An Embodied Predictive Processing Theory of Pain

Abstract

This paper aims to provide a theoretical framework for explaining the subjective character of pain experience in terms of what we will call 'embodied predictive processing'. The predictive processing (PP) theory is a family of views that take perception, action, emotion and cognition to all work together in the service of prediction error minimisation. In this paper we propose an embodied perspective on the PP theory we call the 'embodied predictive processing (EPP) theory. The EPP theory proposes to explain pain in terms of processes distributed across the whole body. The prediction error minimising system that generates pain experience comprises the immune system, the endocrine system, and the autonomic system in continuous causal interaction with pathways spread across the whole neural axis. We will argue that these systems function in a coordinated and coherent manner as a single complex adaptive system to maintain homeostasis. This system, which we refer to as the neural-endocrine-immune (NEI) system, maintains homeostasis through the process of prediction error minimisation. We go on to propose a view of the NEI ensemble as a multiscale nesting of Markov blankets that integrates the smallest scale of the cell to the largest scale of the embodied person in pain. We set out to show how the EPP theory can make sense of how pain experience could be neurobiologically constituted. We take it to be a constraint on the adequacy of a scientific explanation of subjectivity of pain experience that it makes it intelligible how pain can simultaneously be a local sensing of the body, and, at the same time, a more global, all-encompassing attitude towards the environment. Our aim in what follows is to show how the EPP theory can meet this constraint.

Introduction

Pain is essentially subjective: there is no pain unless there is someone, a subject of experience, who is experiencing the pain. What is present to the person in pain does not allow for an appearance-reality distinction. If it seems to someone that they are in pain, this is what it is for them to be in pain. Pain does not have an experience or subject-independent existence (Auvray, Myin & Spence 2010). Our paper proposes to use the predictive processing theory as an explanatory framework for making sense of the subjective nature of pain.

The predictive processing (PP) theory is a family of views that take perception, action, emotion and cognition to all work together in the service of prediction error minimisation (Friston 2010; Hohwy 2013; Clark 2013, 2015; Barrett 2017; Seth 2015, 2021). In this paper we propose an embodied perspective on the PP theory we call the 'embodied predictive processing (EPP) theory. According to the EPP theory, pain is the outcome of unconscious inferential processes that are distributed across the homeostatic processes that make up the person's body.[1] Homeostasis refers to processes that maintain the internal physiological conditions of the body within relatively stable bounds despite undergoing continuous change. These stable physiological states can be thought of as predictions of homeostatic systems.[2] Potential or actual breaches of homeostasis such as noxious stimulation that threaten the ongoing integrity of the body give rise to prediction errors that demand urgent action. The homeostatic processes that protect the body from internal and external threats are not confined to the brain but are distributed across the whole neural axis (the peripheral and central nervous systems) in continuous reciprocal interaction with the autonomic, neuroendocrine, and immune systems. We will henceforth refer to this system as the neural-endocrine-immune (NEI) system.

The embodied predictive processing (EPP) theory of pain claims that pain experiences are generated by the NEI ensemble through processes that maintain the functional integrity of the body. These processes work by predicting the states that must be maintained within a range of values if the integrity of the body is to be preserved, and correcting for prediction errors when they arise. Our aim in what follows is to show how the PP theory of pain can make

[1] The argument of our paper is in the spirit of Colombetti and Zavala's (2019) recent argument against "affective brainocentrism" - the privileging of the brain over other physiological processes in affective neuroscience. Colombetti and Zavala (2019) show how the stress response involves "complex reciprocal influences among brain and bodily systems - endocrine systems in particular but also immune systems, the enteric system, and even the gut microbia" (p. 44). Affective states are not created or produced in the brain *simpliciter*. The bodily changes that occur when a person is stressed are not outputs controlled by the brain. We will make an analogous argument for the multiple homeostatic systems that maintain the integrity of the person's body.

[2] It is important to note that these predictions are not fixed set-points but can be adjusted over time depending on context through processes referred to as allostatic control (Sterling & Eyer 1988). Allostasis refers to processes that anticipate physiological changes before they arise and adapt to meet these challenges in ways that help to restore homeostasis. This process of adaptation comes at a cost referred to as "allostatic load" (McEwen 2000). For instance, blood pressure rises and falls throughout the day as physical demands on the body change. Through allostatic control, blood pressure can be adjusted in advance of these challenges arising. When this adjustment fails to happen and blood pressure is kept high, this results in a harmful allostatic load, and wear and tear on the body.

sense of the essential subjectivity of pain. Our strategy will be to draw upon phenomenological descriptions of pain as a state of the body that situates the person, both spatially and temporally, in its environment. Pain is typically experienced as located in a part of one's body - one's head, for example, in the case of a headache. However, this truism only partly captures how pain is embodied. In addition to sometimes being localisable to a body part, pain can often globally transform how a person relates to their own body, and to the surrounding world. The phenomenologist Minkowski rightly described pain as "an attitude towards the environment" (1958: 134). Consider for instance a person in chronic pain. This person experiences persistent pain in the absence of any measurable damage to the body. Such an experience will typically disrupt the person's habitual, practical bodily immersion in the world, which may come to be replaced by an "all-enveloping" attitude of doubt and distrust towards one's body and the world (Kusch & Ratcliffe 2018; Svenaeus 2015).

It is not only in chronic pain that a person's attitude to the environment is affected. It is widely accepted that pain has different dimensions - sensory-discriminative, affective-motivational, and cognitive-evaluative none of which are sufficient for pain experience (Melzack & Katz 2013). We would further argue that these different dimensions overlap and reciprocally influence each other in ways that precludes treating them as separable components. Thus, the sensation of pain felt in a body part does not suffice to determine the phenomenology of a pain episode. This sensing of the body typically occurs in a wider affective, motivational, and social situation in the world. The subjective character of a pain experience is therefore best characterised as a complex temporally extended process that radically disrupts and ruptures the person's embodied interaction with the world.

We take it to be a constraint on the adequacy of a scientific explanation of subjectivity of pain experience that it makes it intelligible how pain can simultaneously be a local sensing of the body, and, at the same time, a more global, all-encompassing attitude towards the environment. Our aim in what follows is to show how the EPP theory of pain can meet this constraint. More specifically, the PP theory makes sense of how the bodily processes that are responsible for maintaining homeostasis could also constitute the subject's embodied point of view on the world when they are in pain.

1. The Subjective Nature of Pain Experience

We started our paper by noting that pain experiences can be said to be subjective in the sense that pain does not admit of an appearance-reality distinction. If a person experiences pain it does not make sense to tell them: "you are not really in pain, it only seems to you as if you are". The experience of pain, and the reality of being in pain cannot be distinguished. Indeed, in the case of pain, what is experienced arguably depends for its very existence on its being experienced by a subject (Aydede 2006; Auvray, Myin & Spence 2010). One might think that if pain is essentially subjective it must therefore necessarily fail to be explained objectively. Classical arguments for the hard problem of consciousness and the explanatory gap have been premised on such a line of thinking (Nagel 1974; Kripke 1980; Jackson 1982; Levine

1983; Chalmers 1996). It has standardly been supposed that if a property is essentially subjective then it cannot also be explained objectively. Any objective scientific explanation of pain must necessarily leave out the subjective character of pain.

Such debates have a long history in the philosophy of mind. Our aim in this paper is not to argue that the predictive processing (PP) theory can settle this long-standing controversy in the philosophy of mind. We will argue however that the PP theory allows for progress. More specifically the PP theory can make it intelligible how bodily processes, understood objectively and neurobiologically in the terms of the PP theory, could constitute key phenomenological features of the subjective experience of being in pain. Our claim that pain is essentially subjective should therefore not be taken to imply an ontological commitment to pain qualia of the sort that has driven many philosophers to reluctantly embrace dualism. By 'pain qualia' we mean properties that are intrinsic to pain sensations, that are known directly and immediately to their subjects, and that make it the case that pain experiences are unpleasant and hurt for the subject that undergoes them. We agree with Dennett (2015) however that qualia, understood in these terms, cannot exist, since there is no double transduction in the brain (or indeed anywhere else in the body). In the PP theory we outline in this paper, a distinction can be drawn between predictions and the processing of prediction errors. However, at no stage in the processing of these predictions and prediction errors is there a conversion of this electrochemical activity into pain qualia with all of the special properties we just defined.

While we agree with Dennett that there is no double transduction in the brain, we nevertheless maintain that pain experiences are essentially subjective. Thus, we would take distance from Dennett's heterophenomenology according to which the subjectivity of pain experience is fully exhausted by third-person practices of making sense of what subjects say and do (Dennett 1991; 2003). We will argue that the subjectivity of pain is to be understood in terms of its embodiment by a bodily self - the body that each person experiences as their own body (Merleau-Ponty 1945/2012; Legrand 2006; Gallagher 2000; Zahavi 2005; Ciaunica & Fotopoulou 2017; Tsakiris & Fotopoulou 2017). It is in and through my body that the sensing of my pain experiences takes place, and the same is true of you. My body is, more generally, the locus of my perceiving, acting and thinking. My body is my subjective point of view on the world (Merleau-Ponty 1945/2012). Normally, the body opens a person to the world but in pain the opposite can occur. Pain typically transforms the experience of the kinds of possibilities the world has to offer. The body in pain can consume the person's awareness with the consequence that there is a shrinking of the space of possibilities for the person in pain. When this happens the phenomenology of pain comes to resemble what Matthew Ratcliffe has called an "existential feeling" (Ratcliffe 2008) - a kind of bodily feeling that situates the person in the world, orienting them to a space of possibilities. Pain, in common with other existential feelings, structures how the person experiences their current situation (cf. Coninx & Stillwell 2021).[3]

[3] We do not claim that pain *always* shrinks the space of action possibilities for an agent but such an experience is common, particularly in many people living with chronic pain (Coninx & Stillwell 2021). Such experiences of

Pain thus has a dual structure. It is a mode of sensing the body - a person can, for example, feel pain in their shoulder but pain also situates the subject in the world, sometimes contracting the space of possibilities for the subject. Carrying the shopping may become much harder for the person with the painful shoulder. Engaging in social life may also become a significant challenge in a way it was not previous to the onset of pain. When one is without pain, one's future can open out onto a wide range of possibilities. By contrast, when one is in pain, one's future contracts, and one can feel trapped and confined to dealing with the current moment. As Leder (1990) notes, the presence of pain can "render unimportant projects that previously seemed crucial" (p.74). In sum, pain can constrict one's possibilities for living (Heidegger 2001: p.158); it can change the life and everyday experiences of a person (Svenaeus 2015). Pain and suffering typically go together then not only because of the qualia of sensations localisable to body parts but because of how pain situates the subject in the world.

Still one might worry even if the subjectivity of pain experience is not to be explained by positing qualia - intrinsic properties of pain experience - still there must be an unbridgeable gap between the subjectivity of pain experience, and pain experience as objectively described in the science of pain.[4] We take our paper to be a contribution to the burgeoning research programme of naturalised phenomenology that aims to respect this difference between these two ways of understanding pain (Varela et al. 1991; Gallagher 2005; Wheeler 2005; Thompson 2007; Gallagher & Zahavi 2008; Rowlands 2010; Colombetti 2014). On the one hand there is the lived experience of the person in pain. One the other hand, there are neurobiological processes the description of which make it scientifically intelligible what it is for the person to experience pain. The core idea behind the naturalised phenomenology research programme can be described in terms of 'mutual constraints' that hold between these two ways of making sense of pain (Varela 1996; Gallagher & Zahavi 2008; Wheeler 2013).

A naturalising approach to phenomenology entails that both types of understanding (scientific and phenomenological) are necessary to account for the subjectivity of pain experience, but that neither will prove sufficient on its own. Phenomenology tells us what it is for a person to be in pain. It therefore provides a constraint on explanation in the science of pain by providing an understanding of what it is that stands in need of scientific explanation. A scientific explanation of pain experience, if it is to prove adequate, must take the phenomenon of pain as it is articulated and described in phenomenology, and provide an explanation of how this phenomenon is neurobiologically constituted. The constraints also run in the other direction. The science of pain provides constraints on the phenomenological understanding of pain. This much follows from a minimal commitment to naturalism, which

chronic pain are sadly all too common, and they are also revealing. They highlight how pain is not only felt in the body but can also structure how the person relates to their surrounding world, and to other people (Von Mohr & Fotopoulou 2018; Fotopoulou, Von Mohr & Krahé 2021).

[4] We are grateful to an anonymous reviewer for pressing us to further address this point.

we take to require that a phenomenological understanding of pain should not come into conflict with accepted empirical findings in the science of pain.

In what follows we will show how the predictive processing theory can satisfy the mutual constraint requirement of the naturalised phenomenology research programme. From phenomenology we borrow a description of pain experience as the experience of a bodily self situated in a world of meaningful action possibilities. The PP theory describes the causal elements, the organisation of these elements, and the systematic causal interactions among those elements that make intelligible in scientific terms how pain experiences could be neurobiologically constituted.[5] It does so by appealing to one type of state - precision-modulated prediction, and one type of process - error based learning.

2. The Predictive Processing Theory of Nociception

We will build up the predictive processing (PP) theory of pain by considering first how this theory applies to nociception. One should take care however not to conflate pain with nociception.[6] Nociception has the function of registering actual or potential damage to the body. Pain doubly dissociates from nociception: pain can occur in the absence of nociception and nociception can occur in the absence of pain (Baliki & Apkarian 2015). Nociception has the function of protecting the body from potential or actual injury but people do not experience pain each time they encounter a potential threat (Apkarian 2017). Nociception is arguably occurring all the time unconsciously without the person experiencing pain. Pain can also occur in the absence of nociception (Melzack 1999). Think for instance of phantom pain – pain that is felt in a limb that has been amputated or that is congenitally absent. In phantom pain there is no peripheral or spinal nociceptive activity but the person nevertheless experiences normal pain.[7]

The contribution of nociception to pain is anticipatory, occurring in response to the possible threat of tissue damage (Melzack 1996; Wall 1999). This makes nociception ideally suited to being explained in the terms of the PP theory. Nociception does not tell us about the actual state of bodily tissues but about the possible future state of the body, motivating the organism to engage in appropriate avoidance behaviours. Baliki and Apkarian (2015) have suggested that pain "signals the *failure* to protect tissue from injuries or from potential injuries, and as such is coupled with negative affect" (Apkarian 2017: p.74, our emphasis). A person only

[5] See Wheeler 2013: p.143 for a characterisation of enabling explanation in these terms. We have slightly adjusted his phrasing to fit with the case of pain experience we focus on in this paper.

[6] Philosophers of mind sometimes come dangerously close to making such a conflation. Tye (1995) for instance identifies pain experiences with representations of damage in the body but it is the nociceptive system that detects damage. He has suggested that a pain in the leg "is a token sensory experience that represents that something in the leg is damaged, something moreover that is painful or hurts" (Tye 1995: p.228). We take this conflation of pain with nociception to be a legacy of an unfortunate history in which materialist philosophers of mind defending an identity theory identified pain with the firing of c-fibers.

[7] This remains a controversial issue with some authorities suggesting that there is peripheral and spinal activity that arises from activity within the deafferented dorsal root ganglion cells (see e.g. Vaso et al. 2014).

gets to experience pain when the nociceptive system has failed in its function, and the organism has not successfully taken anticipatory action to avoid injury to the body (i.e. when there is prediction error to be resolved). The function of pain experience, one might think following this logic, is to move the person to take urgent action to avoid further, possibly life-threatening, damage to the body (cf. Auvray, Spence & Myin 2010). We will show later how such a perspective on the function of pain experience fits well with the phenomenological descriptions of pain as shaping a person's perception of the world.

The predictive model of nociception can be brought into view by contrasting it with what we will call the 'transduction' model of nociception (see Figure 1). On the transduction model nociception is the bottom-up conversion of external stimuli into an electrochemical signal in reaction to thermal, chemical and mechanical stimulation. Electrical and chemical signals are transmitted bottom-up from the periphery to the spinal cord. The predictive model we are proposing characterises the action potential of a nociceptive cell as a prediction error signal that occurs in response to the perturbation of constant ongoing tonic activity of the cell. This ongoing tonic activity consists of bidirectional flows of electrical and chemical processes along the cell. This bidirectional activity is what we are calling 'prediction' where what is being predicted is the integrity of bodily tissues, or more specifically, the kinds of activities needed to ensure the sensory states associated with tissue integrity. The tonic activity of nociceptive cells can be modelled as a prior distribution over possible states of the body the organism will tend to occupy, irrespective of environmental fluctuations, as long as the organism succeeds in preserving the physical integrity of its bodily tissues. The nociceptive cell is busily engaged in predicting the likely future state of the local tissue milieu top-down, harnessing its history of activity to keep the organism out of harm's way.

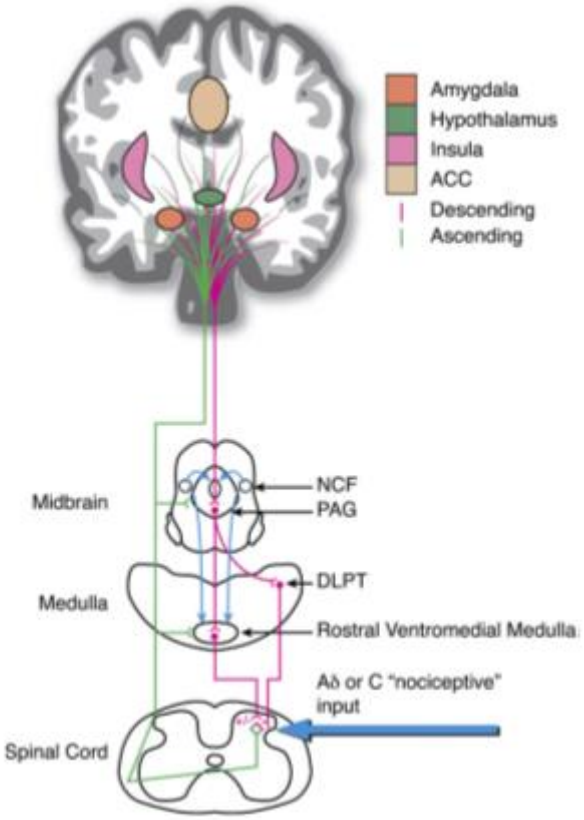

Figure 1. Typical Schematic of the Ascending and Descending Nociceptive Pathway (reproduced from Tracey and Mantyh 2007). The figure illustrates the bidirectional flow of information and the hierarchical organization within the system. In the predictive model of nociception what ascends is a prediction error signal which meets up with descending predictions. (NCF (Nucleaus Cuneiformis); PAG (periaqueductal grey matter); DLPT (dorsolateral pontine tegmentum); ACC (anterolateral cingulate cortex); N.B note original authors note +/- indicates both pro and anti-nociceptive influences respectively.

When the tonic activity of the cell is perturbed by an external stimulus this can be modelled as the process of combining prior predictions with new sensory information. If this new sensory information matches with the prior predictions nothing needs to happen. The prediction of the physical integrity of the body is confirmed. However, if thermal, chemical or mechanical stimulation occurs that threatens the integrity of bodily tissue, the result is prediction error. Prediction error can be modelled as the process of combining prior predictions with a likelihood function. The likelihood is the probability of the new sensory information (thermal, chemical or mechanical stimulation) given some prior beliefs, in this case, in the integrity of the body. Prediction error can be thought of as signalling that sensory information is highly unlikely given the prediction of the integrity of the body. Prediction errors carry the potentially important news of danger or deviation from the organism's ongoing bodily integrity.

We write "potentially important" because whether the prediction error is assigned importance will depend on the weighting that is given to the likelihood in relation to the prior predictions. We can think of the prior predictions in terms of the learning that has already occurred for the nociceptive system about possible threats to the body. The weighting that is given to the

likelihood relative to this 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 the imprecision 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. In the case of nociception, precision weighting has the consequence that only precise error signals that indicate a credible threat to the body get to have an influence on what happens next in the nervous system.

It is common in presentations of the PP theory to find a distinction made between two kinds of processing units in the brain: error units (superficial pyramidal cells) and prediction units (deep pyramidal cells). (See e.g. Friston 2010; Hohwy 2013, 2020; Clark 2013, 2015). Prediction units carry signals top-down conveying the brain's predictions of its sensory input. Error units calculate the difference between predictions and current incoming sensory input. Sensory input provides confirmatory or disconfirmatory feedback on the brain's predictions. Such a distinction makes some sense for canonical cortical circuitry. We suggest it makes more sense to view one and the same cell as doing both prediction and computing error when these predictions fail to match perturbing sensory input. It follows that what each level in the neural axis has access to is not a transduced sensory signal. Sensory neurons are first of all predictors of external stimuli, and what ascends the neural axis is prediction error.[8]

Prediction errors are feedback for the nociceptive system that its predictions of the ongoing integrity of the organism's bodily tissues do not match with the current sensory evidence. This prediction error signal can be used by the organism in two interrelated ways that correspond to perception and action.[9] The first way to resolve a prediction error is to update the predictions of the nociceptive system in such a way as to temporarily accommodate the prediction errors. One might think that this is impossible because the predictions of the systems responsible for nociception should consist of strict set points that do not change over

[8] Cao (2020) has argued that the role of the sensory signal in PP means that there is, what she describes as, an "informational equivalence" between predictive models and more traditional bottom-up models of perception. She writes "Just as predictive theories allow for – and indeed, require – bottom-up feedback from the outside world, traditional views also allow for top-down contributions to perception, whether from memory, context, or attention. Moreover, the idea of starting with perceptual priors and then updating them on the basis of incoming information is compatible with both predictive and traditional theories, as is a conception of vision as an essentially active process involving exploration" (Cao 2020: p.5). Cao concludes on this basis the evidence doesn't decide between predictive and more traditional models. Both are equally able to accommodate the available evidence. In our view Cao's argument is able to get off the ground because of the distinction PP theories typically make between prediction and error units. This leaves room for the sensory signal that error units receive as input to be conceived of along traditional transductive lines. We are proposing a different interpretation of the nervous system in which its default mode of processing is predictive (cf. Buzsáki 2019 on what he calls the "inside-out" view of the brain and its functions). Sensory input only gets to impact on this ongoing tonic activity when important errors are detected.

[9] It should be noted that the PP theory claims that perception and action are co-determining and are therefore not separate processes. This point is sometimes expressed using the control theory of perception as proposed by Powers (1973) according to which action is for the control of perception (Clark 2016; cf. Anderson 2017). In the context of nociception this control of perception can be thought of as the maintaining of the integrity of bodily tissues.

time. However, such an objection is misconceived. The sensitivity of nociceptors is plastic and can change over time with context and as a function of bodily trauma. Following injury for instance the firing thresholds of nociceptors are lowered so that nociceptors that would previously have been silent, respond to what would normally be counted as innocuous stimulation. There is some evidence that chronic pain is in part the result of the failure to readjust the sensitivity of these receptors (Chapman et al. 2009).

The updating of nociceptive predictions occurs as part of the process of controlling actions aimed at harvesting sensory input that, if all goes well, will bring its bodily states back into the range of values consistent with tissue integrity. A simple example is withdrawing your hand from a hot object swiftly terminating contact with a noxious stimulus. This process is called "active inference". These two techniques for resolving prediction errors are interrelated and interdependent in that neither will suffice on its own to minimise prediction error. Active inference is needed to make nociceptive predictions more reliable by better aligning the generative process (i.e. the conditional dependency between actions and outcomes) with prior beliefs. Perceptual inference is needed because successful regulation of action depends upon making good predictions about the outcome of actions (see Hohwy 2013, ch. 4; Hohwy 2020). Active inference however takes centre stage since all of the predictions are organised around controlling actions with the goal of maintaining homeostasis.

So far, our presentation of the PP theory of pain has focused on nociception. However, as we started this section by noting, pain cannot simply be identified with nociception. There can be nociception in the absence of pain, and pain in the absence of nociception. In the next section we show how a variety of non-nociceptive inputs also seem to play a necessary role in the organism's reaching the conclusion of a real and possible threat to the body. The non-nociceptive inputs to this inferential process come from different systems distributed across the body as a whole, in addition to the central and peripheral nervous systems. These systems include the neuroendocrine, neuroimmune systems, and the autonomic nervous system.

3. The Embodied Predictive Processing Theory

Injury to the body disturbs bodily tissues but it also "triggers inflammation, constricts blood vessels, promotes coagulation and stimulates immune response" (Chapman et al 2009: p. 2). The immune system generates a variety of cellular and molecular inflammatory responses that aim to protect the injured area from microbial invasion. The autonomic system is responsible for anticipating potential threats preparing the body for 'fight, flight or freezing'. The endocrine system orchestrates the body's stress response - an allostatic response of mobilising metabolic resources to meet internal and external challenges to the body. The response of the endocrine system begins with arousal in response to a stressor that in collaboration with the autonomic system prepares the body to take adaptive action. A second slower phase promotes recovery bringing the body back to normalcy. We will refer to these systems that respond in a coordinated and coherent manner to injury as forming a neural-endocrine-immune (NEI) ensemble (following Chapman et al. 2009).

The NEI ensemble is a self-organising, complex, adaptive system selecting responses to internal and external stressors with the aim of avoiding catastrophic phase transitions. The result of these self-regulatory actions is that the NEI ensemble ensures the physical integrity of the body is maintained over time. The integrity of the bodily tissues can be thought of as a non-equilibrium steady-state. Injury perturbs the body pushing it (if all goes well) temporarily out of this steady-state. As a complex adaptive system, the NEI ensemble can be modelled using the tools of the PP theory (Friston 2012). The self-organising dynamics of the NEI ensemble predicts the states of the body that need to be kept within a range of values if the organism is to maintain its own homeostasis. The bodily states these systems are predicting are the set of states the NEI ensemble tends to evolve towards from a wide variety of initial states as long as homeostasis is maintained. The NEI ensemble performs allostasis predicting possible challenges to bodily homeostasis before they arise and mobilising the body's resources to meet those challenges (Sterling 2012). When an injury to the body occurs this can be modelled as an increase in uncertainty. What is uncertain is how to deal with a challenge to the body the injury presents. Allostatic processes, mediated by the NEI ensemble, aim at resolving this uncertainty as fast as possible. This can be done in the two ways indicated above, either by changing the predictions of the NEI ensemble (perceptual inference), or by changing the sensory evidence through initiating actions (active inference). These actions can take the form of inflammatory responses of the immune system, and stress responses by the endocrine system and the autonomic nervous system.

We are claiming that pain experience is constituted by the activity of the whole NEI ensemble. Both brain processes and physiological bodily processes work together to constitute a pain experience. For this reason we call our perspective on pain experience an '*embodied* predictive processing' theory. The NEI ensemble exhibits ongoing, endogenous self-generated activity, which we take to be predictive of the states of the body that must remain within a narrow range of values if homeostasis is to be maintained. The perturbation of these systems sometimes takes the form of noxious sensory states that are outside the range of what is predicted. When this occurs the result is a prediction error, which the aforementioned systems will take measures to resolve.

4. Nested Markov Blankets

In this section we show how to think of the systems that make up the NEI ensemble in terms of a nesting of Markov blankets (Kirchhoff et al. 2018; Palacios et al. 2020; Hipólito et al. 2021). The terminology of Markov blankets is borrowed from the literature on causal Bayesian networks (Pearl 1988). This formalism is then applied to the causal dynamics of systems that minimise prediction error.[10] (See figure 2 for a simple depiction of a Markov

[10] Bruineberg et al. 2020 have criticised the use of Markov blankets in the PP literature for conflating a map (the use of the Markov blanket formalism in modelling a system's behaviour) for the territory (the boundary of the system of interest whose behaviour is being modelled) (cf. Andrews 2020). They have argued that the Markov blanket formalism is best viewed as applying to the causal dynamics of a system only under a number of

blanket with full conditionals (i.e., the conditional dependencies between elements that constitute the system and the conditional independencies between internal and external states of a system):

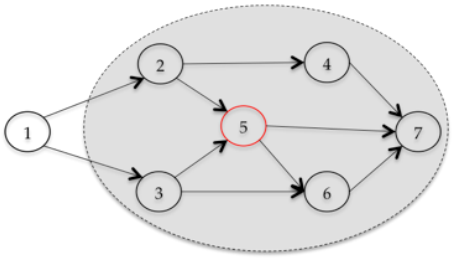

Figure 2. A schematic depiction of a Markov blanket with full conditionals (Kirchhoff et al. 2018). The Markov blanket is the smallest set of nodes {2,3,4,6,7} that renders a target node {5} conditionally independent of all other nodes in the model {1}. The central point to note here is that the behavior of {5} will be predictable by knowing the nodes making up its Markov blanket. This means that any node external to the system in question - in this case, node {1} - will be uninformative vis-a-vis predicting the behavior of {5}. This means that once all the neighbouring variables for {5} are known, knowing the state of {1} provides no additional information about the state of {5}. It is this kind of statistical neighbourhood for {5} that is called a Markov blanket (Pearl 1988). See Kirchhoff & Kiverstein (2019) for additional information.

The Markov blanket for a node in a Bayes network comprises the node's parents, children and parents of its children. The behaviour of the blanketed nodes can be predicted from the states of the blanket without knowing anything about the nodes external to the blanket that are the causes of changes internal to the network. Transposed to the PP theory, the nodes of the Bayes network can be mapped onto the internal states of the generative (the predictive) model. The children of these internal states are taken to be the active states by means of which the organism samples sensory states that over the long run tend to minimise prediction error. The parents of the internal states are the sensory states that are used to drive inference. Thus, we can think of the sensory and active states as forming a boundary for the organism that is produced and maintained through processes of prediction-error minimisation. We will henceforth refer to the boundary of a prediction-error minimising system, that demarcates the internal states of this system from the states that are external to the system, as a 'Markov blanket'.

We have argued that predictive processing takes place at multiple spatial and temporal scales in the NEI ensemble, all the way down to the scale of the individual receptor. The statistical

simplifying assumptions. Thus, the Markov blanket should not be taken to be a boundary for the brain but an explanatory construct that is more or less useful in causal modelling. To fully engage with their critique is beyond the scope of this article but in our view their carefully argued paper misses something important about how the Markov formalism has been applied in the PP literature. The formalism is applied to prediction-error minimising systems where this process of prediction-error minimisation works in the service of maintaining homeostasis. It is this point that justifies the inference from the description of the causal dynamics of the system using the Markov blanket formalism to the use of the terminology of Markov blankets to refer to the boundary of this system. Now one could ask if realism about the description of living systems as prediction-error minimising systems is justified. We are assuming in this paper that such a description is warranted by the scientific literature. See our earlier discussion of the nervous system as fundamentally predictive in its workings, as well as Kiverstein & Kirchhoff (submitted); Kirchhoff, Kiverstein & Robertson (in preparation).

form of the Markov blanket will be the same all the way down to the individual receptor, and all the way up to the scale of the whole organism (and perhaps also at the scale of groups of organisms). Recall we are applying the formalism to systems that are minimising prediction error in order to contribute to the maintenance of homeostasis. Thus, we are claiming that processes of prediction-error minimisation are taking place from the smallest to the largest scale of the nervous system as a whole. The same principles of organisation that apply to the cell - prediction-error minimisation that induces a boundary for the cell, separating what is inside of the cell from what is outside - operate across multiple scales, all the way up to the organisational scale of the whole individual.

For a single cell, the Markov blanket is a boundary that takes the form of alterations across the cell membrane that subsequently mediate the interactions within the cell and with other cells (Friston 2013; Kirchhoff et al. 2018; Palacios et al. 2020). The extracellular environment (i.e. transducible units such as heat, acid, mechanical deformations) are equivalent to what we called sensory states above that influence the interactions within the cell but are not themselves influenced by these interactions. The membrane potential maps onto what we have called the active states, and is influenced by, but does not influence, the states internal to the cell. Cells are also homeostatic processes, reflected by their resting cell membrane potential. They maintain the integrity of their internal organisation through processes of prediction-error minimisation. Thus, external stressors can constitute prediction-error for the cell that threatens its internal integrity (e.g. by alterations in the cell membrane potential). The membrane thus forms and is maintained through the process of prediction-error minimisation. If the membrane potential continues to alter, the eventual consequence is the death of the cell.

Individual cells are parts of larger self-organising processes that can also be described as having their own Markov blanket. These larger scale processes also have boundaries that are produced and maintained through processes of prediction-error minimisation. We talk of the ‘nesting’ of Markov blankets within each of the systems that make up the NEI ensemble because each component of a Markov blanketed system will have its own Markov blanket. The immune system, the neuroendocrine system, and the autonomic system are each composed of cells that also have their own Markov blankets. These systems can be described as networks of cells that maintain their integrity as a whole functional unit under changing conditions. As stable biophysical structures they owe their stability through change to prediction-error minimisation. The Markov blanket formalism can thus be applied to any prediction-error minimising system to describe how the system forms a boundary that distinguishes the states that are internal to the system from those that are external. These kinds of boundaries are not merely between the agent and its environment, but form a series of nested and multiscale boundaries constituted by a multiplicity of Markov blankets, as per figure 3:

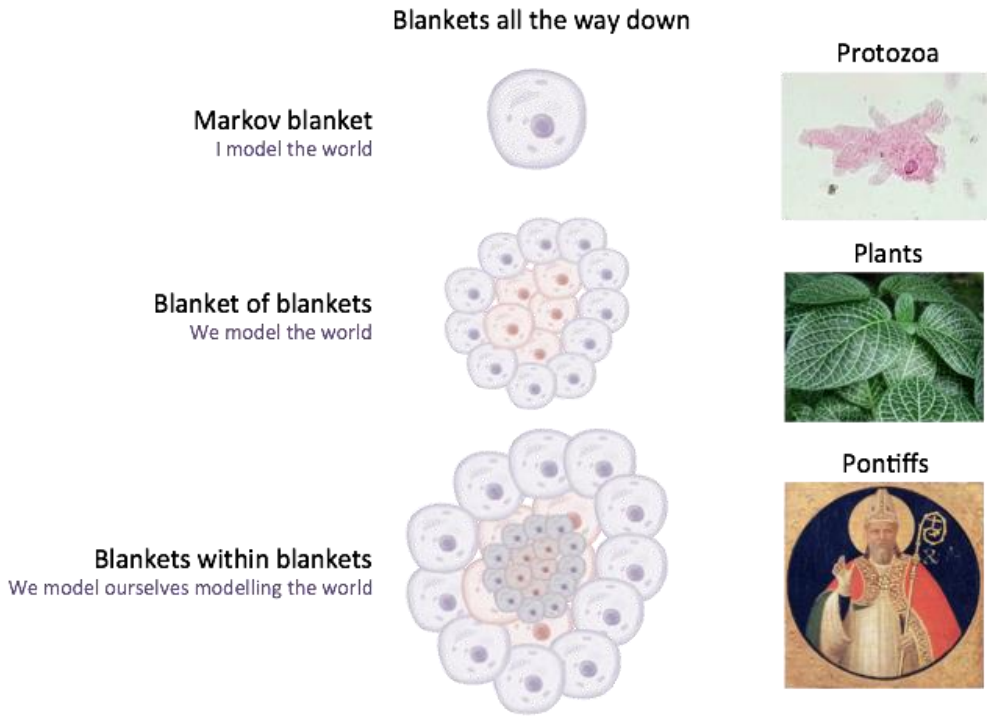

Figure 3: Schematic depiction of Markov blankets. The top figure depicts a single Markov blanket. The middle figure represents a multiscale and nested organisation of Markov blankets. The final figure suggests that cultural practices can envelope a multiplicity of individuals given its nested structure. Figure 2 represents the Markov blanket organisation all the way down to individual cells and all the way up to complex organisms like human beings (Kirchhoff et al. 2018).

Crucially, Markov blankets do not only segregate, but also integrate, the systems that make up the NEI ensemble. The immune, endocrine and autonomic nervous system continuously and reciprocally influence each other through the production of neurotransmitters, peptides, hormones, endocannabinoids and cytokines. We suggest the nesting of the Markov blankets within blankets over multiple spatial and temporal scales allows for these different subsystems to work together as an integrated whole (cf. Palacios et al. 2020).

As a toy example of how this nesting of Markov blankets applies to pain experience, consider Dewey's example of a child that touches a candle flame (Dewey 1896). The candle flame initially looks attractive to the child, which elicits the child's movement towards the flame. The contact with the flame leads to a perturbation of ongoing activity along the neural axis. This perturbation has cascading effects throughout the body such as changes in the electrochemical activity including nociceptors, mechanoreceptors sensitive to heat, and in addition the other systems listed above. Before contact with the flame occurs there are already anticipatory changes in the autonomic system such as change in heart rate, and blood flow. Active states here are not just tied to movement. Release of hormones into the bloodstream or release of catecholamines from the autonomic neuron termini would also count as active states on our account. Activity within these systems can be described in terms of a nesting of Markov blankets with each component of these systems contributing to

maintaining the integrity of the child's body by together orchestrating the swift withdrawal of the child's fingers.

5. Towards an integration of the neurobiological, psychological and social dimensions of pain

Now that we have some of the details of the EPP theory of pain in place, we will return to the challenge of rendering intelligible the complex phenomenology of pain experience in the terms of the science of pain. In the final parts of our paper we will use the EPP theory to provide a naturalised description of this phenomenology. Recall that the naturalistic theory we are aiming for takes the phenomenological understanding of pain as an articulation of the phenomenon that stands in need of explanation. We have described pain experiences as having a dual structure. On the one hand, pain is characterised in terms of a mode of sensing the body. Abdominal pain can, for instance, feel like we are being stabbed in the stomach. The person feels assaulted by their pain, the greater the intensity of the pain, the less able the person is to ignore this assault. On the other hand, it is through my body that I am situated as a bodily self in the world. My body is my manner of relating to the world. The unity the body forms with the world is however disrupted when the person is in pain. The body is thereby transformed from the transparent medium of action into an impediment to acting. The person is no longer at home in the world. If health is in Gadamer's words a "condition of being involved, of being-in-the-world, of being together with one's fellow human beings, of active and rewarding engagement in one's everyday tasks" (Gadamer 1996: p.113), pain can be thought of as a disruption of this being-in-the-world. The result of such a disruption is that the world appears both threatening and alien to the subject.

To make sense of this rich phenomenological profile of pain will call for a model of pain that integrates the biological, psychological and social dimensions to being in pain. Yet so far arguably no model of pain has succeeded in accounting for the dynamic integration of all these elements. In this section, we will show how the EPP theory promises to deliver such an integration. Consider first, the dynamic integration of the biological and psychological. The EPP theory avoids making any sharp separations between the sensory, affective and motivational dimensions of pain that seem to be directed at the biological body and its states, and the cognitive-evaluative dimensions that comprise your beliefs about the body. The sensory-discriminative, affective-motivational, and cognitive-evaluative dimensions of pain should not be thought of as distinct parts or components that make up a whole complex pain experience. Although they can be distinguished analytically, and can on rare occasions come apart in people that have undergone brain damage (Grahek 2007), these dimensions are typically unified in pain experience, and the EPP theory explains how such a unification could be achieved.

Take for instance the cognitive-evaluative dimensions of pain. It is a prediction of the EPP theory that there is no sharp line separating what one believes, how one is affected by pain, and what one is motivated to do. Pain experiences are sensitive to what we say, think, know, and expect (Noë 2016: p.72; cf. Clark 2019). In line with this prediction, it has been shown

that modulating the magnitude of expected pain influences how intense a painful stimulus is experienced to be (Wiech 2016). Medium-intensity stimuli can be experienced as more or less painful depending on what the person expects (Atlas et al. 2010, Leknes et al. 2013). PP theorists have also done work on placebo and nocebo effects (Buchel et al. 2014; Anchisi & Zanon 2015; Ongaro & Kaptchuk 2019). A placebo can for instance induce an expectation of safety resulting in an estimation of lower overall threat to the body and a consequent decreased pain sensation (hypoalgesia, see Moseley 2008; Buchel et al. 2014). In nocebo the opposite can happen – an expectation of harm when combined with a noxious stimulus can result in the estimation of a highly significant threat to the body and thus to a more intense pain experience (Anchisi & Zanon 2015, Benedetti et al. 2013, Buchel et al 2014).

The experience of pain seems to reflect the "overall estimate" of threat that is posed to the body in a particular environment, based upon the integration of relevant information from multisensory sources (Tabor et al. 2017: p.4; based on Moseley and Arntz 2007). Exteroceptive sensory cues for instance have been shown to have an influence on whether a nociceptive cue results in a pain experience (Tabor et al. 2017). When a noxious stimulus is paired with a red light (associated with heat and danger) and with a blue light (associated with safety), the noxious stimulus is perceived as more painful when paired with the red light as compared with the pairing of the same stimulus with the blue light.

Turning now to the social dimension of pain, factors such as social support and empathy have been shown to modulate the intensity of a pain experience (Krahé et al. 2013, 2015; Paloyelis et al. 2016). Von Mohr and Fotopoulou (2018) propose an explanation of this social modulation of pain experience in terms of precision estimation. The more intense the pain experience, the greater attention is captured by the experience. The organism's energetic resources are dedicated to terminating the pain experience. This is to say that the prediction error indicating uncertainty about the continued integrity of the body is assigned a high degree of precision. Von Mohr and Fotopoulou suggest that the support of others can help to reduce the precision estimation. This has the effect of down-weighing the prediction error indicating a credible threat to the body, in favour of the prediction that the body remains in its expected safe condition. In other words, social support provides evidence in favour of the hypothesis that the body is safe, and this leads to prediction errors indicating uncertainty about the continued integrity of the body being treated as less reliable or trustworthy.

Fotopoulou and Tsakiris (2017) argue that the bodily self is rooted in inferential processes that integrate and schematise exteroceptive and interoceptive signals to form a model of the body in its environment (also see Ciaunica & Fotopoulou 2017). They show how such processes of multisensory integration are crucially mediated by interaction with other people. This is to say that the sensory evidence used to form a model of the body is partly based on sensory inputs gathered in interaction with others. Hence "the very first-person experience of my body as mine, as well as the building block of the self-other distinction, are constituted upon the presence of other bodies in proximity and interaction" (Fotopoulou & Tsakiris 2017: p.13; Ciaunica et al. 2021a). In development for example, the contribution of interoception to this inferential process of multisensory integration is mediated by interaction with caregivers.

An infant is largely dependent on others for maintaining homeostasis. The updating of predictions about whether the infant's needs are met or not will often depend on the actions of their caregivers. Infants are in a vulnerable position in this respect: they mostly depend on their caregivers to restore homeostatic balance whenever it is lost. Feelings that signal how well or badly the infant is doing at meeting its needs, such as hunger and satiation, or pain and relief, can therefore be thought of as often originating in social interaction (*Ibid*, p.18).

Consider in this light the finding that the effects on pain of social support varies with attachment styles (Hurter et al. 2014; Krahé et al. 2015). Individuals with high attachment avoidance showed increased pain in the presence of their romantic partner (Krahé et al. 2015). This is in contrast with individuals with a secure attachment style that show diminished pain in the presence of their partners. Moreover, early attachment experiences have been shown to optimise the use of oxytocin for dealing with stressors in adulthood. Oxytocin plays the role of down-regulating HPA (Hypothalamic Pituitary Adrenal) axis reactivity, which we have argued above forms a part of the NEI ensemble that constitutes pain experience (Quirin et al. 2011). If Fotopoulou and Tsakiris are right, this effect of attachment on pain experience is a consequence of the social mediation of the neurobiological processes that anchor the bodily self. The reason that interaction with others can impact on pain experience is because social interaction is a crucial source of sensory input for the inferential processes that form the basis for pain experience. It might be thought that the skin forms a Markov blanket around each individual that separates individual organisms from each other. However, crucially the skin also provides a point of affective contact between organisms. Markov blankets thus do not only segregate individuals, they also connect us to others.[11]

Finally, we note that the EPP theory is also able to integrate the sociocultural dimension of pain.[12] It can for example account for the role of social adversity in sustaining chronic pain. Research suggests that between 10 to 14% of adults in the UK live with moderate to severe chronic pain (Fayaz et al. 2016). Population studies show that:

> "... the prevalence of chronic pain is inversely related to socio-economic factors (Janevic et al. 2017; Blyth 2008; Poleshuck & Green 2008; Maly & Vallerand 2018). Those who are socio-economically depri§ved are not only more likely to experience chronic pain than people from affluent areas, but they are also more likely to

[11] This is a point that has been emphasised in recent work on touch and co-embodiment by Anna Ciaunica and her collaborators (Ciaunica et al. 2021a, b). Ciaunica et al. (2021a) introduces the concepts of co-embodiment, and co-homeostasis to describe the relationship between the mother and foetus in utero.

[12] An additional source of evidence comes from a study of healing prayer on patient's beliefs that they have been cured of a sickness (Paldam & Schjoedt 2016). The study showed that although subjects would occasionally report being cured of visual, hearing and walking impairments, the main effect of healing prayer was to provide patients with pain relief. The authors conclude that healing practices appear to be "specifically associated with pain relief in the musculoskeletal system" (Paldam & Schjoedt 2016: p.224) The analgesic effects seem to occur as a result of the patient's expectation that they will be cured of their pain. The patients expected a miracle and this expectation seems to be sufficient to provide them with pain relief (see also Wiech et al. 2008).

> experience more severe pain and a greater level of pain-related disability (Janevic et al. 2017; Brekke, Hjortdahl, & Kvien 2002, Eachus et al. 1999).” (Mills, Nicolson & Smith 2019: e275)

People who are unemployed as a result of ill health or disability are also more likely to experience chronic pain (Macfarlane et al. 2015). Ethnicity also seems to have a role to play. In a survey of 500,000 people that responded as black, Asian or of mixed-ethnic background were more likely to report chronic pain than people that responded as white (Macfarlane et al. 2015). Persistent exposure to social adversity has as its consequence increased allostatic load (McEwen 2000), leading to changes in the NEI ensemble, specifically in the HPA axis, that we have argued to be partially constitutive of pain experience. The NEI ensemble learns to predict being in a stressful and threatening situation, which the organism must devote metabolic and energetic resources to combatting. This heighted allostatic load may cause wear and tear on the body, and loss of plasticity leading the organism to get stuck with its predictions of stress and urgent threat to the body.

The EPP theory thus holds the potential of providing an integrated account of the biological, psychological, social, and sociocultural dimensions of pain. We began this section by suggesting that such an integrated account of pain experience is a prerequisite for making intelligible how pain experience could be neurobiologically constituted. Still, one might object that the EPP theory still leaves unaddressed the subjectivity of pain experience because it fails to address the ‘what-it-is-like’ of pain, and its characteristic unpleasantness and painfulness. We turn to this objection in the final section of our paper.

6. Making Progress with the Explanatory Gap

What is it about processes of prediction error minimisation orchestrated by the NEI ensemble that explains why pain hurts? Couldn’t these homeostatic processes of error reduction take place entirely in the dark without pain being experienced or felt by the subject? Our paper seems to leave in place the gap between the subjective character of pain, and objective descriptions of processes in the NEI ensemble.

Our initial response is to repeat that we do not claim that the EPP theory fully closes the gap. Our aim has been more modest. We have started from the assumption that phenomenological reflection provides a distinct type of first-person understanding of pain experience from the science of pain. Our aim has been to show that the EPP theory can meet the mutual constraints requirement, as mandated by the naturalising phenomenology research programme. This requires, on the one hand, making sense of how pain experience as described in phenomenology could be neurobiologically constituted. On the other hand, it requires showing that the phenomenological understanding of pain is consistent with the EPP theory. In what remains of our paper we focus on the first constraint: the challenge of rendering intelligible how the phenomenology of pain experiences could be neurobiologically constituted. If this constraint can be satisfied, it will automatically follow that the phenomenological understanding of pain is consistent with the EPP theory.

We have proposed that the systems that work together to form the NEI ensemble serve the function of maintaining homeostasis. Homeostasis forms the basis for biological values that measure how well or badly the organism is doing at meeting its needs. Consider hunger as an illustration: hunger is an affective state that is bad for the organism because it registers a deviation from the organism's current or future energetic needs. Satiating hunger is good for the organism because this state of affairs amounts to the organism's restoring its homeostatic balance. Thus, potential or actual breaches of homeostasis give rise to hedonically charged affective states for the organism. When a potential or actual deviation from homeostasis is registered, a prediction error will occur that is hedonically valenced. Such prediction errors signal that the organism is in a situation of biological significance. They demand the organism take urgent action to change its situation in order to correct the error.

When do such affective states become conscious? Solms (2021) has recently suggested that consciousness arises when an organism's needs are *felt* - that is to say, when they make demands on the organism to put their energetic resources to work to engage in action (Solms 2021: p.99; see also Solms & Friston 2018).[13] Solms proposes to account for feeling in terms of processes of precision-optimisation that the needs that are most urgent are prioritised. It is through precision weighting that the most salient needs come to be felt. Processes of precision estimation evaluate the uncertainty of prediction errors relative to predictions, thereby allowing for the context-dependent calibration of the influence of prediction errors relative to predictions. A precise prediction error signal is a strong or loud signal, one that demands the organism take heed. When the organism is highly confident in the predicted sensory consequences of its actions, there is no need for conscious feelings to be involved in the guidance of action. The actions necessary for meeting the organism's needs and goals can unfold entirely unconsciously and automatically. Recall our earlier discussion of how nociception for instance will typically ensure that the body stays out of harm's way unconsciously, without the necessity for the intervention of pain (Apkarian 2017). This happens when the actions that are needed to keep the body safe have sensory consequences the brain is able to predict with high-confidence. Feelings only arise when there is expected uncertainty about the fulfillment of an organism's needs. A predicted threat to the integrity of bodily tissues is an example of expected uncertainty. The strength of the feeling of pain will depend on how uncertain the organism is about the continued safety of the body. Expected uncertainty leads to a state of arousal, which mobilises the body's energetic resources, making them ready for actions the organism predicts are likely to reduce uncertainty. The uncertainty of your current situation means that the predicted consequences of your action are now highly salient, and you should be on the lookout for opportunities to reduce uncertainty (i.e. deal with the threat) as fast as possible.

[13] Solms is building on a proposal first outlined in Fotopoulou (2013) in which she proposed that consciousness is associated with the optimisation of precision estimation. On Fotopoulou's proposal, consciousness makes it likely that a person will seek out or avoid opportunities of affective significance that they cannot currently predict with sufficient certainty.

Precision estimation is a function that is hypothesised to be carried out by neuromodulators such as dopamine, noradrenaline, acetylcholine, and serotonin (Parr & Friston 2017). These chemicals spread diffusely throughout the brain modulating the strength of neural connections upwards or downwards, in ways that set priorities on what is to be done next. Arousal, and thus feeling, occurs when there is a sudden change in unexpected uncertainty. Physiological arousal has traditionally been measured, following a tradition initiated by Cannon (1929), by activation of the sympathetic nervous system - the so-called fight-or-flight response. However, Colombetti and Harrison (2018) have shown how such a conception of arousal is overly narrow. In particular they draw upon evidence from psychoneuroendocrinology and psychoneuroimmunology that points to bidirectional interactions between the autonomic, neuroendocrine and immune systems along the lines we have proposed in this paper. Given these bidirectional interactions they conclude that physiological arousal is constituted by the combined activity of all these systems.

Summarising our argument so far, we have argued for two claims: we have suggested that physiological arousal tracks changes in unexpected uncertainty. High precision prediction errors form the basis for felt needs. In the case of pain, such felt needs are the need to protect the body from noxious stimulus that threatens its integrity. Second, we have suggested, in line with Colombetti and Harrison (2018), that processes of physiological arousal are constituted by the NEI ensemble. Thus, we take ourselves to have described how the activity of the NEI could generate experiences that feel painful and unpleasant.

We end with a reminder that there is much more to the phenomenology of pain than unpleasant and painful bodily sensations. Pain experiences are not only experiences *of* the internal states of the body. It is in and through the body that a subject experiences the world. Thus, being in pain contributes to how a subject experiences their relation to the world. What is experienced in pain is the relation between body and world. As we started out by suggesting (in section 1), pain can be productively compared to what Ratcliffe calls "existential feelings". Ratcliffe defines an existential feeling as, at one and the same time, a "feeling of the body" and "a way of finding oneself in the world" (Ratcliffe 2008: p. 2). Notice that the body in question here is what we earlier described as the 'body self' - the body as the medium of subjective experience. We will argue this is something that the EPP theory is also well placed to explain. We will conclude on this basis that the EPP theory satisfies the constraint on an adequate explanation of pain that it make intelligible in the terms of the science of pain, how pain as it is described in phenomenology could be neurobiologically constituted (cf. Northoff 2014, 2018; Thacker & Moseley 2012).

To see how the EPP theory meets this constraint, consider a distinction Clark (2019) makes between what he calls "inward" and "outward" direct cycles of prediction, perception and action. He argues that conscious perceptual experience can be understood in terms of what he describes as the "generative entanglement" of these inward and outward cycles of perception and action. The inward looking cycle is aimed at the control of physiological states internal to the body. The result of this type of predictive processing is interoception - perception of the changing physiological conditions of the body from the inside. Action is here understood as

for the control of perception, maintaining the physiological conditions of the body within the range of values consistent with homeostasis. Recall how Fotopoulou and Tsakiris (2017) argue that the bodily self is rooted in inferential processes that integrate and schematise interoceptive and exteroceptive signals, including those arising from social interaction, to generate a model of the body in its environment. Interoceptive signals serve the function of informing the organism how well or badly it is doing at meeting its needs. The inner feelings of arousal combine with exteroceptive and proprioceptive information to "form the basis of subjectivity and the self" (Fotopoulou & Tsakiris 2017: p.15).

Outward looking cycles of prediction do the work of selecting actions that control exteroceptive perception thereby maximising the probability of bringing about the sensory outcomes that are valued and desired. The world that is encountered in perception depends upon an interoceptively mediated sense of a creature's current needs and its embodied state. Salient features of the world are features that, when sampled, minimise uncertainty, and thus contribute to satisfying the creature's needs. Agent's try to sample the world so as to conform their expectation of flourishing. In doing so, they bring forth or enact a world that is significant to them (Varela et al. 1991). The entanglement of inner and outer directed cycles of prediction has the consequence that what a creature perceives is "nuanced by their own bodily needs and states" (Clark 2019: p.4).

We take Clark's notion of generative entanglement of inner and outer-directed cycles of prediction to be a key step that is required to make intelligible how pain could be transformative of a person's being in the world. Once we think of these processes of prediction that take place within the NEI ensemble as structuring the subject's perception of the world, we suggest this makes sense of how pain could transform how a person finds themselves situated in the world. However, we would like to offer one important corrective: Clark presents inner directed cycles of prediction as the brain controlling the physiological conditions of the body in ways that maintain homeostasis. In other words, he retains a brainbound picture of prediction. We have argued by contrast that the whole NEI ensemble works predictively. Processes of prediction span the whole body and are not confined to the brain. The bodily processes that are predicting how the agent is faring in its engagement with the world are also constitutive of the subject's concerned first-person perspective on the world.

Conclusion

We have outlined a perspective on pain based on the embodied predictive processing (EPP) theory that may allow neurobiology to make progress in making sense of the subjectivity of pain. According to the EPP theory, pain is constituted by the whole NEI ensemble that operates to fulfill the prediction of the continued functional and structural integrity of the body. Pain cannot be reduced to nociception or decomposed into sensory-discriminatory, affective-motivational and cognitive-evaluative elements. Instead we have argued pain is the outcome of predictive processing that takes place in the whole neural axis in continuous reciprocal interaction with the immune system, the neuroendocrine and the autonomic

system. All of these systems are working together as an integrated whole. Pain occurs when all of these systems together conclude that the prediction of body integrity is likely to depart from reality. 'Body integrity' refers to the states of the body the organism should return to under a wide variety of different conditions if the organism is to continue to exist. Pain is therefore generated by homeostatic processes that aim to maintain the body in these states under conditions of constant change. Pain provides the organism with feedback that it is diverging in potentially dangerous ways from these (adjustable) setpoints.

We have argued that the neurobiological processes that constitute experience take place in all of the systems that maintain the homeostasis of the body. We have proposed to think of the predictive processing that takes place within each of these systems that make up the NEI ensemble as producing and maintaining a nesting of Markov blankets. At the smallest scale to the largest scale processes of prediction error minimisation play out that separates but also connects and integrates these systems. This nesting of Markov blankets makes sense of how the body could be the subject's point of view on the world. The EPP theory claims that pain is not only in the brain but is a state of the whole body that prioritises the actions the organism needs to undertake to return the body to the state of healthy flourishing that is expected.

References

Anchisi, D., Zanon, M. 2015. A Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia. PLoS One 10:1-20.

Anderson, M. 2017. Of Bayes and bullets: An embodied, situated, targeting-based account of predictive processing. Philosophy and predictive processing. Frankfurt am Main: MIND Group.

Andrews, M. 2020. The math is not the territory: navigating the Free Energy Principle. Preprint, downloaded from: http://philsci-archive.pitt.edu/18315/ Accessed 10.01.2021.

Apkarian, V. 2017. Advances in the neuroscience of pain. In J. Corns (Ed.) *The Routledge Handbook of Philosophy of Pain*

Atlas, L.Y., Bolger, N., Lindquist, M.A., & Wager, T.D. 2010. Brain mediators of the predictive cue effects on perceived pain. *The Journal of Neuroscience* 30(39): 12964-12977.

Auvray, M., Myin, E., & Spence, C. 2010. The sensory-discriminative and affective-motivational aspects of pain. *Neuroscience & Biobehavioural Reviews* 34: 214-223.

Aydede, M. 2006. A critical and quasi-historical essay on theories of pain. In M. Aydede (Ed.) *Pain: New Essays on its Nature and the Methodology of its Study.* Cambridge, MA: MIT Press, pp.1-58.

Baliki, M.N. and Apkarian, A.V. 2015. Nociception, pain, negative moods, and behavior selection. *Neuron* 87: 474–491.

Benedetti, F. 2013. Placebo and the new physiology of the doctor-patient relationship. *Physiological Review* 93: 1207-1246,

Blyth, F. 2008. Chronic pain - is it a public health problem? *Pain* 137: 465-6

Brekke, M., Hjortdahl, P., & Kvien, T. 2002. Severity of musculoskeletal pain: relationship to socioeconomic inequality. *Social Science & Medicine* 54: 221-8

Bruineberg, J., Dolega, K, Dewhurt, J. & Baltieri, M. (2020). The emperor's new Markov blankets. Preprint, dowloaded from http://philsci-archive.pitt.edu/18467/ Accessed 15.01.2021.

Buchel, C., Geuter, S., Sprenger, C., & Eippert, F. 2014. Placebo analgesia: a predictive coding perspective. *Neuron*, 81(6):1223-39. doi: 10.1016/j.neuron.2014.02.042

Buczsáki, G. 2019. *The Brain from Inside Out.* Oxford, UK: Oxford University Press.

Cannon, W. B. 1929. *Bodily Changes in Pain, Hunger, Fear and Rage*, 2nd edn. (Boston, MA: Charles Branford).

Cao, R. 2020. New labels for old ideas: predictive processing and the interpretation of neural signals. *Review of Philosophy and Psychology* 11: 517-46.

Carel, H. 2013. Bodily doubt. *Journal of Consciousness Studies,* 20(7-8): 178-97.

Chalmers, D. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.

Chapman, C.R., Tuckett, R.P., & Song, C.W. 2009. Pain and stress in a system's perspective: reciprocal neural, endocrine and immune interactions. *J. Pain* 9(2): 122-145.

Ciaunica, A., Petreca, B., Fotopoulou, A., & Roepstorff, A. 2021a. Whatever Next and Close to my Self – The Transparent Senses and the 'Second Skin': Implications for the Case of Depersonalisation. *Frontiers in Psychology*.

Ciaunica, A., Constant, A., Priessl, H. & Fotopoulou, K. 2021b. The first prior: from co-embodiment to co-homeostasis in early life. *Consciousness & Cognition*, 91(5): 103117

Ciaunica, A., & Fotopoulou, A. 2017. The Touched Self: Psychological and Philosophical Perspectives on Proximal Intersubjectivity and the Self. In C. Durt, T. Fuchs, & C. Tewes (Eds.), *Embodiment, Enaction, and Culture—Investigating the Constitution of the Shared World*. (Cambridge MA: MIT Press), pp.

Clark, A. 2019. Consciousness as generative entanglement. *The Journal of Philosophy*, 116 (12): 645-662. ISSN 1939-8549

Clark, A. 2016. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. New York: Oxford University Press.

Clark A. 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioural & Brain Sciences* 36(3): 181-204. doi: 10.1017/S0140525X12000477

Colombetti, G. & Zavala, E. 2019. Are emotional states based in the brain? A critique of brainocentrism from a physiological perspective. *Biology & Philosophy,* 34(45): 1-20.

Colombetti, G. & Harrison, N. 2018. From physiology to experience: enriching existing conceptions of "arousal" in affective science. In M. Tsakiris & H. De Preester (Ed's) *The Interoceptive Mind: From Homeostasis to Awareness.* (Oxford, UK: Oxford University Press), pp. 245-258.

Colombetti, G. 2014. *The Feeling Body: Affective Science Meets the Enactive Mind*. (Cambridge, MA: MIT Press).

Coninx, S. & Stilwell, P. 2021. Pain and the field of affordances: an enactive approach to acute and chronic pain. *Synthese,* https://doi.org/10.1007/s11229-021-03142-3

Dennett, D. C. 2015. Why and how does consciousness seem the way it seems? In T. Metzinger & J. M. Windt (Eds). Open MIND: 10(T). Frankfurt am Main: MIND Group.

Dennett. D.C. 2003. Who's on first? Heterophenomenology explained. *Journal of Consciousness Studies,* 10(9-10): 19-30.

Dennett, D.C. 1991. *Consciousness Explained*. Boston: Little Brown.

Dewey, J. 1896. The reflex arc concept in psychology. *Psychological Review* 3: 357-370.

Eachus, J., Chan, P., Pearson, N., Propper, C., & Davey-Smith G. 1999. An additional dimension to health inequality: disease severity and socioeconomic position. *Journal of Epidemiology and Community Health* 53: 603-11

Engel, G. L. 1977. The need for a new medical model: A challenge for biomedicine. *Science*, 196(4286), 129–136.

Fayaz, A., Croft, P., Langford, R., Donaldson, J., & Jones, G. 2016. Prevalence of chronic pain in the UK: a systematic review and meta-analysis of population studies. *BMJ Open* 6 (6) e010364; DOI: 10.1136/bmjopen-2015-010364

Fotopoulou, A., Von Mohr, M., & Krahé, C., Y. 2021, preprint. Affective regulation through touch: homeostatic and allostatic mechanisms. *PsyArXiv*, 10.31234/osf.io/ksj3x, Accessed 04/07/2021.

Fotopoulou, A., Tsakiris, M., 2017. Mentalizing homeostasis: The social origins of interoceptive inference. *Neuropsychoanalysis*, 19, 3-28.

Fotopoulou, A, 2013. Beyond the reward principle: consciousness as precision seeking. *Neuropsychoanalysis*, 15, 33-38.

Friston, K.J. 2013. Life as we know it. *Journal of the Royal Society Interface*, 10(86). doi:10.1098/rsif.2013.0475

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

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

Gadamer, H-G. 1996. *The Enigma of Health: The Art of Healing in a Scientific Age*. (Translate J. Gaiger and N. Walker. Stanford, CA: Stanford University Press).

Gallagher, S. & Zahavi, D. 2008. *The Phenomenological Mind: An Introduction to Philosophy of Mind and Cognitive Science.* (London and New York, NY: Routledge, 2008).

Gallagher, S. 2005. *How the Body Shapes the Mind*. (Oxford, UK: Oxford University Press).

Gallagher, S. 2000. Philosophical Conceptions of the Self: Implications for Cognitive Science. *Trends in Cognitive Sciences,* 4 (1): 14–21.

Grahek, N. 2007. *Feeling Pain and Being in Pain.* (Cambridge, MA: The MIT Press),

Heidegger, M. 2001. *Zollikon Seminars: Protocols—Conversations—Letters*. (Translated F. Mayr and R. Askazy. Evanston, IL: Northwestern University Press).

Hipólito, I., Ramstead, M.J.D., Convertino, L., Bhat, A., Friston, K. & Parr, T. 2021. Markov blankets in the brain. Neuroscience and Biobehavioural Reviews, 125: 88-97

Hohwy, J. 2020. New directions in predictive processing. *Mind & Language* 35(2): 209-223.

Hohwy J. 2013. *The Predictive Mind.* Oxford, UK: Oxford University Press.

Hohwy, J., Roepstorff, A., & Friston, K. 2008. Predictive coding explains binocular rivalry. *Cognition* 108(3): 687-701.

Hurter, S., Paloyelis, Y., Williams, A. C., and Fotopoulou, A. 2014. Partners' empathy increases pain ratings: Effects of perceived empathy and attachment style on pain report and display. *Journal of Pain*, 15(9), 934– 44.

Jackson, F. 1982. Epiphenomenal qualia. *Philosophical Quarterly*, 32: 127-136.

Janevic, M.R., McLaughlin, S.J., Heapy, A.A., Thacker, C., & Piette, J.D. 2017. Racial and socioeconomic disparities in disabling chronic pain: findings from the Health and Retirement Study. *Journal of Pain* 2017; 18: 1459-67

Kiverstein, J. & Kirchhoff, M. 2022 (submitted): Scientific realism about Friston blankets without literalism. Commentary on Bruineberg et al. The emperor's new blankets. *Behavioural and Brain Sciences*

Kirchhoff, M, Robertson, I., & Kiverstein, J. (In progress). The literalist fallacy and the free energy principle: on model-building, scientific realism and instrumentalism.

Kirchhoff, M.D. & Kiverstein, J. 2019. How to Determine the Boundaries of the Mind: A Markov Blanket Proposal. *Synthese*, 1–20.

Kirchhoff, M., Parr, T., Palacios, E., Friston, K., & 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

Kusch, M., & Ratcliffe, M. J. 2018. The world of chronic pain: A dialogue. In K. Aho (Ed.), *Existential Medicine*. Rowman & Littlefield International, pp.61-80.

Krahé, C., Y. Paloyelis, H. Condon, P. M. Jenkinson, S. C. R. Williams, and A. Fotopoulou. 2015. Attachment style moderates partner presence effects on pain: A laser-evoked potentials study. *Social Cognitive and Affective Neuroscience* 10: 1030–1037.

Krahé, C., Springer, A., Weinman, J., Fotopoulou, A., 2013. The social modulation of pain: others as predictive signals of salience – a systematic review. *Frontiers Human Neuroscience*, doi: 10.3389/fnhum.2013.00386

Kripke, S. A. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Leder, D. 1990. *The Absent Body.* Chicago, IL: University of Chicago Press.

Legrand, D. 2006. The bodily self: the sensorimotor roots of pre-reflective self-consciousness. *Phenomenology and Cognitive Sciences* 5: 89–118.

Leknes, S., Berna, C., Lee, M.C., Snyder, G.D., Biele, G., & Tracey, I. 2013. The importance of context: when relative relief renders pain pleasant. *Pain*, 154: 402-410.

Levine, J., 1983. Materialism and qualia: the explanatory gap problem. *Pacific Philosophical Quarterly* 64, 354–361.

Macfarlane, G., Beasley, M., Smith, B., Jones, G., & Macfarlane, T. 2015. Can large surveys conducted on highly selected populations provide valid information on the epidemiology of common health conditions? An analysis of UK Biobank data on musculoskeletal pain. *British Journal of Pain*, 9: 203-12

Maly, A. & Vallerand, A.H. 2018. Neighborhood, socioeconomic, and racial influence on chronic pain. *Pain Management Nursing* 19: 14-22

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

Melzack, R. & Katz, J. 2013. Pain. *WIREs Cognitive Science*, 4:1–15. doi: 10.1002/wcs.1201

Melzack, R. 1999. From the gate to the neuromatrix. *Pain* 6 (Suppl): S121-S126.

Melzack, R. 1996. Gate control theory: on the evolution of pain concepts. *Pain Forum* 5: 128-138.

Merleau-Ponty, M. 1945/2012. *Phenomenology of Perception*. Translated by D. Landes, London, UK: Routledge

Mills, S.E.E., Nicolson, K.P & Smith, B.H. 2019. Chronic pain: a review of its epidemiology and associated factors in population-based studies. British Journal of Anaesthesia 123(2): e273-e283

Minkowski, E. 1958. Findings in a case of schizophrenic depression. (Translated by B. Blis, In R. May, E. Angel, and H. Ellenberger (Ed's.), *Existence.* New York: Simon and Schuster), pp. 127–138.

Moseley, G.L. 2008. Placebo effect: reconceptualising placebo. *BMJ,* 336 (7653): 1086.

Moseley G.L. & Arntz, A. 2007. The context of a noxious stimulus affects the pain it evokes. *Pain*, 133 (1):64-71.

Nagel, T., 1974. What is it like to be a bat? *The Philosophical Review* 83, 435–450.

Noë, A. 2016. Sensations and situations: a sensorimotor integrationist approach. *Journal of Consciousness Studies* 23 (5-6): 66-80.

Northoff, G. 2018. *The Spontaneous Brain: From the Mind-Body to the World-Brain Problem.* (Oxford, UK: Oxford University Press).

Northoff, G. 2014. *Unlocking the Brain. Volume I: Coding*. (Oxford, UK: Oxford University Press.)

Ongaro, G. & Kaptchuk, T.J. 2019. Symptom perception, placebo effects, and the Bayesian brain. Pain, 160(1): 1-4.

Palacios, E.R., Razi, A., Parr, T., Kirchhoff, M.D., & Friston, K. 2020. On Markov blankets and hierarchical self-organisation. *Journal of Theoretical Biology*, 486 110089-1-110089-14.

Paldam, E. & Schjoedt, U. 2016. Miracles and pain relief: experienced health effects of charismatic prayer healing in a large collection of Christian testimonies. *Archives for the Psychology of Religion*, 38: 210-231.

Paloyelis, Y., Krahé, C., Maltezos, S., Williams, S. C., Howard, M. A., and Fotopoulou, A. (2016). The analgesic effect of oxytocin in humans: A double-blind, placebo-controlled cross-over study using laser-evoked potentials. *Journal of Neuroendocrinology*, 28(4).

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

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* San Francisco: Morgan Kaufmann Publishers.

Poleshuck, E. & Green, C. 2008. Socioeconomic disadvantage and pain. *Pain* 136: 235e8

Quirin, M., Kuhl, J., & Düsing, R. 2011. Oxytocin buffers cortisol responses to stress in individuals with impaired emotion regulation abilities. *Psychoneuroendocrinology*, 36(6): 898-904.

Ratcliffe, M. 2008. *Feelings of Being: Phenomenology, Psychiatry and the Sense of Reality*. Oxford University Press.

Rowlands, M. 2010. *The New Science of the Mind: from Extended Mind to Embodied Phenomenology.* (Cambridge, MA: MIT Press).

Solms, M. 2021. *The Hidden Spring: A Journey to the Source of Consciousness.* (London, UK: Profile Books).

Solms, M. & Friston, K. 2018. How and why consciousness arises: considerations from physics and physiology. *Journal of Consciousness Studies,* 25: 202-38.

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

Sterling, P. and Eyer, J. 1988. Allostasis: A new paradigm to explain arousal pathology. In: K. Fisher and J. Reason (Ed's), *Handbook of Life Stress, Cognition and Health*. (Hoboken, NJ: John Wiley & Sons), pp. 629– 49.

Svenaeus, F. 2015. The phenomenology of chronic pain: embodiment and alienation. *Continental Philosophy Review* 48: 107-22.

Tabor, A., Thacker, M.A., Moseley, G.L., & Körding, K.P. 2017. Pain: a statistical account. *PLoS Computational Biology* 13(1): e1005142. doi:10.1371/journal.pcbi.1005142

Tracey, I. & Mantyh, P.W. 2007. The cerebral signature for pain perception and its modulation. *Neuron* 55(3)*:*377-91. doi: 10.1016/j.neuron.2007.07.012.

Thompson, E. 2007. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind.* (Cambridge, MA: Harvard University Press).

Tye, M. 1995. *Ten Problems of Consciousness.* Cambridge, MA: MIT Press.

Vaso, A., Adahan, H-M., Gjika, A., Zahaj, S., Zhurda, T., Vyshka, G & Devor, M. 2014. Peripheral nervous system origin of phantom limb pain. *Pain*, 155(7): 1384-1391. doi: 10.1016/j.pain.2014.04.018.

Varela F., E. Thompson, and E. Rosch. 1991. *The Embodied Mind: Cognitive Science and Human Experience.* (Cambridge, MA: MIT Press).

Von Mohr, M. & Fotopoulou, A. 2018. The cutaneous borders of interoception: active and social inference of pain and pleasure on the skin. In M. Tsakiris & H. De Preester (Ed's) *The Interoceptive Mind: From Homeostasis to Awareness.* (Oxford, UK: Oxford University Press), pp.102-120.

Varela, F. 1996. Neurophenomenology: a methodological remedy to the hard problem. *Journal of Consciousness Studies*, 3, 330–350.

Wall, P.D., 1999. *Pain: The Science of Suffering*. Weidenfeld & Nicolson, London.

Wiech, K. 2016. Deconstructing the sensation of pain: The influence of cognitive processes on pain perception. *Science* 354(6312): 584-587.

Wiech, K., Farias, M., Kahane, G., Shackel, N. Wiebke, T., & Tracey, I. 2008. An fMRI study measuring analgesia enhanced by religion as a belief system. *Pain* 139(2): 467-76.

Wheeler, M. 2013. Science friction: phenomenology, naturalism and cognitive science. *Royal Institute of Philosophy Supplement*, 72: 135-167.

Wheeler, M. 2005. *Reconstructing The Cognitive World: The Next Step* (Cambridge, MA: MIT Press).

Zahavi, D. 2005. *Subjectivity and Selfhood: Investigating the First-Person Perspective*. (Cambridge, MA: MIT Press)