and a generative model that fails to get a good grip on incoming sensory information. In this section we will show how precision predictions are maintained in part by tracking the rate of change in error reduction.

According to PPF, the outcomes of actions that are preferred and valued are highly expected, and the agent selects actions that fulfill those expectations (den Ouden et al., 2010, Friston et al., 2009, Clark, 2015, FitzGerald et al., 2014, Friston et al., 2012; Kiverstein, Miller & Rietveld 2017). Dopaminergic discharges (and other neuromodulatory chemicals such as serotonin, oxytocin, and norepinephrine) weigh the precision of a belief that an action policy will bring about expected outcomes (Friston et al., 2012; Schwartenbeck et al. 2014; Linson et al., 2018; Parr & Friston 2017). When we do worse than expected, the unexpected sensory and physiological states are punishing because they are states the outcomes which were not well predicted by the agent, perhaps because the agent does not have a good grip on the volatility of the environment or because they are acting on a high risk policy. Doing better than expected at reducing error indicates by contrast that there is less volatility or risk than one expected. One is therefore able to do better than expected at bringing about the valuable sensory states that one predicts to be the consequences of one's actions.

In recent work we have highlighted the role of doing better than expected at error reduction in contributing to precision estimation (Kiverstein, Miller & Rietveld 2017, 2020; see also Hesp et al. 2021). Unexpected increases or decreases in volatility are good information for the agent about how confident they can be that an action policy will lead to expected outcomes. Unexpected decreases in the rate of error reduction informs the organism that a belief in an action policy should be assigned lower confidence. An unexpected increase in rate of error reduction informs the organism that things are going better than expected. Precision, then, is adjusted on action policies not only based on the amount of error or error reduction occurring in the system, but also the rate at which error is managed over time (Kiverstein, Miller & Rietveld 2017, 2020; Hesp et al. 2021)

Error dynamics - the rate of change in error reduction - are registered by the organism as embodied affective states (Kiverstein, Miller & Rietveld 2017, 2020; Joffily & Coricelli 2013; Van de Cruys 2017; Hesp et al. 2021; Haar et al., 2020). We can think of an agent's performance in reducing error in terms of a slope that plots the various speeds that prediction errors are being accommodated relative to their expectations. Positively and negatively valenced affective states are a reflection of better than or worse than expected error reduction, respectively. Valence refers to the organism's evaluation of how it is faring in its engagement with the environment (i.e., how well or badly things are going for the organism). Think, for example, of the frustration and agitation that commuters feel when their train is late, and they have an urgent meeting to attend. These negative feelings are, in part, the body informing the system that some relevant source of error was expected to have been reduced by now but is not. The unexpected rise in error at the train's tardiness is felt in the body as an unpleasant tension. That tension may provoke the agent to check the transit authority for delays or find an alternative (more reliable) means of transport such as a taxi in order to reduce the felt tension - to catch back up to their previous slope of error reduction. We will henceforth describe optimally functioning agent's as being motivated to seek out good slopes of error reduction.

From this perspective, momentary subjective happiness is the result of unexpectedly reducing prediction error. This feels good because we have done better than expected at improving our predictive grip on the environment, something our very health depends upon (Sterling 2018, 2020). There are already a number of well-established approaches to understanding the neurobiology of momentary happiness that point in a similar direction. For example, Ruttledge and colleagues have, over a number of brain imaging experiments, demonstrated a strong relationship between subjective feelings of happiness and better than expected performance (2014, 2015).[9] Positively charged affect plays an important role in the predictive system. It ups the learning rates for situations in which there is a prime opportunity to learn how to adapt to the demands of the environment more efficiently, which is the modus operandi of the predictive system. We will see in the next section, however, that while this is no doubt an important part of

[9] Rutledge and colleagues had subjects engage in a probabilistic reward task, where they selected between various risky monetary options. Participants were asked between trials: "how happy are you right now?". Rutledge and colleagues showed that the feeling of happiness comes not when participants received a monetary payoff but when they did better than expected relative to their previous performance. This tracking of better-than-expected gains shows up in the brain in reward-related midbrain dopaminergic activity (Rutledge et al. 2014, p.12255). Instead of taking dopamine to track reward prediction errors, we have suggested that dopamine may track the rate of change in reduction of prediction error.

what it is to be well as a human being, it is not the whole story PP has to offer. Addicts can maximise their momentary subjective happiness but still find themselves in sub-optimal modes of engaging with their environment.

4. Bad Bootstraps and Sub-Optimal Grip

There are various dangers and difficulties that can arise in the optimisation of a generative model. The central role that prediction plays in generating perception and action means that hidden biases have tremendous power to direct behaviours in ways that tend to produce the outcomes that confirm just those biases. Relative to a predictive model, the agent can find themselves acting in ways that confirm their predictions, thus allowing them to minimise prediction error. Thus, having a generative model that succeeds in minimising prediction error is thus no guarantee of optimal psychological functioning.

Take as an instructive example long term substance addiction. Substances of addiction impact on the midbrain dopaminergic systems in the same way as unexpected rewards.[10] This has the effect of training expectations about the rate of error reduction both in the present moment, and over the longer term (Miller et al. 2020). The drug user comes to expect a tremendous reduction in error each time they use a substance. The continued release of dopamine that accompanies the use of the substance makes it seem as if the addictive substance is always and endlessly rewarding. The agent learns that nothing else in their life can reduce error in such a dramatic fashion. As a consequence, the agent neglects other policies that could serve the agent's goals. They get caught in a vicious cycle in which they act to fulfill the prediction that the drug seeking and drug using action policies are the best opportunity for realising their preferred and valued outcomes. So strong is the pull of the policy to use the addictive substance that the person pays no attention to other action policies that may also be of relevance to them. As they lose touch with their other cares and concerns error inevitably begins to build (e.g., health begins to degrade, relationships fall apart, jobs are lost), which in turn motivates the drug seeking and taking behaviours as a means of regulating the increasingly unmanageable levels of error.

A recent agent-based model showed that in order to optimise a model of the environment an agent must strike the right balance between epistemic actions that explore the environment for new policies, and pragmatic actions that exploit existing policies (Tschantz et al. 2020). A model that generates only pragmatic actions, like we see in the addition example above, will lead an agent to an overly rigid, sub-optimal course of behaviour we will henceforth refer to as a "bad bootstrap" (following Tschantz and colleagues). A model that generates only epistemic actions will be accurate and comprehensive, but it will fail to guide behaviour towards relevant

[10] Psychostimulants (e.g., cocaine, and amphetamines) act directly on this system producing a burst of dopamine as if the organism was encountering something which is needed. Opiates (e.g., heroin and morphine) inhibit GABAergic neurons leading to the disinhibition of dopamine neurons (Khoshbouei et al 2003).

possibilities for action in a dynamically changing environment. Agents learn an optimal model through strategies for balancing exploratory epistemic actions with exploiting what is already known for the purpose of pragmatic action. One way that organisms strike this optimal balance is by setting precision over action policies using their sensitivity to error dynamics. We will suggest *it is negotiating this explore-exploit trade-off by means of sensitivity to error dynamics that is key to well-being*. First, we use substance addiction to provide an illustration of how the prediction-minimising agent can get trapped in bad bootstraps.

Substance addiction is an example of a bad bootstrap because precision estimation over action policies is context-insensitive. Addicts choose the familiar option of seeking and using the drug, and continue to do so even when the outcomes are negative. In order to learn an optimal generative model an agent must flexibly update the estimation of precision on action policies with changes in context. PP theorists see addiction as a problem that arises when the higher-levels of the hierarchy (which is where the person's longer-term goals are encoded) are no longer assigned precision) (Pezzulo, Rigoli & Friston 2015; Clark 2019). Addiction, then, can be thought of as the result of a loss of contextualization between higher (cortical) and lower (subcortical) neural behavioural controllers. Goal-directed control at higher cortical levels provides the context for simpler habit-based and sensorimotor forms of control at lower-levels of cortical hierarchy. As drug-related habits become increasingly powerful, all the other goals that matter to the agent such as going to the gym or pursuing a promotion at work come to be neglected. Pathological forms of addiction arise when goal-directed and habit-based control come into conflict. The result of this conflict is a buildup of error in the person's life. Predictions related to goal-directed control at higher layers in the cortical hierarchy are trumped by highly precise prediction errors associated with drug-seeking and using behaviours. Instead of habit-based forms of control working in the service of fulfilling predictions arising from longer-term goals and concerns, habit-based control comes to drive action in isolation from goal-based predictions.

The key question the brain must settle is whether the agent is in a context in which habits can be relied upon to bring about valuable outcomes. Should the agent instead invest effort to explore for more valuable outcomes that do a better job of fulfilling long-term goals? To settle this question, however, requires the context-sensitive updating of precision estimation, which is exactly what fails to happen in pathological cases of addiction. People struggling with addiction tend not to gather more evidence that might lead them to change their behaviour. At least, they fail to do so until they are able to see through the illusion of error reduction induced by the effects of substances of addiction on the systems that estimate the precision of action policies.

The failure of this context-sensitive adjustment of precision leads the global dynamics of the brain to get trapped in fixed-point attractors that lead to a single attractive outcome. Fixed point attractors are contrasted with itinerant policies that allow for epistemic actions, and the exploration of sets of attractive states (Friston 2012; Zarghami & Friston 2020). Any given neural region can perform multiple functions over time depending on the patterns of effective connectivity it forms with other neural regions.[11] This multifunctional profile allows for task-specific coalitions to be configured on the fly as and when they are needed in a context-dependent manner (Anderson 2014; Clark 2016, ch.5). Recall that it is by means of the constant adjustment of precision estimations that patterns of effective connectivity in the brain emerge and change from moment to moment (Zarghami & Friston 2020). We've suggested above that neurotransmitters track the rate of change in error reduction (amongst other things). Positive and negative changes in the rate of error reduction are sensed in the body as positive and negatively charged affective states. We suggest these affective states (when all is going well) serve as an endogenous source of instability ensuring that neural coalitions form, dissolve, and reform in the brain in a context and task-dependent manner. In bad bootstraps rigid affect can have the opposite effect, trapping the global dynamics of the brain in sub-optimal patterns of engagement.

Bad bootstraps can be conceived of in dynamical systems terms as the loss of metastable dynamics.[12] Metastability is the consequence of two competing tendencies of the parts of a system to separate and express their intrinsic dynamics and to integrate and coordinate to create new dynamics (Kelso 1995; 2012). In a metastable system, there is "attractiveness but, strictly speaking, no attractor" (Kelso & Engström, 2006, p. 172; cf. Araújo et al. 2014). Attractor states describe the states in a system's phase space that the system tends to converge on when contextually perturbed. Metastable systems transit between regions of their state space spontaneously without the need for external perturbation. The organisation of a metastable system is therefore transient. For short periods, coordination among the parts emerges reflecting the tendency of the parts of the system to integrate. However, due to the tendency of the same parts to segregate, a recurring destruction of this coordination can also be observed as the behaviour of the component parts escapes from each other's orbit of influence. In the brain we see this creation and destruction of coordination in large-scale global patterns of synchronous and desynchronised activation across neuronal ensembles (Friston 1997, 2000; Varela 1999; Varela et al. 2001; Deco & Kringelbach 2016; Zarghami & Friston 2020). The brain as a metastable system is typically poised between stability (coordination of parts) and instability (segregation of parts) remaining close to a critical state from which the system can spontaneously shift from a coordinated to a disordered state and back again. We will close our paper by explaining why this poise between stability and instability might be necessary for well-being.

5. Metastable Attunement and Wellbeing

[11] "Effective connectivity" refers to the short-term moment to moment patterns of causal influence between neurons modelling by Dynamic Causal Modelling (Kiebel et al. 2009).

[12] Friston (2012) proposes that "metastability" is jeopardized in addiction by precision weighing being set too high on a certain set of sensory errors. This in turn specifically impedes itinerant wandering policies characteristic of metastable dynamics - the visiting of a succession of unstable fixed points in a phase space (Rabinovich et al. 2008). For more on this point, see section 5 below.

Agents like us that live in complex dynamic environments will benefit from remaining at the edge of criticality between order and disorder, between what is well known (and reliable) and the unknown (and potentially more optimal). Frequenting this edge of criticality requires that predictive organisms are prepared to disrupt their own fixed-point attractors (habitual policies and homeostatic setpoints) in order to explore just-uncertain-enough environments that are ripe for learning about their engagements. When things are going well, and they are on good slopes of error reduction, they should continue on the same path. When, however, a niche is so well predicted that there ceases to be good slopes of error reduction available, agents should begin to explore for opportunities to do better. Rate of error reduction is continuously changing. We will argue that if an agent uses error dynamics to set precision on action policies this will have the consequence that they avoid getting stuck in any attractor state. We will refer to this dynamical state of remaining metastably poised as a state of "metastable attunement". By tracking the changing rate of error reduction, such an agent will be attuned to opportunities to continually improve in error reduction.

Metastable attunement moves the agent in such a way that they find the balance between exploiting existing action policies and performing information-seeking epistemic actions that aim at reducing uncertainty. We have seen above how slower dynamics at higher layers of the hierarchical generative model provide the context that constrains the faster changing dynamics at lower layers of the generative model (Friston et al. 2020). The patterns of effective connectivity that form between higher and lower layers of the model are transient, changing each moment on the basis of precision assigned to policies. These patterns form, we have suggested, because of the role of valence in sculpting patterns of effective connectivity. Given the connection between valence and error dynamics, large-scale neural coalitions change from moment to moment in ways that reflect changes in the rate of error reduction. When a particular niche ceases to yield productive error slopes negative valence signals to the agent that they ought to destroy their own fixed-point attractors in favor of more itinerant wandering policies of exploration. Patterns of effective connectivity emerge and dissolve due to both environmental conditions and changes in our own internal states and behaviours. However, we also have a tendency to actively destroy these attractor states, thereby inducing instabilities and creating peripatetic or itinerant (wandering) dynamics (Friston, Breakspear, and Deco 2012). Alternatively, when errors accumulate, due to our frequenting spaces where there is an unmanageable complexity or volatility, the negative valence then tunes the agent to fall back on opportunities for action that are already well known and highly reliable. Notice, when all goes well such slope-chasing agents will be constantly moved by their valenced affective states (via changes in error dynamics) towards this edge of criticality, where error is neither to complex nor too easily predicted that the agent no longer has anything to learn (Kiverstein, Miller & Rietveld 2017; Anderson et al. 2020).[13]

Being attuned in this way to the edge of criticality makes for a resilient agent, one that can readily adapt to environmental challenges in a way that we have seen is necessary for allostasis. Systems that frequent this edge of criticality have fitness advantages over other more strictly ordered or chaotic systems because they strike an optimal balance between efficiency and degeneracy (Sajid et al. 2020). Such systems are able to respond efficiently to particular contexts of activity *while also* remaining open to exploring a wide variety of other possible contexts to bring about their goals (degeneracy) (Roli et al. 2018). This is precisely what people suffering from long term addiction tend to fail at - highly precise drug seeking and taking behaviours overwhelm the system leading it to inflexibly select those drug related policies even when other more beneficial policies may be available. Bad bootstraps like addiction create fragility in a dynamical system due to their making the system rigid and so less adaptable to a changing environment.

We have seen that metastable attunement allows the agent to remain poised over a multiplicity of possible actions. To put this in a different vocabulary from ecological dynamics: agents that are metastably attuned are able to maintain grip on a field of affordances as a whole (Bruineberg & Rietveld 2014; Rietveld, Denys & van Westen, 2018). This is because an agent that is able to remain at the edge of order and disorder will combine flexibility with robustness. Think of the boxer finding an optimal distance from the boxing bag where she is ready for all the relevant affordances the bag offers (Chow et al. 2011; Hristovski et al. 2009). She is ready to make jabs, uppercuts and hooks based on her distance from the bag. Given this bodily readiness, a random fluctuation of the bag then contributes to the selection of which action unfolds and which affordance she engages first. Systems that maintain metastable attunement are poised in a way that allows them to make the most of the affordances relevant to them, and to learn the most about the environments they frequent (see for example, Shew & Plenz 2013; Shew et al. 2011; Gautam et al 2015).

We suggest a distinction is therefore needed between local error dynamics that allow for the tuning of precision in relation to a particular action policy, and global error dynamics that track how well the agent is doing overall given the many affordances that are relevant to them.[14] Local

[13] Prediction errors that are neither too complex for a model to resolve nor too simple for the model to learn anything from we have called "consumable errors" (Miller et al. forthcoming). As slope-chasers we are motivated to seek out just the right quantities of manageable error that allow for the improvement of a model's predictions (Oudeyer & Smith 2016; Oudeyer, Kaplan & Hafner 2007; Kidd et al 2012; Andersen & Roepstorff, *under review*; cf. Berlyne 1970). Too many error signals that an environment is unmanageably volatile, while too little error means the environment is too well known for the predictive mind to learn.

[14] See Sandved-Smith et al 2020 for discussions about the structure of higher-level policies governing the allocation of precisions over lower-level tasks. In this 'deep parametric' generative modelling framework it becomes possible

success in error reduction is not sufficient for overall well-being. To see why not consider how a teenager might achieve this kind of improvement in their skills by spending their days playing computer games.[15] The computer game could provide them with just enough of a challenge to ensure that they are continually making progress in reducing prediction errors. We can suppose that the computer game would be designed to provide the player with just the right amount of prediction error - neither too much so that they find themselves frustrated, nor too little so that they quickly master the game and become bored of playing it. We can imagine that the game would create just enough novelty to keep the player engaged. But as with the example of substance addiction, this continued engagement would come at the expense of everything else in their lives. They may begin to neglect their friendships, schoolwork, and overall fitness in order to spend more time playing the game. Such an individual could not reasonably be said to be flourishing even though they may experience positive affect so long as they are playing the game.

Given that the agent has many cares and concerns, there will, on any given occasion, be multiple affordances of relevance to them. An important part of the optimisation of the generative model are apt predictions about how best to deploy precision in relation to any relevant affordances of concern to them. Changes in how well these predictions about precision fare can be used in much the same way as local error dynamics, helping to tune the agent in ways that keep them in touch with the best slopes of error reduction. However, instead of the slopes of prediction error management having to do with improvements in a specific domain, the high levels of the generative model that track global error dynamics pertain to the system's overall ability to manage volatility across multiple domains. The time scale of global error dynamics is longer than local error dynamics pertaining to how the general trend of error reduction is going into the future. For this reason, we suggest that the levels of the hierarchical generative model that control the deployment of precision are likely to be higher levels that deal with processes that unfold over long intervals of time.

Global error dynamics are important for psychological well-being because they allow an agent to maintain metastable poise over the field of relevant affordances as a whole. So long as the agent uses global error dynamics to adjust precision estimations, they will tend to act in ways that reflect their multiple cares and concerns. When an activity does not go as anticipated (say you are learning a musical instrument and struggling to play a piece of music) you can fall back on other projects or concerns that you also care about (such as your family relationships). You can switch from one activity to doing something else that is also expected to lead to valued outcomes. The result is that an agent can be failing to predict well in some local activity, but succeeding at predicting how to get into valued sensory states elsewhere, thus resulting in overall

---

to appreciate the precisions over these higher-level policies themselves, creating a nested hierarchy of error dynamics corresponding to local vs global considerations. We thank Sandved-Smith for discussions on this point.

[15] Our thanks to Andy Clark for pressing us on this point. For discussion of related examples see Clark (2018).

predictive success. Such an agent will continually make progress in learning, growing and broadening their field of relevant affordances, which will, in turn, increase their confidence in managing unexpected volatility as it arises over the whole of their lives. Since agents that make use of global error dynamics will do best at reducing error in the long run, they will tend also to occupy positively valenced affective states. (this follows from the explanation we have given of positive valence in terms of error dynamics.) This is to say they will tend to experience a positive hedonic sense of wellbeing over the course of their lives. They will experience a background mood of positive well-being - feedback that they are succeeding at deploying precision in an optimal way.

A key component of psychological well-being is therefore continual progress in learning that metastable attunement makes possible (cf. Kaplan & Oudeyer 2007; Oudeyer & Smith 2016; Kidd et al. 2012; Clark 2018). Metastable attunement doesn't just underwrite resilience, it also allows for the additional possibility of growth or improvement. Finding the right balance between pragmatic and epistemic actions, which is made possible by metastable attunement, is key. Doing so means that the agent will be able to optimally reduce long-term uncertainty. The result is an agent that will sometimes actively induce temporary stress in the form of increased uncertainty so that they can grow and improve in their skills.

There are certain human activities that increase the likelihood of metastable attunement. Interestingly these are also arguably activities that contribute to eudaimonic well-being. There are well established correlations between increased well-being over a lifetime and a focus on non-zero-sum goals and activities such as altruism, the development of virtue, social activism, a commitment to family and friends (Headey 2008; Garland et al. 2010). In contrast, pursuit of zero-sum activities, such as purely financial gains, has been found to be detrimental to life-long well-being (Headey 2005, 2008). The development of skills and abilities for engaging in non-zero-sum activities seems to be especially important for creating and sustaining lifelong satisfaction - or what is traditionally referred to as eudaimonia.[16] Why is this the case? Consider someone who approaches life as a zero-sum game. They will tend to develop skills and abilities that are socially antagonistic (Różycka-Tran et al. 2019). One side effect of this approach to life is that it can lead to missed opportunities for collaboration and social complexifications that often support long term success or happiness. A zero-sum approach to life tends to reduce or restrict one of our richest sources for reducing meaningingful prediction-errors: other people. In contrast, non-zero-sum activities encourage cooperation and collaboration, and therefore conducive to metastable attunement. These sorts of activities support a continuous opening to new possibilities and affordances. While the goal of buying a car comes to an end upon purchasing that car, the goal of being more mindful or compassionate, of being a better partner,

[16] Garland and colleagues (2010, 2015) have developed an account of how eudaimonic activities support well-being by encouraging upward spirals of psychological resilience and flourishing through forwardly-progressing and self-reinforcing cycles of positive affect and cognition.

or serving one's community are all goals that are potentially never finished. These are activities that allow for the continuous broadening of the field of relevant affordances we described above. The more one engages with non-zero-sum activities the more opportunities for development emerge - new skills to hone, new qualities to develop, new people to engage and collaborate with.

Conclusion

For prediction error minimizing agents like ourselves, optimality refers to our development of a generative model capable of successfully managing the volatility of our environments over the long term. Part of that optimization relies on the continual development and refinement of our various niche-appropriate skills and abilities. As we've seen, agents that are behaviourally tuned by changes in how well or poorly they are doing at reducing prediction error will be attracted to that critical edge where the most error can be resolved. The most resolvable error tends to be encountered just above the level of our current skillfulness - not so complex that we cannot get a good predictive grip and not so well known that there are no productive errors left to resolve. Momentary subjective happiness signals that our generative model is improving in its predictions. A system that is tuned by momentary subjective happiness, as we are, naturally becomes a better predictor of its environment over time. However, while this continuous progression in prediction is necessary for optimal well-being, it is not sufficient. We only have to reflect on the various ways that our current designer culture has manufactured for generating local predictive successes while diminishing our longer-term optimizations. Addictive activities as a whole are examples of this.

Optimal psychological functioning requires that we are able to continually develop in our various local projects *and* balance our metabolic expenditures between those activities in ways that provide good predictive dividends. Computationally speaking, this balancing occurs when the predictive system is able to make good predictions about how precision is being allocated to beliefs about action policies. When those predictions are good the agent is able to optimize the balance between exploiting well-learned policies and exploring new policies (even when doing so temporarily leads to increases in error).

We have proposed that optimal psychological function should be thought of as emerging from maintaining a metastable poise. A system that is sensitive to how it deploys precision, and so is able to juggle multiple cares and concerns in an optimal way, will also be a system that is best able to meet and resolve unexpected uncertainty. It is this continual growth of skills and abilities and the optimal balancing of resources between those domains of learning that produces this optimal control. And it is this optimal control that is experienced by the agent as a background feeling of well-being - the felt experience that the system is set up to handle life's many challenges.

References:

Adler, N. E., & Ostrove, J. M. (1999). Socioeconomic status and health: what we know and what we don't. *Annals of the New York academy of Sciences*, *896*(1), 3-15

Alexandrova, A. (2017). *A philosophy for the science of well-being*. Oxford University Press.

Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain*. MIT Press

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

Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning* (Vol. 135). Cambridge: MIT press

Barrett, L. F., & Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature reviews neuroscience*, *16*(7), 419-429.

Boorse, C. (1997). A rebuttal on health. In *What is disease?* (pp. 1-134). Humana Press, Totowa, NJ

Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in human neuroscience*, *8*, 599;

Cacioppo, J. T., Cacioppo, S., Capitanio, J. P., & Cole, S. W. (2015). The neuroendocrinology of social isolation. *Annual review of psychology*, *66*, 733-767

Chow, J.Y., Davids, K., Hristovski, R., Araújo, D., Passos, P. (2011). Nonlinear pedagogy: Learning design for self-organising neurobiological systems. *New Ideas in Psychology,* 29: 189-200.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, *36*(3), 181-204

Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press

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

Conant, R. C., & Ross Ashby, W. (1970). Every good regulator of a system must be a model of that system. *International journal of systems science*, *1*(2), 89-97

Constitution of the World Health Organisation 1948
Conference, I. H. (2002). Constitution of the World Health Organization. 1946. *Bulletin of the World Health Organization*, *80*(12), 983.

Corcoran, A. W., & Hohwy, J. (2018). Allostasis, interoception, and the free energy principle: Feeling our way forward. In Tsakiris, M., & De Preester, H. (Eds.). (2018). *The interoceptive mind: from homeostasis to awareness*. Oxford University Press

Corlett, P. R., & Fletcher, P. C. (2014). Computational psychiatry: a Rosetta Stone linking the brain to mental illness. *The Lancet Psychiatry*, *1*(5), 399-402

Corlett, P. R., Taylor, J. R., Wang, X. J., Fletcher, P. C., & Krystal, J. H. (2010). Toward a neurobiology of delusions. *Progress in neurobiology*, *92*(3), 345-369

Deco, G., & Kringelbach, M. L. (2016). Metastability and coherence: extending the communication through coherence hypothesis using a whole-brain computational perspective. *Trends in neurosciences*, *39*(3), 125-135;

Edwards, M.J, Adams, R.A., Brown, H., Pareés, I. & Friston, K. (2014). A Bayesian account of 'hysteria'. *Brain* 135 (11): 3495-3512.

Fletcher, P. & Frith, C. (2009). Perceiving is believing: A Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews: Neuroscience,* 10: 48-58.

Friston, K. (2010). The free-energy principle: a unified brain theory?. *Nature reviews neuroscience*, *11*(2), 127-138

Friston, K., Breakspear, M., & Deco, G. (2012). Perception and self-organized instability. *Frontiers in computational neuroscience*, *6*, 44

Friston, K. J., Fagerholm, E. D., Zarghami, T. S., Parr, T., Hipólito, I., Magrou, L., & Razi, A. (2021). Parcels and particles: Markov blankets in the brain. *Network Neuroscience*, *5*(1), 211-251

Friston, K. J., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: the brain as a phantastic organ. *The Lancet Psychiatry*, *1*(2), 148-158.

Garland, E. L., Fredrickson, B., Kring, A. M., Johnson, D. P., Meyer, P. S., & Penn, D. L. (2010). Upward spirals of positive emotions counter downward spirals of negativity: Insights from the broaden-and-build theory and affective neuroscience on the treatment of emotion dysfunctions and deficits in psychopathology. *Clinical psychology review*, *30*(7), 849-864

Garland, E. L., Farb, N. A., R. Goldin, P., & Fredrickson, B. L. (2015). Mindfulness broadens awareness and builds eudaimonic meaning: A process model of mindful positive emotion regulation. *Psychological inquiry*, *26*(4), 293-314

Gautam, S. H., Hoang, T. T., McClanahan, K., Grady, S. K., & Shew, W. L. (2015). Maximizing Sensory Dynamic Range by Tuning the Cortical State to Criticality. PLOS Computational Biology, 11(12), e1004576. https://doi.org/10.1371/journal.pcbi.1004576

Haar, A.J.H., Jain, A., Schoeller, F., & Maes, P. (2020). Augmenting aesthetic chills using a wearable prosthesis improves their downstream effects on reward and social cognition. *Nature: Scientific Reports*, 10: 21603.

Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K. J., & Ramstead, M. J. (2021). Deeply felt affect: The emergence of valence in deep active inference. *Neural Computation*, *33*(1), 1-49

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

Hristovski, R., Davids, K., & Araujo, D. (2009). Information for regulating action in sport: metastability and emergence of tactical solutions under ecological constraints. *Perspectives on cognition and action in sport*, 43-57

Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLoS Comput Biol*, *9*(6), e1003094

Kahneman, D., Krueger, A. B., Schkade, D., Schwarz, N., & Stone, A. (2004). Toward national well-being accounts. *American Economic Review*, *94*(2), 429-434

Kelso, J. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. MIT press

Kelso, J.S. (2012) Multistability and metastability: understanding dynamic coordination in the brain. *Philosophical Transactions of the Royal Society* B: 367: 906-918.

Kelso, J. S., Engstrom, D. A., & Engstrom, D. (2006). *The complementary nature*. MIT press

Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS one*, *7*(5), e36399

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

Khoshbouei H, Wang H, Lechleiter JD, Javitch JA, Galli A (2003) Amphetamine-induced dopamine efflux. A voltage-sensitive and intracellular Na+-dependent mechanism. J Biol Chem 278(14):12070–12077

Kringelbach, M. L., & Berridge, K. C. (2017). The affective core of emotion: linking pleasure, subjective well-being, and optimal metastability in the brain. *Emotion Review*, *9*(3), 191-199

Lawson, R. P., Rees, G., & Friston, K. J. (2014). An aberrant precision account of autism. Frontiers in human neuroscience , 8 , 302. doi:10.3389/fnhum.2014.00302

Lewis, M. D., & Todd, R. M. (2007). The self-regulating brain: Cortical-subcortical feedback and the development of intelligent action. *Cognitive Development*, *22*(4), 406-430.

Lightman, S. L., Birnie, M. T., & Conway-Campbell, B. L. (2020). Dynamics of ACTH and cortisol secretion and implications for disease. *Endocrine reviews*, *41*(3), 470-490.

McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, *38*(2), 339-346

McEwen, B. S. (1998). Protective and damaging effects of stress mediators. *New England journal of medicine*, *338*(3), 171-179

McEwan, B.S. (2000). Allostasis and allostatic load: Implications for neuropsychopharmacology. *Neuropsychopharmacology* 22(2): 108-124.

Miller, M., & Clark, A. (2018). Happily entangled: prediction, emotion, and the embodied mind. *Synthese*, *195*(6), 2559-2575

Montague, P. R., Dolan, R. J., Friston, K. J., & Dayan, P. (2012). Computational psychiatry. *Trends in cognitive sciences*, *16*(1), 72-80

Oudeyer, P. Y., & Smith, L. B. (2016). How evolution may work through curiosity-driven developmental process. *Topics in Cognitive Science*, *8*(2), 492-502

Palmer, C. J., Lawson, R. P., & Hohwy, J. (2017). Bayesian approaches to autism: Towards volatility, action, and behavior. Psychological bulletin, 143(5), 521

Parr, T., & Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of The Royal Society Interface*, *14*(136), 20170376

Pellicano, E., & Burr, D. (2012). When the world becomes 'too real': a Bayesian explanation of autistic perception. *Trends in cognitive sciences*, *16*(10), 504-510

Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in neurobiology*, *134*, 17-35

Pezzulo, G., Rigoli, F., & Friston, K. J. (2018). Hierarchical active inference: a theory of motivated control. *Trends in cognitive sciences*, *22*(4), 294-306

Power, M. L., & Schulkin, J. (2012). Maternal obesity, metabolic disease, and allostatic load. *Physiology & Behavior*, 106 (1), 22–28.

Rabinovich, M., Huerta, R., & Laurent, G. (2008). Transient dynamics for neural processing. *Science*, 48-50

Rietveld, E., Denys, D., & Van Westen, M. (2018). Ecological-enactive cognition as engaging with a field of relevant affordances: The skilled intentionality framework (SIF). In A. Newen, L. De Bruin, & S. Gallagher (Eds.), *The Oxford handbook of 4E (embodied, embedded, extended, enactive) cognition*. Oxford: Oxford University Press, pp.41-70.

Roli, A., Villani, M., Filisetti, A., & Serra, R. (2018). Dynamical criticality: overview and open questions. *Journal of Systems Science and Complexity*, *31*(3), 647-663

Różycka-Tran, J., Piotrowski, J. P., Żemojtel-Piotrowska, M., Jurek, P., Osin, E. N., Adams, B. G., ... & Maltby, J. (2019). Belief in a zero-sum game and subjective well-being across 35 countries. *Current Psychology*, 1-10

Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2015). Dopaminergic modulation of decision making and subjective well-being. *Journal of Neuroscience*, *35*(27), 9811-9822

Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences*, *111*(33), 12252-12257

Ryan, R. M., & Deci, E. L. (2001). On happiness and human potentials: A review of research on hedonic and eudaimonic well-being. *Annual review of psychology*, *52*(1), 141-166

Ryff, C. D., & Keyes, C. L. M. (1995). The structure of psychological well-being revisited. *Journal of personality and social psychology*, *69*(4), 719

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593-1599

Schwartenbeck, P., FitzGerald, T. H., Mathys, C., Dolan, R., Wurst, F., Kronbichler, M., & Friston, K. (2015). Optimal inference with suboptimal models: addiction and active Bayesian inference. *Medical hypotheses*, *84*(2), 109-117

Seeman, T. E., & McEwen, B. S. (1996). Impact of social environment characteristics on neuroendocrine regulation. *Psychosomatic medicine*, *58*(5), 459-471

Seeman, T., Epel, E., Gruenewald, T., Karlamangla, A., & McEwen, B. S. (2010). Socio-economic differentials in peripheral biology: Cumulative allostatic load. *Annals of the New York Academy of Sciences*, *1186*(1), 223-239

Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *371*(1708), 20160007

Seth, A. K., Suzuki, K., & Critchley, H. D. (2012). An interoceptive predictive coding model of conscious presence. *Frontiers in psychology*, *2*, 395

Shew, W. L., & Plenz, D. (2013). The Functional Benefits of Criticality in the Cortex. The Neuroscientist, 19(1), 88–100. https://doi.org/10.1177/1073858412445487

Shew, W. L., Yang, H., Yu, S., Roy, R., & Plenz, D. (2011). Information Capacity and Transmission Are Maximized in Balanced Cortical Networks with Neuronal Avalanches. Journal of Neuroscience, 31(1), 55–63. https://doi.org/10.1523/JNEUROSCI.4637-10.2011

Smith, L. S., Hesp, C., Lutz, A., Mattout, J., Friston, K., & Ramstead, M. (2020). Towards a formal neurophenomenology of metacognition: modelling meta-awareness, mental action, and attentional control with deep active inference

Sterling, P. (2012). Allostasis: a model of predictive regulation. *Physiology & behavior*, *106*(1), 5-15

Sterling, P. (2018). Point of View: Predictive regulation and human design. *Elife*, *7*, e36133

Sterling, P. (2020). *What is Health?: Allostasis and the Evolution of Human Design*. MIT Press

Sterling, P., Eyer, J., Fisher, S., & Reason, J. (1988). Handbook of life stress, cognition and health. *Allostasis; A new paradigm to explain arousal pathology. New York: Wiley*, 629-649

Tschantz, A., Seth, A. K., & Buckley, C. L. (2020). Learning action-oriented models through active inference. *PLoS computational biology*, *16*(4), e1007805

Van de Cruys, S. (2017). *Affective value in the predictive mind* (pp. 0-0). MIND Group; Frankfurt am Main.

Varela, F.J., Lachaux, J-P., Rodriguez, E. & Martinerie, J. (2001) The brain web: Phase-synchronization and large brain integration. Nature Reviews Neuroscience 2, pp.229-239

Wilkinson, R. & Pickett, K. (2010) *The Spirit Level: Why Equality is Better for Everyone.* London: Penguin.

Young, E. A., Abelson, J., & Lightman, S. L. (2004). Cortisol pulsatility and its role in stress regulation and health. *Frontiers in neuroendocrinology*, *25*(2), 69-76.

Zarghami, T. S., & Friston, K. J. (2020). Dynamic effective connectivity. *Neuroimage*, *207*, 116453