
Embodied Decisions and the Predictive Brain

Christopher Burr

A cognitivist account of decision-making views choice behaviour as a serial process of deliberation and commitment, which is separate from perception and action. By contrast, recent work in embodied decision-making has argued that this account is incompatible with emerging neurophysiological data. We argue that this account has significant overlap with an embodied account of predictive processing, and that both can offer mutual development for the other. However, more importantly, by demonstrating this close connection we uncover an alternative perspective on the nature of decision-making, and the mechanisms that underlie our choice behaviour. This alternative perspective allows us to respond to a challenge for predictive processing, which claims that the satisfaction of distal goal-states is underspecified. Answering this challenge requires the adoption of an embodied perspective.

Keywords

Action-oriented representation | Active inference | Decision-making | Distributed consensus | Embodied decisions

Acknowledgements:

A special thanks to Max Jones and Richard Pettigrew for their comments on earlier drafts, and to two anonymous reviewers and the editors for their invaluable feedback. I believe the paper is in a far better form as a result of their help. Thanks also to Andy Clark and Mark Miller for insightful discussions and helpful recommendations during the earliest stages of this work.

1 Introduction

Consider the following situation. You are busy writing a paper, which has a deadline that is fast approaching. It is important that your time is used productively to ensure that this deadline is met. You have already been writing for a couple of hours without a break and notice that your productivity is decreasing. What should you do? As it's almost lunchtime, you consider the fact that your tiredness is a product of your hunger. However, there is a chance that if you take too long a break you will be unable to regain your train of thought. Perhaps you should continue working for a bit longer, and have lunch once the word count has been reached, or maybe it's better to compromise and take a short break now to make a coffee. What about doing something else entirely? Far from it being a simple matter of choosing between a few well-delineated options, it appears that you must also determine what options are available to you. A complete account of how we make decisions, among other things, should be able to explain these challenges. The main concern of the paper is the possibility of situating such an account of the mechanisms that underlie decision-making within the framework of predictive processing (PP).

It has been claimed that the PP framework has the capacity to unify a wide range of different phenomena, such as perception, cognition and action (see [Clark 2016](#); [Hohwy 2013](#), for introductions). According to Hohwy, “prediction error minimization is the only principle for the activity of the brain” ([Hohwy 2016](#), p.2). This principle can be couched in terms of statistical inference, and provides a functional ground for unification—one which Hohwy sees as resisting an embodied reading ([Hohwy forthcoming](#)). Alternatively, Clark has also claimed that PP may offer a unifying perspective on many of the brain's capacities, including perception, action, learning, inference, and cognitive control, but opts for a less neurocentric perspective ([Clark 2016](#)). Favouring the latter approach, this paper con-

cerns itself with exploring how the PP framework can contribute to our understanding of the process of *decision-making* from an embodied perspective, thus supporting Clark's view.

We also focus on a specific problem that has been raised by (Basso 2013). This is the problem of the *underspecification* of action sequences, which are initially represented by some distal goal-state. In PP, a future goal-state is essentially a higher-level prediction used as a means of enabling action through the reduction of proprioceptive prediction-error (i.e. Active Inference) (see Clark 2016; Hohwy 2013, for introductions). However, as Basso states:

[...] the future goal state created in the beginning is accurate only in some particular circumstances (i.e., when both the task and algorithm are well-defined). In most cases, people are used to facing underspecified tasks in which a future goal state cannot be employed to derive the intermediate states (Basso 2013, p. 1)

This challenge is also important for an account of decision-making, which is traditionally assumed to operate as a deliberation over representations of expected goal-states associated with the performance of some action. However, if the sequence of actions is underspecified this raises a challenge of what constrains the selection of one set of possible actions over another. For example, imagine you decide to take a break from writing and go and make lunch. There are myriad ways this initial (perhaps vaguely specified) goal-state could be satisfied, and some may be less effective than others (e.g. deciding to make a long lunch that keeps you from returning to writing for a prolonged period of time).

We will argue that recent work on embodied decision-making (Cisek and Pastor-Bernier 2014; Lepora and Pezzulo 2015), offers both an interesting perspective on the neural mechanisms that underlie decision-making, and may also be able to provide important insights for developing PP. However, we will also argue that doing so requires adopting an embodied perspective if we are to respond to the above challenge. We will argue that a neurocentric perspective is too narrow to acknowledge the many important ways that our choice behaviour is constrained by our bodies and the world.

2 Embodied Decisions

Consider another decision. You must choose between two routes to work where Route A takes you through a city that has a high risk of heavy traffic but is short in distance. The other route is less likely to be affected by the increased congestion, but is longer than the former. Suppose you know from previous experience that, given the time you are leaving, it is more likely that the traffic will be light, and your preference is always for the shortest time spent travelling. What should you do?

Table 1: A simple representation of a decision under risk.

	Heavy Traffic (30%)	Light Traffic (70%)
Route A	24 minutes	14 minutes
Route B	18 minutes	17 minutes

Table 1 represents a *decision under risk*, as the agent has full knowledge of the available options and probabilities attached to the relevant states. In situations like this, deciding what to do is relatively straightforward, and a number of decision rules exist as proposals for norms of rationality. For example, the *principle of maximizing expected utility* would suggest taking Route A, as the following demonstrates that it has the shortest expected duration (and therefore the greatest expected utility, assuming that utility is a negative linear transform of duration):

$$\text{Expected Duration of Route A} = 0.3 \times 24 + 0.7 \times 14 = 17$$

$$\text{Expected Duration of Route B} = 0.3 \times 18 + 0.7 \times 17 = 17.3$$

Savage famously referred to these situations as ‘small worlds’, where it is possible to “look before you leap”, by which he meant an agent has knowledge of the states of the world and all of the options available to her (Savage 1954/1972, p. 16). Even in the case where the probabilities attached to the states are unknown (decisions under uncertainty), many decision-theoretic norms (e.g. dominance and subjective expected utility maximization) exist to help guide this process. However, unlike small worlds, the real world is not so neatly circumscribed. In contrast, everyday decisions can be viewed as ‘large worlds’, where agents may be unaware of all the relevant information, including an uncertainty of what options are available. This level of uncertainty is a challenge for decision theory, as the possibility of framing a genuine decision problem requires that an agent already have options to deliberate over. Even hallmarks of rationality such as Bayesianism have been criticised as inapplicable in these types of large worlds (Binmore 2008).

It may be argued that this is not really a problem for decision theory *per se*. The issue of determining options is a problem for the perceptual system to solve. Perception is faced with the task of specifying and constructing a representation of features of the environment, which can then be used as the basis for making decisions (along with abstract representations of related decision variables such as potential risk). According to this traditional view of decision theory, decision-making is seen as a prototypical cognitive task, which can be decomposed into a process of *deliberation* (i.e. calculating the values of the relevant decision variables) and *commitment* (i.e. selecting an action). Furthermore, motor behaviour is simply the manner in which a decision is reported, and can be used to reveal an agent’s preferences (Sen 1971).




Figure 1 - The Classical Sandwich Model

This account of decision-making is based on a number of *cognitivist* assumptions, which are nicely captured by Hurley’s critique of what she termed the “classical sandwich model” of the mind (Hurley 1998, p. 401). In this model (Figure 1), the outer slices of perception and action are peripheral to the inner filling of cognition, and thus separate from one another. They are also separate from cognition, which interfaces between perception and action. First, perception builds a reconstructive representation of features of the external world. These discrete, abstract representations are then transformed by cognitive processes into a motor plan for action, according to the system’s beliefs and desires, and subsequently carried out by the system’s effectors. Within this model, decision-making would reside within the middle box, and deliberation and commitment could take place in some ‘central executive region’ such as the prefrontal cortex, which integrates relevant information from other systems such as working memory (Baddeley 1992).

Hurley saw a number of problems or limitations with this account. For example, this serial process would be insufficiently dynamic to cope with the time pressures of a constantly changing environment. In the time taken to construct a representation and plan an action by integrating the necessary information, the environment may have changed, which would render the current model (and any actions based on it) inaccurate. This worry about the urgency of performing an action in ecologically-valid scenarios is particularly pressing when applied to the case of decision-making. In traditional decision theory, models of decision-making do not incorporate the time constraints of agents, and therefore fail to account for a number of additional pressures that agents face.

The cognitivist view of decision-making then, highlighted by the classical sandwich model, leads to a tendency to think of sensorimotor control in terms of the transformation of input representations into output representations through a series of well-demarcated, encapsulated processing stages. It also often leads to the assumption that key decision variables are encoded as an abstract value in some central executive region separate from sensorimotor processes (Levy and Glimcher 2012; Padoa-Schioppa 2011). Cisek has argued that this picture is hard to reconcile with a growing body of neurophysiological data (Cisek 2007; Cisek and Kalaska 2010). He claims that key functions of decision-making, which the cognitivist would expect to be neatly delineated, in fact appear widely distributed throughout the brain, including key sensorimotor regions. To accommodate this otherwise anomalous data he proposes a notion of *embodied decisions*, which have a number of properties that are quite different to the kinds of decisions modelled by traditional decision theory (Cisek and Pastor-Bernier 2014). The following sections will explore each of them.

2.1 Decision-Making as a Distributed Consensus

Cisek proposes the *affordance competition hypothesis* (ACH) as a model that aims to explain both the cognitive and neural processes implicated in decision-making, and one which attempts to make sense of emerging neurophysiological data that conflicts with traditional decision theory (Cisek 2007; Cisek and Kalaska 2010). According to the ACH, decisions emerge from a distributed, probabilistic competition between multiple representations of possible actions, and importantly overlaps with sensorimotor circuits (see Figure 1 of Cisek 2007, p. 1587). To expound this view, a number of components require clarification—evidence supporting the claims are reviewed in (Cisek and Kalaska 2010; Cisek 2012).

Cisek's focus on the *distributed* manner of decision-making stands in obvious contrast to the earlier cognitivist framework, and also to other models that propose that decision-making occurs downstream of the integration of multiple sources of information, which yields a common representation of abstract value (Padoa-Schioppa 2011). Instead, according to the ACH, the sensorimotor system is continuously processing sensory information in order to *specify* the parameters of potential actions, which compete for control of behaviour as they progress through a cortical hierarchy, while at the same time, other regions of the brain provide biasing inputs in order to *select* the best action (Cisek and Kalaska 2010). These processes of *specification* and *selection* occur simultaneously and continuously, and are not localisable to a specific region. Rather, the *competition* occurs by way of mutual inhibition of neural representations, which specify the parameters of potential actions, until one suppresses the others. At such a time, a global, distributed consensus emerges.

Integral to this process is the role of continuously biasing influences (e.g. rule-based inputs from prefrontal regions, reward predictions from basal ganglia, and a range of further biasing variables from sub-cortical regions). Each of these biasing inputs contribute their votes to the selection process. As the authors state:

[...] the decision is not determined by any single central executive, but simply depends upon which regions are the first to commit to a given action strongly enough to pull the rest of the system into a 'distributed consensus'. (Cisek and Pastor-Bernier 2014, p. 4)

Again, this idea stands in stark contrast to the cognitivist picture, where the perceptual system merely processes information in order to construct a perceptual representation, which provides the evidence about the environment needed to make decisions. Rather, here we have the *beginnings* of an account that explains how the relevant options of a decision problem are being selected in parallel with the specification of sensorimotor information:

[...] although traditional psychological theories assume that selection (decision making) occurs before specification (movement planning), we consider the possibility that, at least during natural

interactive behavior, these processes operate simultaneously and in an integrated manner. (Cisek and Kalaska 2010, p. 277)

One of the specific claims made by Cisek and Pastor-Bernier is that, as part of the competitive process, the brain is simultaneously specifying and selecting among representations of multiple *action opportunities* or affordances, which compete within the sensorimotor system itself (Cisek and Pastor-Bernier 2014).¹ These representations serve as indications of the possible actions available in the agent's environment, rather than as objective, organism-independent properties of the world.

For example, Cisek and Kalaska discuss recordings taken from the dorsal premotor cortex (PMd) in monkeys during a reaching task (Cisek and Kalaska 2010). In the experiment, monkeys were presented with two potential reaching actions by way of spatial cues, where one would later be indicated as the correct choice (using a non-spatial cue). During a memory period, where the spatial cues were removed and the future correct choice was uncertain, recorded activity in the PMd continued to specify both directions simultaneously, suggesting an anticipatory nature for the neural activity. When the information specifying the correct choice was eventually presented, activity relating to the respective action was strengthened, and the unwanted action was suppressed.

Importantly, this process occurs within the *same system* that is used to prepare and execute the movement associated with the action representations. Furthermore, Cisek and Kalaska state that the task design allowed for the monkeys to exploit a different (cognitivist) strategy, where the target locations are stored in a more general-purpose working memory buffer, distinct from motor representations, and converted to a motor plan after a decision has been made. However, though conceptually possible, the findings did not seem to support this latter view. Instead, the study seems to point to a need for representations that encode predictive (or anticipatory) action opportunities, rather than abstract representations that specify the state of the world independently from an agent's particular goals and capacities.

The ACH also makes key predictions that can be tested in future experiments. For example, it predicts that actions that are further apart from one another will show stronger mutual inhibition than those that are closer together. This is because action representations are specified in terms of spatial parameters (Cisek and Kalaska 2010), which means that a decision between similar actions (with overlapping representations) can be encoded using a weighted average. This weighted average could evolve over time, initially tolerating some uncertainty between two future actions, whereas drastically different options could not. A prediction made by the ACH is, therefore, that if one records from neural cells related to a given option, while modulating the desirability of a different option, the gain of that modulation will be strongest when the other option is most dissimilar to the one coded by the recorded cell (Cisek and Pastor-Bernier 2014, p. 5).

Importantly, this view stands in contrast to decision-theoretic accounts that model humans as making decisions between different options by integrating the relevant factors into a single variable, such as subjective utility (Levy and Glimcher 2012). For example, some have argued that the orbitofrontal cortex (OFC) and ventromedial prefrontal cortex (vmPFC) could integrate the relevant information and encode such an abstract value (Padoa-Schioppa 2011). However, this view is hard to reconcile with neurophysiological data (Cisek and Kalaska 2010). Given these findings, as well as others (cf. Klaes et al. 2011; Pastor-Bernier and Cisek 2011), it seems that at least part of a prototypical cognitive process (decision-making) is inextricably intertwined with sensorimotor control, suggesting a blurring of the boundaries between perception, action and cognition. However, as Cisek himself notes,

¹ The use of the term 'affordances' will be familiar to those with an understanding of Gibson's work in ecological psychology (Gibson 1979). However, Cisek states his account is directly influenced by a wider commitment to the tradition of ecological psychology, and relies on different notions (i.e. action opportunity, affordance, pragmatic representation and sensorimotor plan) throughout his work to emphasise similar points (Cisek and Kalaska 2010). Given the controversy on how best to understand the metaphysical commitments of Gibson's notion (cf. Chemero 2003; Chemero 2011), we will focus on the less demanding notion of action opportunities instead of affordances.

“we are capable of making decisions that have nothing to do with actions, and in such situations the decision must be abstract.” (Cisek 2012, p. 927) Whether we have to accept this limitation of the ACH, or whether PP can develop on its findings and scale up to include purportedly abstract decisions, is a primary focus of this article, and will be dealt with specifically in section 4.

2.2 Simultaneous Decisions

Consider the earlier example of attempting to write a paper, where deciding what to do was plagued by an indeterminacy in knowing whether all possible actions had been considered. Cisek and Pastor-Bernier claim these types of everyday decisions are archetypical kinds of *simultaneous* decisions. By this they mean that there are a multitude of possible actions an agent could take at any one time, and therefore, we are continuously deciding what to do.

To attempt to explain this, Cisek and Kalaska review research on the pervasive effect of attentional modulation, which supports the idea that activity in the visual system is strongly influenced by attentional modulation, even in familiar and stable environments (Cisek and Kalaska 2010). This is usually recorded as an enhancement of activity correlated with the attended regions of space, and a suppression of activity from the unattended regions. For example, studies by Stefan Treue show the ubiquitous effects of attentional modulation in primate visual cortex (Treue 2001; Treue 2003). This attentional modulation results in the enhancement of activity towards behaviorally relevant stimuli, along with a corresponding suppression of those cells tuned to non-attended spatial features. Attending only to those features of the world that are deemed salient may not appear to be a rational strategy. However, echoing the sentiments of work in ecological rationality, Treue acknowledges that it is nevertheless “an effective use of limited processing resources.” (Treue 2003, p. 428)

Despite the attractiveness of appealing to saliency and attention on ecological grounds, flexible and adaptive choice behaviour requires reciprocal communication between affective and sensorimotor regions. This is because determining what is salient requires an awareness of the changing demands of both the *external* and *internal* environment, in order to respond to the homeostatic demands of the agent. At present this is one area that is left underdeveloped by Cisek, Kalaska and Pastor-Bernier. In section 3.2, we will see how this aspect of embodied decisions can be developed further, by exploring the unique roles of attention within predictive processing, and its emphasis on interoceptive inference.

2.3 Dynamic Choice Behaviour

Finally, Cisek and Pastor-Bernier argue that the notion of embodied decisions requires a dynamic account of the decision-making process. They claim that the continual processing of noisy or uncertain sensory information after commitment, suggests that agents continue to deliberate during the overt performance of a task. This means the agent constantly monitors the overt performance of their actions through sensory feedback (e.g. proprioception). The existence of this evidence, they argue, requires the revision of some commonly used formal models in decision theory that are unable to account for this post-selection monitoring and alteration.

Lepora and Pezzulo also acknowledge this requirement, and claim that *action performance* should be considered a proper part of a dynamic model of decision-making, rather than being understood as merely the output of the decision process (Lepora and Pezzulo 2015). As a proof of principle to support this claim, they develop a computational model, which they call the embodied choice (EC) model. This is compared against two serial evidence-accumulation models, based on the well-known drift-diffusion model², and their respective performance is evaluated (see Lepora and Pezzulo 2015 for details).

² The drift-diffusion model aims to capture how a subject integrates (noisy) accumulating evidence, for multiple distinct options, in a forced choice task. The model assumes that evidence is integrated at various time steps, until some threshold is reached and a commitment is made to one of the options.

The first of these models is represented by a simple *serial process*, in which deliberation fully precedes a choice that commits the agent to the preparation and performance of the chosen action—much in the same way that the ‘classical sandwich model’ highlights (Hurley 1998). A *parallel model* develops this by connecting the decision process to action *preparation*. This speeds up the agent’s performance by anticipating what action will be most likely given the incoming sensory evidence. As evidence in support of one option increases, the agent can begin to make preparations for the respective action, before fully committing to it. Though the latter model gains a speed increase, it does so at the expense of accuracy. It could easily turn out that evidence that initially supports one option is overshadowed by later competing evidence, leading to inaccurate or clumsy actions.

To deal with this speed-versus-accuracy trade-off, Lepora and Pezzulo develop the EC model, which, in addition to the parallel feed-forward connection, has a feedback connection that allows action dynamics (e.g. current trajectory and kinematics) to influence the decision-making process. Whereas the previous models consider decisions to be independent of ongoing action (only allowing for influence from previous experiences), embodied choices consider ongoing action as an integral part of the decision-making process, with proprioceptive signals feeding into the deliberative process to provide information about the evolving biomechanical costs of associated actions.

This is important for real-world decisions. To illustrate, Lepora and Pezzulo (Lepora and Pezzulo 2015, pp. 4-5, emphasis added) give the following example of a lion that has begun tracking a gazelle, deciding to switch and track another:

[...] if the lion waits until its decision is complete, it risks missing an opportunity because one or both gazelles may run away. The lion faces a decision problem that is not stable but dynamic. In dynamic, real-world environments, costs and benefits cannot be completely specified in advance but are defined by various situated factors such as the relative distance between the lion and the gazelles, which change over time as a function of the geometry of the environment (e.g. a gazelle jumping over an obstacle can follow a new escape path) and the decision maker’s actions (e.g. if the lion approaches one gazelle the other can escape).

They continue:

[...] action dynamics in all their aspects (i.e. both their covert planning and their overt execution) have a backwards influence on the decision process by changing the prospects (the value and costs of the action alternatives). For example, when the lion starts tracking one of the gazelles, undoing that action can be too costly and thus the overall benefit of continuing to track the same gazelle increases. This produces a *commitment effect* to the initial choice that reflects both the situated nature of the choice and the cognitive effort required for changing mind at later stages of the decision.

A couple of comments are necessary. First, being receptive to ongoing action means the EC model can consider the evolving biomechanical costs that are salient to the current decision. Although the serial and parallel models can incorporate action costs as well, they must do so *a priori*, as there is no way for the ongoing action to feedback into the deliberative process. Critics may argue that part of the developmental process for any organism is learning about the body, and associated biomechanical costs, which are not going to change that drastically, given the limited number of states that the body can be in. Therefore, prior knowledge of biomechanical costs can be incorporated through learning. This is surely correct, but is also incomplete. As the gazelle example should highlight, biomechanical costs are also partly dependant on the evolving state of the environment, and where other agents are involved, will be difficult to precisely evaluate in advance. Second, commitment effects make it harder to change your mind once an action is performed, because the later sensory information must outweigh the initial commitment that arises from having started an action. Situated agents that are

receptive to subjective commitment effects may gain an important adaptive advantage, especially if the agent is able to learn about them for future interactions (see section 5).

As well as dovetailing nicely with the embodied decisions account, Lepora and Pezzulo found their EC model to perform better in terms of speed and accuracy than the alternative models. Initially, the models were evaluated in two simulation studies representing a two-alternative forced choice task (see [Lepora and Pezzulo 2015](#), for details), which on its own stands as an interesting proof-of-principle. However, they also compared their models with empirical evidence from human studies, and found that the EC model was a good fit with human behaviour.

Taken together, the aforementioned properties of embodied decisions stand in contrast to the cognitivist assumptions of traditional decision theory. To reiterate, the cognitivist perspective on decision-making is strictly separated from evidence accumulation in perceptual systems, and the control of action in motor systems. However, embodied decisions view deliberation as a continuous competitive process within sensorimotor circuits, modulated by relevant biases from cortical and sub-cortical regions as well as from ongoing action, and from which a distributed consensus emerges. It is hard to maintain the traditional functional separation of perception, cognition and action if we are to appreciate this process fully.

A number of issues remain. First, although there is mention of ‘continuously biasing influences’ in the embodied decisions research, there is little explicit mention of the role of affective signals in the aforementioned work. This is of vital importance, as an agent should have some way of determining which action opportunities it cares about most. Second, some important questions remain about whether the concept of embodied decisions can scale up and accommodate more explicit, goal-directed decision-making. Before addressing these issues directly, we will explore how predictive processing shares many of the same motivations as embodied decisions. By doing so, we hope to uncover where the two frameworks can offer mutual development.

3 Decision-Making in the Predictive Brain

Like embodied decision-making, an embodied account of predictive processing eschews the idea that perception is a passive accumulation of evidence with the purpose of reconstructing some detailed inner model of the world ([Clark in press](#)).³ Instead, perception is in the service of guiding actions that keep the organism within homeostatic bounds and maintain a stable grip upon its environment ([Clark 2015](#); [Friston et al. 2010](#)). In this manner, Clark sees the PP story as a contemporary expression of the *active vision* framework ([Ballard 1991](#); [Churchland et al. 1994](#)), and argues that many of the latter’s motivations are also present in an *action-oriented* account of predictive processing. He emphasises the following role for prediction error minimisation (PEM):

[...] it is the guidance of world-engaging action, not the production of ‘accurate’ internal representations, that is the real purpose of the prediction error minimizing routine itself. ([Clark 2016](#), p. 168)

This leads to a different relation between the key notions of *perceptual inference* and *active inference* from other proponents of the PP framework (e.g. [Hohwy 2013](#)). These terms refer to the two ways that

3 The exact manner in which PP should accommodate a notion of embodiment is still an open question, and is unlikely to be resolved without greater clarity regarding the notion of embodied cognition. One may worry that failing to resolve this issue would impede a satisfactory unification of the two ideas being considered here.

In response to this worry, we can offer two remarks. First, we echo the sentiments expressed by Shapiro ([Shapiro 2011](#)), who argues that it is acceptable to acknowledge the varied (sometimes nebulous) methodological practices that operate under the ‘embodied cognition’ banner, and that for the time being we should refer to embodied cognition as a research programme, rather than as a theory, to avoid the appearance of dogmatic unity. Second, the two accounts considered here both acknowledge the existence and explanatory importance of action-oriented representations, which is perhaps one of the most hotly debated areas in embodied cognition. Therefore, given this common ground, for the time being we believe it best to pursue the possible unification, while being careful to acknowledge any theoretical disagreements.

prediction-error can be minimised: either the system can update the parameters of its inner models, in order to generate new predictions about what is causing the incoming sensory data (perceptual inference), or it can keep its generative model fixed, and resample the world such that the incoming sensory data accords with the predictions (active inference).

Although both play an important role in PP, for Clark, the primary role of perceptual inference is to “prescribe action”, and as such, he states, that our percepts, “are not action-neutral ‘hypotheses’ about the world so much as ongoing attempts to parse the world in ways apt for the engagement of that world.” (Clark 2016, p. 124) This is a thoroughly action-oriented account, and importantly, it is this shift in emphasis that exposes a unity between Clark’s account of predictive processing and the insights of the ACH, which claims that neural processes represent action opportunities, rather than organism-independent, objective properties of the world (Clark 2016, p. 181). In addition, Clark views ‘active inference’ as a more-encompassing label for the combined mechanisms whereby the perceptual and motor systems cooperate in a dynamic and reciprocal manner to reduce prediction-error by exploiting the two strategies highlighted above.⁴ Active inference is accomplished using a combination of perceptual and motor systems rather than being confined to the latter that are traditionally associated with action. This view is also supported by recent neuroanatomical evidence that suggests a close relationship in the functional anatomy of the sensorimotor systems (Adams et al. 2013; Shipp et al. 2013). As Friston et al. argue:

The primary motor cortex is no more or less a motor cortical area than striate (visual) cortex. The only difference between the motor cortex and visual cortex is that one predicts retinotopic input while the other predicts proprioceptive input from the motor plant. (Friston et al. 2011, p. 138)

This supports one of the core claims of PP, which states that action is accounted for by a downwards cascade of predictive signals through motor cortex, which elicit motor activity in much the same way as predictions descend through perceptual hierarchies. In short, motor control is just more prediction, albeit about proprioceptive signals (see chapter 4 of Clark 2016). Rather than updating the generative model in response to error signals, these control-states predict subjunctive sensory trajectories that *would* ensue *were* the agent performing some desired action. It is this development of the active inference framework that allows PP to provide an account of choice behavior.

In PP, choices are made between competing higher-level predictions about expected sensory states. The formal basis for this perspective is based on the *free-energy principle* (Friston 2010). In (Friston et al. 2014) this is extended to account for decision-making in terms of hierarchical active inference. Friston describes choices as ‘beliefs about alternative policies’, where a policy is defined as a *control sequence* (i.e. sensory expectations associated with a sequence of descending proprioceptive predictions) that determines which action is selected next. Policies are selected under the prior belief that they minimise the prediction error between attainable and desired outcomes across multiple, hierarchically-nested levels. Importantly, these policies are selected on the basis of a belief in both their expected outcome, and also their expected precision (see section 3.2). Recently, Pezzulo and Cisek (Pezzulo et al. 2016) have argued that the ACH can be thought of in similarly hierarchical terms, and argue that this is further reason for adopting a control-theoretic, or action-oriented approach to cognition.

As applied to decision-making, this line of thought bears a close resemblance to work by Daniel Wolpert, who has demonstrated the close ties between motor control and decision-making (Wolpert and Landy 2012). One key difference between the views expressed by Wolpert and PP, however, is the latter’s rejection of the need for the separate, explicit representation of *cost functions*.

⁴ We have also argued elsewhere that it is misleading to simply equate *perceptual inference* with perception and *active inference* with action (Burr and Jones 2016). Instead, we take perception to be an active exploration of the environment, involving a continuous (and simultaneous) unfolding of *both* perceptual inference *and* active inference. Similarly, action involves both altering the environment by changing one’s bodily state, and monitoring the ongoing changes. In this manner, perception and action, understood in folk psychological terms, involve a combination of *both* perceptual *and* active inference at the level of underlying cognitive processing (see Vance 2017, for further discussion on these points).

In PP, cost functions are absorbed into the generative models harboured by the brain, and continuously updated by prediction error signals, becoming intertwined with the expectations of some policy. As these expectations will have been shaped by learning, there is already a *prior belief* about a policy's probability—a probability based on previous experience and captured by the extent to which it minimises prediction error through action (Friston et al. 2014). Some may worry that this view eliminates too much, or is too deflationary, and that the need for encoding some measure of the value associated with an outcome is necessary to explain why certain behaviors are preferred over others. Several responses can be offered to placate such worries.

Firstly, Clark explains that many working roboticists have already turned away from the explicit encoding of separate value/cost functions, arguing that they are too inflexible and biologically unrealistic due to their computational demands (Clark 2015). Instead they favour approaches that exploit the complex dynamics of embodied agents (such as the computational approach of (Lepora and Pezzulo 2015)), as they are computationally less demanding. These approaches acknowledge that the physiological constraints of an agent provide implicit means of understanding the value associated with dynamic action performance, without the need for positing additional abstract neural representations (see section 5.1).

Secondly, there are a number of debates in decision theory about whether the brain does in fact calculate value, with some arguing in favour of some abstract form of a neural 'common currency' (Levy and Glimcher 2012). However, as Vlaev et al. (Vlaev et al. 2011) argue, these views are beset with difficulties both from behavioural studies that explore contradictory, empirically-observed context effects (e.g. preference reversals and prospect relativity), as well as competing neurophysiological studies (see section 5.3). Vlaev et al. review a range of theories and models and provide the following positions to help capture these commitments:

- Value-first position: the brain computes the value of different options and simply picks the one with the highest value.
- Context-dependent value: the brain computes values, but the choice is heavily context-dependent on the set of available options.
- Comparison with value computation: the brain computes how much it values options, but only in relation to other values.
- Comparison-only: choice depends on comparisons without any computation of value.

Micro-debates exist within each of these positions. For example, is value represented on some ordinal, interval or ratio scale, and what objects are represented? Regardless of how these debates turn out, it should be clear that the *value-first* position is incompatible with the embodied decisions perspective. This is because value-first positions maintain that the value of an option is stable, and explicitly represented. We have already seen that the embodied decisions account is opposed to such a view, due to conflicting neurophysiological evidence. In addition, we have seen how PP eschews the explicit representation of value/cost functions altogether. However, it is unclear which of the alternative positions would best describe an embodied account of PP.

As a possible solution, den Ouden et al. (den Ouden et al. 2012), review the neurophysiological evidence relevant to an understanding of prediction errors, and argue that there is support for multiple kinds of prediction errors (PEs) in the brain: perceptual PEs, cognitive PEs, and motivational PEs. The first two types are referred to as *unsigned* PEs. These do not reflect the valence of any sensory input, but simply the surprise of its occurrence. The final kind of PEs, however, are known as *signed* PEs, for they reflect whether an outcome was better or worse than expected. They state:

Signed PEs play a central role in many computational models of reinforcement learning. These models describe how an agent learns the value of actions and stimuli in a complex environment,

and signed PEs that contain information about the direction in which the prediction was wrong, serve as a teaching signal that allows for updating of the value of the current action or stimulus. (den Ouden et al. 2012, p. 4)

Having access to multiple kinds of PEs, including those with affective significance, may provide the brain with the means to implicitly *compare* and *evaluate* which policy is most desirable based on prior learning. If value is determined indirectly through the comparison of multiple PEs, this would allow the agent to assess which of the myriad possible action opportunities is most salient given its current needs. The comparison could take the form of a distributed competition, in line with the proposal offered by Cisek, with no need for an abstract encoding of value that is generated downstream of sensorimotor processing (Cisek 2012).

In addition, adopting the suggestion of multiple PEs seems to frame PP as either an example of the ‘context-dependent value’ view or the ‘comparison with value computation’ view—depending on which additional mechanisms are posited to co-ordinate or integrate the options based on the type of PE considered. For example, although PP eschews talk of explicit cost functions, there is nevertheless a non-trivial sense in which the brain is comparing the expected values of the predictions that stand in place of the cost functions. Given the uncertainty regarding the precise implementational details of an exact architecture for PP (see Clark 2016, pp. 298-299, for a list of possible schemas), this could be a possibility. It is also one area where a synthesis between the work on embodied decisions and PP could be mutually beneficial—the former is presently developing novel computational methods that may help specify architectural details, whereas the latter provides a wider framework that unifies perception, cognition, action, and as we will see shortly emotion. However, there is another, possibly more radical approach that makes use of *precision-weighting* mechanisms, which may frame PP as an example of the ‘comparison-only’ view. This approach may be able to answer the underspecification challenge posed at the start of the paper, but doing so requires bringing the body more closely within the remit of PP.

4 Scaling Up

One challenge is particularly pressing given what we have so far discussed. The work on embodied decisions has focused primarily on exploring the neural mechanisms that underlie decision-making in simple, *visually-guided* motor tasks, such as grasping an object or pressing a button—so called habitual (or situated) decisions. Though this may be sufficient for explaining a wide variety of simple behaviors across a number of different species, humans (and some non-human animals) possess far more complex decision-making capacities. It is possible that the embodied decisions approach will be unable to account for the rich, and seemingly heterogeneous practices, that traditional decision theory tends to concern itself with.

For example, when you decide to buy a house, or choose where to go on holiday, it is not immediately obvious how a notion of embodied decisions could be of any use. Buying a house or going on holiday are both activities that require long-term planning. This implies the prolonged maintenance of a desired goal-state (e.g. buying a house), in order to coordinate and constrain relevant behaviours (e.g. acquiring a mortgage and communicating with solicitors). It is not immediately clear how the predictive brain handles the representation of distal goal-states by making solely embodied decisions of the kind hitherto discussed.

To respond to this challenge, we need to turn to a distinction often made within the decision-making literature between *deliberative* and *habitual* forms of decision-making. Competing *model-based* and *model-free* accounts are respectively put forward to try and capture the associated phenomena (Daw et al. 2011; Doll et al. 2012; Lee et al. 2014). In deliberative cases, where choice depends on the evaluation of various options, model-based methods use richly structured internal models (often rep-

representing conditional, probabilistic relations between states of the relevant domain) to simulate future action possibilities and their associated values—such accounts are increasingly studied in neuroeconomics (Glimcher and Fehr 2014). These general methods are flexible enough to apply to a wide range of circumstances. In contrast, habitual decisions rely on previously-learned “cached” or “heuristic” strategies to choose between actions and are frequently used in reinforcement learning. Although less flexible than model-based accounts, the benefit of model-free methods as Clark acknowledges, is that they nonetheless “implement “policies” that associate actions directly with rewards, and that typically exploit simple cues and regularities while nonetheless delivering fluent, often rapid, response” (Clark 2013b, p. 5).

Initially, the model-free strategy seems incompatible with PP due to the latter’s strict adherence to the existence of generative models. However, Clark argues that precision-weighting may explain how the brain flexibly switches between these two strategies on the basis of expected precision and accuracy (Clark 2013b; Clark 2016). With regards to the model-free, or heuristic strategy he claims:

[...] the use (when ecologically apt) of simple cues and quick-and-dirty heuristics is not just compatible with prediction-based probabilistic processing: it may also be actively controlled by it. (Clark 2013b, p. 8)

Here, Clark appeals to work in reinforcement learning (Daw et al. 2005; Gläscher et al. 2010) that shows how model-free strategies can implicitly learn (and embody) probabilities associated with certain action sequences or policies through trial and error, without the need to retain an explicit value or construct a detailed representation. These “cached” policies can then be redeployed at a later stage by the agent if they are estimated to be more reliable than the alternatives.

Recently, a number of studies have argued that the brain decides between these two modes by employing some form of arbitration mechanism (so called “neural controllers”) that predicts the respective reliability of the respective policies and chooses between them (Daw et al. 2005; Lee et al. 2014; Pezzulo et al. 2013). Clark believes these controllers could be accounted for in PP by appealing to *precision-weighting*, allowing them to be brought within the scope of the generative models (Clark 2013b; Clark 2016).

In PP, precision-weighting is considered to be a process by which the brain increases the gain on the prediction errors that are estimated to provide the most reliable sensory information, conditional on the higher-level prediction (Feldman and Friston 2010; Hohwy 2012). These precision-estimations work in close unison with higher-level predictions to provide context for the incoming prediction-error. For example, a high-level prediction carrying contextual information about the environment (i.e. at a high spatiotemporal scale) can also provide contextual information about which sensory inputs are most reliable (e.g. noisy environments mean auditory information is unreliable). It is claimed that the mechanisms behind this precision-weighting involve altering the post-synaptic gain on prediction-error units, and may also provide a way to reconcile the competing effects of signal suppression and signal enhancement (Clark 2016). If so, this would provide a more unified account for accommodating findings such as those mentioned earlier in the experiments of Cisek and Kalaska, where multiple choices are selected between on the basis of competing suppression and enhancement (Cisek and Kalaska 2010).

As well as providing a way of balancing the influence between top-down and bottom-up signals, it has also been argued that the mechanisms behind precision-weighting could provide a means of altering the brain’s *effective connectivity* (i.e. the influence that one region exerts over another) (see Clark 2013b). For example, (den Ouden et al. 2010) found evidence that striatal prediction errors play a modulatory role on the large-scale coupling between distinct visuomotor regions. Additional research has also explored the context-dependent, transient changes in patterns of cooperation and competition between control systems, which result from higher-level cognitive control (Cocchi et al.

2013). Some have argued that these transiently assembled networks, are formed in response to the task demands faced by the situated agent, and that key neuromodulatory mechanisms (e.g. volume transmission) may be responsible (Anderson 2014).

This additional role for precision-weighting allows PP to accommodate the ubiquitous effects of attentional modulation that Cisek and Pastor-Bernier discuss (see chapter 6 of Clark 2016). For example, (Friston et al. 2012) investigated the neuromodulatory role of the dopaminergic system, with a specific focus on decision-making and reinforcement learning. They argue that dopamine controls the precision of incoming sensory inputs, which engender action, by balancing the respective weight of top-down and bottom-up signals during active inference. This balancing means, crucially, that the predictions that drive action, also determine the context in which the movements are made and provide a way of balancing top-down expectations with bottom-up prediction errors to flexibly manage higher-level goals (based on urgency and saliency).

Given these more global influences of precision-weighting, some have argued that the distinction between model-free and model-based methods is too coarse-grained to be usefully applied at the neural level (Gershman and Daw 2012; Clark 2016). Instead, the distinction becomes a difference in degree rather than in kind. In line with an earlier suggestion regarding the brain's maintenance of a careful balance predictions and prediction-error, Clark (Clark 2016) suggests that the same strategy should be applied to the balancing between model-free and model-based methods. As such, the former would be associated with a greater degree of bottom-up processing (i.e. driven by sensory information and sensorimotor coupling), and the latter would be associated with more knowledge-driven, top-down processing:

The context-dependent balancing between these two sources of information, achieved by adjusting the precision-weighting of prediction error, then allows for whatever admixtures of strategy tasks and circumstances dictate. (Clark 2016, p. 253)

This is an interesting claim, and one that PP appears well-equipped to handle, but it is unlikely on its own to satisfy the challenge of underspecification introduced at the start of the paper. How does the brain select, from the wide range of action opportunities, the sequence that most effectively leads to the satisfaction of some distal (possibly abstract) goal representation?

A speculative proposal offered by (Pezzulo et al. 2016) argues that different types of policies can be distinguished according to whether they are associated with *extrinsic value* (i.e. the expected physical reward for completing the action) or *epistemic value* (i.e. the additional information gain or resolution of uncertainty). This distinction may be useful in explaining how offline forms of motor planning—resembling earlier theories of motor simulation (Grush 2004)—evolved progressively as elaborations on earlier sensorimotor control loops (Pezzulo 2012). These loops could then be utilised as a sort of *epistemic action*, simulating the expected precision of certain policies, and reducing the epistemic uncertainty associated with overt behaviour. In cases like this, there may be a payoff for considering actions with high epistemic value in order to ascertain whether there are other options that have not yet been considered (Friston et al. 2015). Such cases represent a sort of best-guess for the agent, based on prior knowledge of how similar situations have played out in the past. This suggestion, if valid, is likely to play an important explanatory role for more distal explanations of how the problem of underspecification is resolved in biological organisms. As PP and the ACH both claim that probabilistic neural representations of action-opportunities are in continuous competition—biased by sub-cortical and cortical influences—the above suggestion seems like a worthwhile area for future investigation.

However, an embodied account can appeal to further constraints that may help to provide a more complete answer to the underspecification challenge (Basso 2013). It is to these constraints that we now turn.

5 Constraining Effective Choice Behaviour

Consider the decision of whether to go for a run. The distal goal-state of ‘go for a run’ is satisfied once you begin your workout. However, there is a series of more fine-grained causal events that exists between the time when you purportedly “decide” to go for a run and the satisfaction condition of having gone for a run. We wish to argue that the decision to go running, should include the full series of fine-grained causal microstructures—beginning with the mental representation of the goal-state considered, and ending with the overt performance of the necessary behaviour.⁵ As such, the decision of whether to ‘go for a run’ is temporally extended over time, and as we will see, is partially constituted by events that extend beyond the brain and body.⁶ This reconception of decision-making from a diachronic perspective allows us to more effectively appreciate the dynamic nature of embodied decisions, and the continuous effects of active inference. However, it requires a number of claims to be explored and defended.

The first is that goals do not exist independently of their being enacted through an organism’s interactions with the world; that is, they have no independent objectivity (Gallese and Metzinger 2003). Only *goal representations* have a physical existence, realised by particular patterns of neural activity. Secondly, although we speak of *goal representations*, as we use the term, they differ from traditional notions of representation in a number of ways: a) they have no truth-conditions, only conditions for satisfaction that are directed towards the deployment of certain actions that minimise prediction error through active inference (Gallese and Metzinger 2003), and b) they are strictly grounded in facts about the agent’s embodiment—although possibly multimodal at some high-level of abstraction, they are not amodal in the sense used by the cognitivist (Burr and Jones 2016).

The motivation behind (a) follows from the truth of the first claim, defended by (Gallese and Metzinger 2003 p. 371), that “no such things as goals exist in the objective order of things”, therefore, “a goal representation cannot be true or false.”⁷ In PP, goal representations (in the form of higher-level predictions) are required by active inference, and thus have satisfaction (or fulfilment) conditions based on the imperative to minimise sensory prediction error. This leads to consideration of (b), and to the question of whether the existence of goal representations require more than can be provided by an embodied account of PP.

For example, the distal goal-state to go for a run appears to be abstract, despite being decomposable into more fine-grained sub-events (e.g. put on trainers; warm-up muscles; fill water bottle; lock door on leaving house; spend 20 minutes attempting to get your GPS (Global Positioning System) watch to detect your location). It is this abstract nature of the initial goal-state that leads to the underspecification challenge. This is because each of these multi-functional events can be considered independent of the specific goal-state—I may fill my water-bottle because I am thirsty and require a drink; I will lock my door whenever I leave my house (irrespective of whether I am going for a run). This fact regarding the multi-functionality of sub-events doesn’t appear to change even when the series of sub-events is so frequently performed that I rarely deviate from the order of performance. Alternatively, another decision (e.g. whether to buy a house) may be performed so infrequently, and contain a wide diversity of sub-events, that I will have very little idea of the order of events in advance.

With the aforementioned in mind, we will show how an embodied account of PP is able to appeal to a wide-range of explanatory factors, in order to demonstrate how the neural mechanisms that underlie decision-making constrain our choice behaviour (considered here from a diachronic perspective) in important, and perhaps adaptive ways.

5 For simplicity we put no requirements on how far you travel, or how fast you run in order for the statement, “I went for a run” to obtain. This vagueness is likely to be a characteristic of many choice behaviours, and we believe that a certain flexibility is necessary to account for differences in an individual’s own satisfaction conditions (e.g. less than 1km may not suffice for a professional athlete).

6 Obviously some decisions will be extended over shorter or longer period of times dependent on the framing of the decision (e.g. deciding between two sandwiches at a shop versus deciding on a new career path).

7 This point seems to fit with the subjunctive nature of higher-level predictions as employed in active inference accounts of motor control.

5.1 Physiological Constraints

In PP, predictions arise from generative models in the brain. These models are encoded as probability density functions, which are structured according to an increasing level of spatiotemporal scale. The predictions at the lowest levels correspond to the activity of sensory receptors encoding perturbations at small and fast spatiotemporal scales, whereas the higher-level models provide more general contextual information concerning larger and slower structures in the environment. The theoretical and empirical support for this picture has already been documented extensively in work by Karl Friston and colleagues (Friston et al. 2010), who argue that the formal similarities of their hierarchical models to the hierarchical structure of the motor system lends them biological plausibility (Kanai et al. 2015). Here, they argue, there will be a highly restricted set of possible parameters, which specify the range of possible actions, given the limited ways in which parts of the body could be configured. These parameters further restrict the set of actions, and may allow for automated or simple reflexive patterns resembling the sorts of habitual decisions we saw in the previous section. Far from being a hindrance to an agent, these restricted features can have adaptive value, allowing the agent to more easily detect and learn about the relevant features that emerge in the course of interacting with the world. This will in turn help with efficient action selection, as the specification of the relevant parameters can be reliably constrained by relevant factors of their embodiment.

For instance, as the eyes saccade from left to right, the visual scene will shift from right to left in a predictable manner, relative to the speed and direction of saccadic motion. An active perceiver can exploit regular relations between sensory input and motion of this kind in order to detect objective structural and causal features of the environment. These predictable relationships between bodily movement and sensory input are known as sensorimotor contingencies (SMCs) (O'Regan and Noe 2001), and are commonly discussed within the embodied cognition literature. We have elsewhere argued that this aspect of embodied cognition is implied by the PP framework (Burr and Jones 2016), and Seth has also proposed utilising SMCs to extend the PP framework to account for phenomena such as perceptual presence, and its absence in synaesthesia (Seth 2014).

Recent work by Cos et al. provides an interesting development to this idea (Cos et al. 2014). They argue that human subjects make a rapid prediction of biomechanical costs when selecting between actions. For instance, when deciding between actions that yield the same reward, humans show a preference to the action that requires the least effort, and are remarkably accurate at evaluating the effort of potential reaching actions as determined by the biomechanical properties of the arm. Cos et al. argue that their study (a reach decision task) supports the view that a prediction of the effort associated with respective movements is computed very quickly. Furthermore, measurements taken of cortico-spinal activity initially reflects a *competition between candidate actions*, which later change to reflect the processes of preparing to implement the winning action choice. Although there may be a possible disagreement concerning the exact manner in which cost functions are encoded and modelled, studies like this provide further reasons for taking the work of Lepora and Pezzulo (Lepora and Pezzulo 2015) seriously, due to the close connection with the aforementioned *commitment effects*, and the dynamic unfolding of decision-making.

Learning about the average biomechanical costs associated with performing certain actions could be a useful first-step in the formation of simple heuristics that stand in lieu of rational deliberation, and may also explain the presence of purportedly maladaptive decisions (e.g. sunk-cost fallacy). For example, some tasks undoubtedly require too much effort to properly deliberate over (e.g. choosing between a pair of socks), but a failure to properly identify these types of situations based on environmental markers, may lead to the misapplication of a strategy that is maladaptive in the current environment. For researchers, in cases where this strategy leads to undesired commitment effects, there may also be an opportunity to learn about the cognitive architecture of the agent in question. This is

because commitment effects reflect both biomechanical costs, as well as cognitive costs associated with changes of mind.

Being receptive to these changes in context is therefore of the utmost importance, as the value of many actions will vary contextually, dependent on factors such as fatigue, injury and environmental resistance (e.g. hill-climbing). However, rather than attempting to internalise all of the environmental variables, an alternative strategy is to simply allow the constraints of the body and environment to stand-in as a constituent part of the decision-making process. This is where dynamic, responsive feedback from the body, as input back into an ongoing decision is so important, and where work in situated cognition can provide constructive assistance. As Lepora and Pezzulo note:

In situated cognition theories, the current movement trajectory can be considered an external memory of the ongoing decision that both biases and facilitates the underlying choice computations by offloading them onto the environment. (Lepora and Pezzulo 2015, p. 16)

This work also connects with a further topic explored in the embodied decisions literature concerning decision-making in situations of increasing urgency.

5.2 Urgency

Accommodating urgency exposes another important connection between PP and embodied decisions. Given the level of urgency of an agent's higher-level goal states, the gain of incoming sensory information should be adjusted accordingly. Higher-level goals should therefore encode more abstract expectations regarding the optimal amount of time taken to deliberate in any given decision. Cisek and Pastor-Bernier point to the importance of an urgency signal in their work:

[...] in dynamically changing situations the brain is motivated to process sensory information quickly and to combine it with an urgency signal that gradually increases over time. We call this the 'urgency-gating model'. (Cisek and Pastor-Bernier 2014, p. 7)

When the urgency of a decision is low, only an option with strong evidence will win the probabilistic competition, and the agent may seek out alternative options (i.e. exploration). However, as the urgency to act increases, the competition between the options can increase, such that a small shift may be sufficient to alter the distribution. Cisek and Pastor-Bernier highlight a number of neuroimaging studies that support the existence of such an urgency signal, and argue that evidence accumulation may therefore not be the only cause of the build-up of neural activity seen during decision-making experiments.

By emphasising the importance of precision-weighting as a neuromodulatory mechanism for altering the brain's effective connectivity, PP may be able to further develop this line of thought in a more unified framework, which demonstrates the closely intertwined nature of perception, action, cognition and emotion, learning, and decision-making. This is because PEM is receptive to causes in the environment across a number of spatiotemporal scales. For example, perhaps some perturbing influence happens regularly at the order of milliseconds, but is also nested within a further perturbing influence that occurs on the timescale of minutes. The hierarchical structure of the brain is well-suited to accommodate these changes, but it is also well-suited to regulate additional factors such as the biomechanical costs involved with certain actions, which themselves may differ across spatiotemporal scales. Anyone who has done long-distance running and suffered as a result of inadequate pacing will attest to the importance of being receptive to the body's changing demands across extended timescales. This connects to a further important constraint, which has so far been overlooked—the importance of affective information originating in the body.

5.3 Interoceptive Inference

In PP, the predictions generated by the inner models of the brain do not merely attempt to anticipate the flow of sensory input from the outside world, but also the flow of interoceptive inputs (i.e. pertaining to endogenously produced stimuli, e.g. bodily organs). These inputs further constrain the set of viable actions in important ways. For example, deciding to quench one's thirst or sate one's hunger is often more important than allowing oneself to be distracted by alternative action opportunities. Being receptive to the current state of your body is fundamental to making adaptive decisions, as it allows us to determine which options have the greatest value *relative to our present needs*.

Tracking this type of sensory information requires incorporating *interoceptive information* into the PP framework. Seth has argued that Active Inference can be extended to accommodate interoception, and that key areas such as the anterior insular cortex (AIC) are well-suited to play a central role as both a comparator that registers top-down predictions against error signals, and as a source of anticipatory visceromotor control (i.e. the regulation of internal bodily states) (Seth 2013).

This is important for integrating decision-making within the PP framework, as it allows for a consistent understanding of the role that affective information (and possibly emotions) play in guiding our actions. Such a view, often attributed to the likes of William James, has seen a resurgence of interest, with many theories being proposed for how we are able to integrate emotions into the models of the mechanisms that underlie decision-making (e.g. Lerner et al. 2015; Phelps et al. 2014; Vuilleumier 2005). Some of these studies (Phelps et al. 2014) echo the sentiments of the earlier embodied decisions work, but go further in demonstrating how specific biasing inputs, such as affective information, play a fundamental modulatory role in the competitive process of action selection. An increasingly widespread claim, is that affective signals provide a basis for determining the salience of potential actions (Barrett and Bar 2009; den Ouden et al. 2012; Lindquist et al. 2012).

More specifically, the notion of *core affect*, which Lindquist et al. (Lindquist et al. 2012) define as “the mental representation of bodily sensations that are sometimes (but not always) experienced as feelings of hedonic pleasure and displeasure with some degree of arousal”, provides a way for an agent to know if some action is salient (i.e. good or bad for it). Importantly, this evaluation need not be considered as a separate step in a computational process. Barrett and Bar argue that activity in OFC is reflective of ongoing integration of sensory information from exteroceptive cues, with interoceptive information from the body (Barrett and Bar 2009). They claim that this supports the view that perceptual states are “intrinsically infused with affective value”, such that the affective significance (or salience) of an object (or action opportunity) is intertwined with its perception. This is one area where a synthesis between the work on embodied decisions and PP could be most beneficial, as the latter provides a more developed account of the importance of interoceptive information, with which to flesh out the notion of biasing inputs in distributed decision-making.

A further suggestion may come from research that postulates a more dynamic, action-oriented account of emotional episodes. For example, Lewis and Todd (Lewis and Todd 2005) view emotional episodes as the self-organising synchronization of neural structures, which help consolidate and coordinate neural activity throughout the nervous system. The emotional episode is posited to explain how an agent selectively attends to certain perceptual states, and is not perturbed by alternative goal obstructions (e.g. alternative action opportunities). An emotional episode thus acts as an important coordinating process for distributed neural activity, importantly including the sort of biases involved in the action selection required in embodied decision-making.⁸

⁸ It is important to note that an emotional episode is differentiated from “core affect”. An emotional episode is typically associated with an intentional object that is considered to elicit the agent's attention, and a corresponding cognitive appraisal of the stimulus. Core affect refers to a neurophysiological state that is consciously *accessible* as a simple primitive non-reflective feeling, and is an important component in emotional episodes (see Russell and Barrett 1999, for an analysis). However, even when an emotional episode is not present, there is always some felt core affect in the background, potentially ready to evolve into an intentional emotional episode.

The dynamic approach to emotions is important. It treats emotions as evolving states, rather than simply end-points, and can thus play a coordinating role that biases decision-making. As such, emotions are well-suited to perform long-term, action-guiding roles, as according to Lewis and Todd, they are directly concerned with “improving our relations with the world through some action or change of action.” Therefore, an emotional episode can direct attention away from obstructions that prevent the agent from obtaining some goal, and towards an associated intentional object. Lewis and Todd take this to be a fundamental factor that (partially) defines an emotional episode, and means that an emotion can assist an agent in overcoming and responding to goal obstruction, perhaps explaining why emotional episodes persist over time. This is important for understanding how distal goal-states, which require coordinated actions can be fulfilled.

These constraints help the predictive brain to recruit the relevant neural systems, best suited to respond to the current challenges it faces, based on current expectations defined by its ongoing activity. It also dovetails with the increasing attention being paid to the relation between emotions and decision-making (Phelps et al. 2014; Lerner et al. 2015), and emotions and cognition more generally (Pessoa 2013). Action opportunities are thus selectively attended to, based partly on the needs of the organism as determined by affective information, and may be modulated by ongoing dynamics that can be associated with action-guiding emotional episodes. Far from being mere limitations, emotional episodes can be seen as playing a coordinating role, the absence of which would likely result in unmanageable disorder. Although this picture is not sufficient to account for how all distal goal-states are obtained through co-ordinated action selection, it appears to be an important contributing factor—as suggested by the work on embodied decisions. The final suggestion, which could also bear on the social nature of emotional learning, is to look beyond the brain and the body, to acknowledge the constraints of the external environment.

5.4 Enculturation

An embodied account of PP embraces our cognitive limitations, and looks to explore how culture and the external world have been shaped to enable us to smoothly interact with the world. The need to move beyond a neurocentric perspective to an embodied (or situated) perspective is nicely expressed in a quote by anthropologist Clifford James Geertz :

Man’s nervous system does not merely enable him to acquire culture, it positively demands that he do so if it is going to function at all. Rather than culture acting only to supplement, develop, and extend organically based capacities logically and genetically prior to it, it would seem to be an ingredient to those capacities themselves. A cultureless human being would probably turn out to be not an intrinsically talented, though unfulfilled ape, but a wholly mindless and consequently unworkable monstrosity. (quoted in Lende and Downey 2012, pp. 67-68)

Of interest to this project is recent work by (Lende and Downey 2012) on *The Encultured Brain*: an exploration of recent interdisciplinary work in the fields of neuroscience and anthropology. The relevance of this interdisciplinary work (known as *neuroanthropology*) to an embodied account of PP is captured in the following:

A central principle of neuroanthropology is that it is a mistake to designate a single cause or to apportion credit for specialized skills (individual or species-wide) to one factor for what is actually a complex set of processes. (Lende and Downey 2012, p. 24)

Like embodied PP, neuroanthropology realises that exploring the brain alone (a form of methodological solipsism, cf. Fodor 1980) is insufficient to explain the myriad skilful interactions that define adaptive life, and instead requires turning to the notion of *enculturation*. Enculturation can be defined as the idea that certain cognitive processes emerge from the interaction of an organism situated in a

particular environmental, or socio-cultural niche. Neuroanthropology forcibly claims that many neurological capacities, such as language or skills, simply do not appear without the immersion of an organism within a particular culture (i.e. enculturation). In fact, Lende and Downey even state that “embodiment constitutes one of the broadest frontiers for future neuroanthropological exploration”, and that neuroanthropology is interested in “brains in the wild”, to appropriate a phrase from (Hutchins 1995). This requires understanding how our brain’s support skillful activity, and also how this activity has in turn re-wired our brains. Initial evidence points to: differences of neural structure and function between East Asian and Western cultures that may account for differences in notions of self (Park and Huang 2010); cross-cultural differences in subject’s ability to accurately judge relative and absolute size of objects (Chiao and Harada 2008), as well as evidence for differences in spatial representation of time (Boroditsky and Gaby 2010).

The notion of enculturation is often appealed to by those most accurately described as enactivists. For example, (De Jaegher and Di Paolo 2007) appeal to a notion of *participatory sense-making* to account for how social meaning can be generated and transformed through the interactions of a group of individuals collectively participating in collaborative activities. The notion of *participatory sense-making* is an extension of the enactivist notion of *sense-making*, which is the process that describes how an autopoietic system creates meaning through its lived experiences (Thompson 2007). For the enactivist, meaning does not exist independently of a system, but is defined by the selective interactions that are specific to (and defining of) certain phenotypes. These selective interactions could create a source of emergent meaning (Steels 2003), and also alter the functional and structural properties of our brains by retuning existing motor programs, which facilitate adaptive action selection and performance (Soliman and Glenberg 2014).

As our cognitive capacities have become increasingly advanced, we can begin to appreciate how our ability to shape our environment led to ways of simplifying it, in order to meet the requirement of minimising prediction error (i.e. making our environment more predictable). Hutchins (Hutchins 2014) offers a nice example of restructuring our material environment through interactive behaviours, which can be understood as a case of dimensionality reduction (an important component of PEM). For example, he offers the case of queueing as an instance of enabling a more straightforward perceptual experience. This is because the experience of a one-dimensional line, is more predictable than the experience of a two-dimensional crowd, and in turn the experience of a queue has a lower entropy (and thus a lower source of surprise) than the experience of a crowd. He states, “[t]his increase in predictability and structure is a property of the distributed system, not of any individual mind.” (Hutchins 2014, p. 40)

Of interest to this discussion, Fabry (Fabry 2017) considers the relation between enculturation and PP arguing in favour of a complementary approach. Her focus is on the acquisition and development of cognitive capacities, which she argues falls naturally out of the basic principles of PP (i.e. updating of model parameters via PEM). Her exploration of cognitive development considers the role of *learning driven plasticity*, which she defines as “the idea that the acquisition of a certain cognitive capacity is associated with the changes to the structural, functional, and effective connectivity of cortical areas.” (Fabry 2017, p. 4) Echoing the sentiments of the current discussion, she states this plasticity is constrained by anatomical properties of the agent, functional biases of the underlying neural circuits (Anderson 2014), and the environmental (and cognitive) niche. Despite the wide scope, Fabry argues that learning driven plasticity, which governs the acquisition of cognitive capacities (e.g. decision-making), is realised by ongoing prediction-error minimisation mechanisms. This *predictive acquisition of cognitive capacities* requires an understanding of the ontogenetic development of the situated agent, and also consideration of the more distal phylogenetic development. Both of these pursuits must consider the interrelated development of both the body of the organism⁹, and the environmental niche.

⁹ See Fabry’s (Fabry 2017, p. 7) suggestion regarding learning dependent bodily adaptability, which considers the developmental trajectory of skilled motor action alongside neural plasticity.

As Fabry argues, this is best approached from an embodied perspective. We agree with Fabry that, by connecting with relevant research in enculturation (and neuroanthropology), PP may find a complementary approach, which provides a way of answering the ultimate ‘why’ questions behind how the brain co-evolved alongside our body and external environment. This could supplement the focus on the ‘how’ questions that PP already seems well-suited to explain (cf. [Clark 2013a](#)).

Finally, we can reflect on how the social environment provides constraints on the sorts of action sequences we take when making decisions. Consider the decision between whether to go to Tokyo or Lima when booking a holiday online. Rather than seeing this as a synchronic choice, reported by clicking a button, we could instead view the choice behaviour as the first in a series of successive actions, which are subsequently constrained by virtue of the agent’s previously acquired knowledge. In this instance, that a significant financial cost (and corresponding feeling of regret) would be incurred were she not to go ahead with the subsequently implied actions (e.g. pack bags, head to airport etc.). In addition, in cases where social costs would be incurred (e.g. backing out of a verbally agreed arrangement), it is possible to view the decision as constrained by social commitment effects, akin to the sorts earlier proposed by Lepora and Pezzulo.

The various constraints considered above demonstrate a possible way in which the flexible precision-weighting of both habitual and deliberative decision mechanisms in the brain can be importantly constrained to lead to adaptive choice behaviour. Each of these constraints offers an intriguing explanatory step towards settling the underspecification challenge raised at the start of this section. By doing so, it pushes the explanatory scope of the PP framework outside of the neural mechanisms, while retaining an important foothold on how the neural mechanisms (as described by PP) flexibly interact with external factors to reduce uncertainty and maximise the probability of desired goal-states in as efficient a manner as possible. The body is an indispensable explanatory factor in this account: it allows us to uncover ways in which overt behaviour progressively tunes sensorimotor knowledge over developmental timescales; provides constraints on urgent decisions that will help to provide a more continuous evolutionary picture; offers affectively-significant information into the decision process that guides the organism to salient action selection; and it offers us myriad ways to mutually-interact with our sociocultural environment.

6 Conclusion

Herbert Simon famously claimed that choice behaviour should be understood as constrained by a pair of scissors, where the blades represent the limitations of the environment and the cognitive capacities of the agent in question ([Simon 1990](#)). Although his ecological approach to (bounded) rationality was cognitivist in nature, the core truth of his statement remains valid.

PP has much to offer for the second of the blades, and will likely place important theoretical constraints on the first if we acknowledge the lessons of neuroanthropology and enculturation and seek a complementary approach. Furthermore, by developing on the work in embodied decisions, we seek to gain a greater understanding of how decision-making is a dynamically-unfolding, situated process, rather than simply a process of deliberation and commitment. It seems as though developing the PP framework, in line with the myriad ways suggested, calls for an approach that is capable of explaining how the body and the world constrain the selection of action opportunities in the brain, in order to flexibly and efficiently support the adaptive choice behaviour of a situated agent.

References

- Adams, R. A., Shipp, S. & Friston, K. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure and Function*, 218 (3), 611–643.
- Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain*. Cambridge, MA: MIT Press.
- Baddeley, A. (1992). Working memory. *Science*, 255 (5044), 556–559.
- Ballard, D. H. (1991). Animate vision. *Artificial Intelligence*, 48 (1), 57–86.
- Barrett, L. F. & Bar, M. (2009). See it with feeling: Affective predictions during object perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364 (1521), 1325–1334.
- Basso, D. (2013). Planning, prospective memory, and decision-making: Three challenges for hierarchical predictive processing models. *Frontiers in Psychology*, 3, 1–2.
- Binmore, K. (2008). *Rational decisions*. Princeton: Princeton University Press.
- Boroditsky, L. & Gaby, A. (2010). Remembrances of times east: Absolute spatial representations of time in an Australian aboriginal community. *Psychological Science*, 21 (11), 1635–1639.
- Burr, C. & Jones, M. (2016). The body as laboratory: Prediction-error minimization, embodiment, and representation. *Philosophical Psychology*, 29 (4), 586–600.
- Chemero, A. (2003). An outline of a theory of affordances. *Ecological Psychology*, 15 (2), 181–195.
- (2011). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
- Chiao, J. & Harada, T. (2008). Cultural neuroscience of consciousness: From visual perception to self-awareness. *Journal of Consciousness Studies*, 15 (10–11), 58–69.
- Churchland, P. S., Ramachandran, V. S. & Sejnowski, T. J. (1994). A critique of pure vision. In C. Koch & J. L. Davis (Eds.) *Large-scale neuronal theories of the brain*. Cambridge: MIT Press.
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362 (1485), 1585–1599.
- (2012). Making decisions through a distributed consensus. *Current Opinion in Neurobiology*, 22 (6), 927–936.
- Cisek, P. & Kalaska, J. (2010). Neural mechanisms for interacting with a world full of action choices. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 33, 269–298.
- Cisek, P. & Pastor-Bernier, A. (2014). On the challenges and mechanisms of embodied decisions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369 (20130479), 1–14.
- Clark, A. (2013a). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36 (03), 181–204.
- (2013b). The many faces of precision. *Frontiers in Psychology*, 4 (270), 1–9.
- (2015). Embodied prediction. In T. K. Metzinger & J. M. Windt (Eds.) *Open Mind*. Frankfurt am Main: MIND Group. <https://dx.doi.org/10.15502/9783958570115>.
- (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.
- (in press). Busting out: Predictive brains, embodied minds, and the puzzle of the evidentiary veil. *Noûs*. <https://dx.doi.org/10.1111/nous.12140>.
- Cocchi, L., Zalesky, A., Fornito, A. & Mattingley, J. B. (2013). Dynamic cooperation and competition between brain systems during cognitive control. *Trends in Cognitive Sciences*, 17 (10), 493–501.
- Cos, I., Duque, J. & Cisek, P. (2014). Rapid prediction of biomechanical costs during action decisions. *Journal of Neurophysiology*, 112 (6), 1256–1266.
- Daw, N. D., Niv, Y. & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8 (12), 1704–1711.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69 (6), 1204–1215.
- De Jaegher, H. & Di Paolo, E. (2007). Participatory sense-making. *Phenomenology and the Cognitive Sciences*, 6 (4), 485–507.
- den Ouden, H. E. M., Daunizeau, J., Roiser, J., Friston, K. & Stephan, K. E. (2010). Striatal prediction error modulates cortical coupling. *Journal of Neuroscience*, 30 (9), 3210–3219.
- den Ouden, H. E. M., Kok, P. & de Lange, F. P. (2012). How prediction errors shape perception, attention, and motivation. *Frontiers in Psychology*, 3 (548), 1–12.
- Doll, B. B., Simon, D. A. & Daw, N. D. (2012). The ubiquity of model-based reinforcement learning. *Current Opinion in Neurobiology*, 22 (6), 1–7.
- Fabry, R. E. (2017). Predictive processing and cognitive development. In T. Metzinger & W. Wiese (Eds.) *Philoso-*

- phy and predictive processing*. Frankfurt am Main: MIND Group.
- Feldman, H. & Friston, K. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4 (215), 1–23.
- Fodor, J. A. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences*, 3 (1), 63–109.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11 (2), 127–138.
- Friston, K., Daunizeau, J., Kilner, J. & Kiebel, S. J. (2010). Action and behavior: A free-energy formulation. *Biological Cybernetics*, 102 (3), 227–260.
- Friston, K., Mattout, J. & Kilner, J. (2011). Action understanding and active inference. *Biol Cybern*, 104 (1–2), 137–160.
- Friston, K., Shiner, T., Fitzgerald, T., Galea, J. M., Adams, R., Brown, H., Dolan, R. J., Moran, R., Stephan, K. E. & Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS Computational Biology*, 8 (1), 1–20.
- Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T. & Dolan, R. J. (2014). The anatomy of choice: Dopamine and decision-making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369 (1655), 1–12.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T. & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 1–28. <https://dx.doi.org/10.1080/17588928.2015.1020053>.
- Gallese, V. & Metzinger, T. (2003). Motor ontology: The representational reality of goals, actions and selves. *Philosophical Psychology*, 16 (3), 365–388.
- Gershman, S. J. & Daw, N. D. (2012). Perception, action and utility: The tangled skein. In M. Rabinovich, K. Friston & P. Varona (Eds.) *Principles of brain dynamics: Global state interactions* (pp. 293–312). Cambridge, MA: MIT Press.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Glimcher, P. W. & Fehr, E. (2014). *Neuroeconomics*. Academic Press.
- Gläscher, J., Daw, N., Dayan, P. & O’Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, 66 (4), 585–595.
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27 (03), 377–396.
- Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, 3 (96), 1–14.
- (2013). *The predictive mind*. Oxford: Oxford University Press.
- (2016). The self-evidencing brain. *Noûs*, 50 (2), 259–285. <https://dx.doi.org/10.1111/nous.12062>.
- (forthcoming). The predictive processing hypothesis and 4E cognition. In A. Newen, L. Bruin & S. Gallagher (Eds.) *The Oxford handbook of 4E cognition*. New York: Oxford University Press.
- Hurley, S. (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- (2014). The cultural ecosystem of human cognition. *Philosophical Psychology*, 27 (1), 34–49.
- Kanai, R., Komura, Y., Shipp, S. & Friston, K. (2015). Cerebral hierarchies: Predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370 (1668), 20140169–20140169.
- Klaes, C., Westendorff, S., Chakrabarti, S. & Gail, A. (2011). Choosing goals, not rules: Deciding among rule-based action plans. *Neuron*, 70, 536–548.
- Lee, S. W., Shimojo, S. & O’Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, 81 (3), 687–699.
- Lende, D. H. & Downey, G. (2012). *The encultured brain*. Cambridge, MA: MIT Press.
- Lepora, N. F. & Pezzulo, G. (2015). Embodied choice: How action influences perceptual decision making. *PLoS Computational Biology*, 11 (4), 1–22. <https://dx.doi.org/10.1371/journal.pcbi.1004110>.
- Lerner, J. S., Li, Y., Valdesolo, P. & Kassam, K. S. (2015). Emotion and decision making. *Annual Review of Psychology*, 66 (1), 799–823.
- Levy, D. J. & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22 (6), 1027–1038.
- Lewis, M. D. & Todd, R. M. (2005). Getting emotional. *Journal of Consciousness Studies*, 12 (8–10), 210–235.
- Lindquist, K. A., Wager, T. D., Kober, H., Bliss-Moreau, E. & Barrett, L. F. (2012). The brain basis of emotion: A meta-analytic review. *Behavioral and Brain Sciences*, 35 (03), 121–143.
- O’Regan, J. K. & Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24 (05), 939–973.

- Padoa-Schioppa, C. (2011). Neurobiology of economic choice: A good-based model. *Annual Review of Neuroscience*, 34, 333–359.
- Park, D. C. & Huang, C. M. (2010). Culture wires the brain: A cognitive neuroscience perspective. *Perspectives on Psychological Science*, 5 (4), 391–400.
- Pastor-Bernier, A. & Cisek, P. (2011). Neural correlates of biased competition in premotor cortex. *The Journal of Neuroscience*, 31 (19), 7083–7088.
- Pessoa, L. (2013). *The cognitive-emotional brain: From interactions to integration*. Cambridge, MA: MIT Press.
- Pezzulo, G. (2012). An active inference view of cognitive control. *Frontiers in Psychology*, 3 (478), 1–2.
- Pezzulo, G., Rigoli, F. & Chersi, F. (2013). The mixed instrumental controller: Using value of information to combine habitual choice and mental simulation. *Frontiers in Psychology*, 4 (92), 1–15.
- Pezzulo, G., Cartoni, E., Rigoli, F., Pio-Lopez, L. & Friston, K. (2016). Active inference, epistemic value, and vicarious trial and error. *Learning & Memory*, 23 (7), 322–338.
- Phelps, E. A., Lempert, K. M. & Sokol-Hessner, P. (2014). Emotion and decision making: Multiple modulatory neural circuits. *Annual Review of Neuroscience*, 37 (1), 263–287.
- Russell, J. A. & Barrett, L. F. (1999). Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology*, 76 (5), 805–819.
- Savage, L. (1954/1972). *The foundations of statistics*. Dover: John Wiley and Sons.
- Sen, A. (1971). Choice functions and revealed preference. *Review of Economic Studies*, 38, 307–317.
- Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, 17 (11), 565–573.
- (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5 (2), 97–118.
- Shapiro, L. (2011). *Embodied cognition*. London: Routledge.
- Shipp, S., Adams, R. A. & Friston, K. (2013). Reflections on agranular architecture: predictive coding in the motor cortex. *Trends in Neurosciences*, 36 (12), 706–716.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41 (1), 1–20.
- Soliman, T. & Glenberg, A. M. (2014). The embodiment of culture. In L. Shapiro (Ed.) *The routledge handbook of embodied cognition* (pp. 207–219). London: Routledge.
- Steels, L. (2003). Intelligence with representation. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 361 (1811), 2381–2395.
- Thompson, E. (2007). *Mind in life*. Cambridge, MA: Harvard University Press.
- Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences*, 24 (5), 295–300.
- (2003). Visual attention: The where, what, how and why of saliency. *Current Opinion in Neurobiology*, 13 (4), 428–432.
- Vance, J. (2017). Predictive processing and the architecture of action. In T. Metzinger & W. Wiese (Eds.) *Philosophy and predictive processing*. Frankfurt am Main: MIND Group.
- Vlaev, I., Chater, N., Stewart, N. & Brown, G. D. A. (2011). Does the brain calculate value? *Trends in Cognitive Sciences*, 15 (11), 546–554.
- Vuilleumier, P. (2005). How brains beware: Neural mechanisms of emotional attention. *Trends in Cognitive Sciences*, 9 (12), 585–594.
- Wolpert, D. M. & Landy, M. S. (2012). Motor control is decision-making. *Current Opinion in Neurobiology*, 22 (6), 1–8.