Predictive Minds Can Be Humean Minds

Frederik T. Junker, Jelle Bruineberg and Thor Grünbaum

The predictive processing literature contains at least two different versions of the framework with different theoretical resources at their disposal. One version appeals to so-called optimistic priors to explain agents' motivation to act (call this optimistic predictive processing). A more recent version appeals to expected free energy minimization to explain how agents can decide between different action policies (call this preference predictive processing). The difference between the two versions has not been properly appreciated, and they are not sufficiently separated in the literature. They constitute two different theories with strikingly different accounts of motivation and action. By reducing all desire-like constructs to belief-like constructs, optimistic predictive processing entails a substantial revision of standard accounts of motivation and action in philosophy and cognitive science. By contrast, preference predictive processing introduces desire-like constructs that play Humean motivational roles in the explanation of action. In this Humean iteration, predictive processing resembles other prominent computational frameworks implementing a distinction between beliefs and desires, such as reinforcement learning and Bayesian decision theory. Ultimately, predictive processing faces a dilemma between parsimony of mental constructs and completeness of its explanations of agency and the mind.

ORCID: Junker, orcid.org/0000-0001-9785-6793; Bruineberg, orcid.org/0000-0001-5675-3376; Grünbaum, orcid.org/0000-0002-6464-8098

1. Introduction

Over the last decade, predictive processing (PP) theories have enjoyed rapidly growing interest in both philosophy and cognitive science. According to these theories, the agent navigates its environment by engaging in a form of model-based inference aimed at minimizing prediction error. PP's two main selling points are its universality, which is its potential to provide a unified and mechanistically plausible theory of all neurocognitive phenomena, and its parsimony, which is its potential to explain all neurocognitive and behavioural phenomena in terms of the same fundamental process: precision-weighted prediction error minimization, or Bayesian inference (Friston [2009], [2010]).

A central appeal of PP to some theorists is its radical conception of agency. Many contemporary philosophical theories of agency are Humean in nature: they presuppose a basic distinction between beliefs and desires. Beliefs tell you how the world is, while desires motivate actions. If PP delivers on its promise, then it seems that an account of agency can be given without reference to desires: all that is needed is a web of probabilistic beliefs that guides perception and action. In other words, PP is anti-Humean.

The anti-Humean commitments of PP, combined with its universalist aspirations, have been the source of considerable debate recently (Klein [2018], [2020]; Clark [2020]; Sun and Firestone [2020]; Van de Cruys et al. [2020]; Yon et al. [2020]). Critics have pointed out that because of its anti-Humeanism, PP cannot be a universal account of the mind: at some point you need desires to explain motivation and action. The main intuition pump to motivate this challenge is the dark room problem: why does a prediction error minimizing agent not simply seek out highly predictable environments, such as a dark room? In such a highly predictable environment, the PP agent would continuously predict pitch darkness that would always match the sensory input perfectly (Friston et al. [2012]; Klein [2018]; Sun and Firestone [2020]). Clearly, we are not such agents. The challenge for PP is to explain why we are not dark-room-seeking creatures and why we are instead motivated to go out and explore the world. Such motivation is standardly understood in terms of the agent having mental states, such as desires, with a world-to-mind direction of fit.

Over the years, PP theorists have repeatedly responded that PP has the theoretical resources to handle the dark room problem, as well as other challenges that might seem to require invoking distinct desires. Traditional responses attribute stubborn optimistic prior expectations to the agent that it occupies states that satisfy its bodily needs and curiosity about novel information (Bruineberg et al. [2018a]; Yon et al. [2019]; Van de Cruys et al. [2020]), while more recent responses appeal to more sophisticated models of policy selection via expected free energy minimization (Clark [2020]; Seth et al. [2020]). The implicit assumption among both proponents and critics is that PP remains anti-Humean: if PP turns out to be true, we would have succeeded in providing a universal account of agency exclusively in terms of probabilistic beliefs and expectations.

In this paper, we take issue with the common assumption that all versions of PP are incompatible with Humeanism and standard philosophical theories of motivation and action more generally. We argue that PP has recently developed a Humean branch. We pinpoint the origin of this branch to the introduction of models that make use of so-called expected free energy minimization (Friston et al. [2015]). Expected free energy models involve both state-estimation, the estimation of the most likely state of the environment given current sensory input, and policy selection, which is the selection of a policy that is expected to lead to preferred outcomes. We argue that an expected free energy minimizing creature can believe it is in one state while desiring to be in another. Hence, this is a Humean creature.

We identify two distinct branches in the PP literature, the difference between which has not been properly appreciated. According to one branch, which we will refer to as optimistic PP, agents are equipped with optimistic priors that make them predict that they will observe outcomes that are favourable to them. Optimistic predictions and prediction error minimization drive actions towards such outcomes. According to the other, which we will refer to as preference PP, agents score action policies on how well they minimize expected free energy and select the ones that strike the best balance between reducing uncertainty and bringing about preferred outcomes. The theories entail strikingly different accounts of motivation and action. We will conclude that while optimistic PP is anti-Humean and limited in explanatory scope, preference PP has wider explanatory scope, but also comes with straightforwardly Humean commitments, specifically, the acceptance of desire-like constructs. Consequently, the PP theorist needs to choose between parsimony (there are only belief-like states) and universality (explanation extends to all aspects of agency and mental function).

To make the discussion more approachable for non-experts in the PP literature or action theory, we have laid out our argument as follows. In section 2, we present Humeanism, as well as the main challenges for anti-Humean PP accounts of motivation and action. In section 3, we introduce the distinction between optimistic PP and preference PP. In section 4, we show that preference PP is a Humean theory. In section 5, we articulate their respective implications for a theory of action. If optimistic PP and preference PP entail different commitments with respect to the Humean nature of motivation, they imply different theories of action. We argue that while preference PP is compatible with certain versions of standard accounts of action, optimistic PP entails a radical revision. In section 6, we discuss the assertion that adopting the free energy principle (FEP) eliminates all distinctly motivational constructs from our ontology. That is, one might object that preference PP is, in fact, firmly anti-Humean. We argue that the FEP entails no such elimination. In the end, PP is confronted with a choice between universality and parsimony. In its ambition to explain all aspects of agency and mental function, PP has had to invoke desire-like constructs playing Humean motivational roles in the explanation of action.

2. Predictive Processing and Motivation

According to the Humean theory of motivation (Davidson [1963]; Smith [1987]), beliefs and desires are distinct types of mental states. Desires have motivational force, whereas beliefs have none. We follow Davidson ([1963]) in understanding desires as standing for the broader category of pro-attitudes. Pro-attitudes in the literature on practical reasoning, action, and ethics are attitudes in favour of something (for example, approval, admiration, liking, preference, and esteem) and include evaluative judgements that an action has some positive characteristic (for example, being desirable, reasonable, admirable, or dutiful). The notion of pro-attitudes thereby goes beyond a simple form of Humeanism restricted to primitive motivational states, such as drives or urges. This becomes important later when discussing the notion of preferred outcomes found in preference PP models, which could include evaluative judgements that play a motivational role and cannot be reduced to a purely doxastic register. Though used interchangeably, we will primarily use the more common term 'desires', except where the broader connotations of pro-attitudes become important.

According to the Humean, reason on its own is never sufficient to motivate agents to act. This idea is often understood as the claim that beliefs on their own are motivationally inert. Only with the addition of some desire will the agent be motivated to act. This difference between motivationally inert beliefs and motivationally active desires is often captured by defining desires in terms of dispositions to act. What we call the Humean theory of motivation is not committed to any particular theory of desire (Schroeder [2004]) or any strict form of motivational externalism (Williams [1979]). The Humean theory is simply committed to the claims that explanation of instrumental action requires both beliefs and desires, that desires are primitive or irreducible to beliefs, and that only desires have motivational force (which can be spelled out in terms of dispositions to act or in some other way).

One common way of describing the difference between beliefs and desires is in terms of a difference in direction of fit (Searle [1983]). Beliefs have a mind-to-world direction of fit, such that in case of a mismatch, they ought to be revised to fit the world. Desires, on the other hand, have a world-to-mind direction of fit, such that in case of a mismatch, the world should be changed to fit the desire. Given Humeanism, if there is a mismatch between the desired situation and the state of the world, the agent should be motivated to act to change the state of the world. Anti-Humeans, on the other hand, reject the existence of distinct desires and their unique role in motivating action. It is often claimed that PP theories are anti-Humean in this sense.

2.1. Humean challenges to predictive processing

According to PP, in all its guises, the agent navigates its environment by engaging in a form of model-based inference (Friston [2009], [2010]; Clark [2013], [2016]; Hohwy [2013], [2016]; Parr et al. [2022]). The brain generates a stream of top-down predictions about what sensory signals it expects to receive given its current best model of the world and the situation in which the agent finds herself. The predicted input is compared to the input the agent actually receives. This comparison leads to prediction errors, which are used to update the generative model or to change the sensory input by acting on the world. The combined process is sometimes called active inference. Since the theory's only primitive is prediction, many have pointed out that it cannot sustain a belief–desire distinction. The lack of distinct desires makes PP anti-Humean.

The fact that an anti-Humean theory of motivation denies the existence of distinct motivational, action-disposing mental states does not entail a denial of the distinction between directions of fit. The anti-Humean could still accept that mental states are sometimes made to fit the world, and sometimes the world is made to fit our mental states. Anti-Humeans often argue that beliefs can have motivational power, which is to say, they can have a world-to-mind direction of fit. Thus, not only do anti-Humeans need to provide reasons for denying the Humean distinction between beliefs and desires, but they also need to offer an alternative explanation of how agents resolve various types of conflict between mental states and the world without resorting to a primitive notion of desire or pro-attitude. Several Humean objections to PP have highlighted this explanatory challenge.

First, the need for action-disposing mental states with a distinctly world-to-mind direction of fit seems to be driving the dark room problem: without a mental state that disposes the agent to change the world and leave the dark room, how can the PP theorist explain the agent's ability to leave the dark room (Klein [2018]; Sun and Firestone [2020])? This is one way to motivate the Humean challenge to PP. Without distinct belief-like and desire-like states, core aspects of cognition and behaviour seem to be left unexplained (Yon et al. [2020]).

Second, predictions alone seem to be motivationally inert, and so PP does not seem to have the necessary primitives to explain what motivates action, or how and why we select the actions that we do (Klein [2018]). Frequently, agents must choose between multiple actions and rank them against each other. If the agent does not have distinct motivational states representing the value or reward of different actions, such selection problems quickly become intractable. Representing actions in terms of both probabilities and value allows for simpler comparison than if each action must be specified purely in terms of long series of conditional predictions ('when conditions C_1, and C_2, and C_3, … obtain, I will perform actions A_1, or A_2, or A_3 …').

Third, other computational frameworks, including reinforcement learning and Bayesian decision theory, are consistent with Humeanism. In these frameworks, value signals are always necessary for action, and probabilities and values are represented and computed independently. The empirical success of these frameworks provides some empirical support for Humeanism (Colombo [2017]).

Finally, a fourth challenge to anti-Humean PP has recently been presented by Klein ([2020]). According to Lewis ([1988], [1996]), desires are contingent. We might have uncommon desires or lack common ones. There is no necessary relationship between desire and belief. As Lewis ([1996], p. 304) puts it: 'Any values can go with any credences'. Building on this insight from Lewis, Klein ([2020]) has argued that PP fails to simultaneously explain action and value learning. To explain action under PP, the predictions driving actions must be effectively unrevisable to ensure that the ensuing prediction error can only be minimized through action instead of by simply revising the predictions.[1] However, sometimes action-guiding predictions ought to be revised when the agent learns that an action is no longer desirable. For instance, when the agent learns that the water is contaminated, she ought to revise the prediction that 'when I am thirsty, I drink water'. In short, to explain action within PP, the predictions driving action need to be unrevisable—but to explain value learning, they need to be revisable. This problem could be resolved by adding a rule that allows the agent to update the prediction when the water is observed to be contaminated.

However, to ensure that the prediction remains unrevisable in normal circumstances (where the water is not contaminated), the updating rule cannot be Bayesian updating. If the prediction were generally susceptible to Bayesian updating, the agent could update it, even in normal circumstances. For instance, the agent could take prolonged periods of thirst without drinking as evidence against the prediction that 'when I am thirsty, I drink water'. Beliefs, therefore, need to include lots of conditionals so that the updating rule only applies in the right circumstances (for example, 'unless the water is contaminated, when I am thirsty, I will drink water'). But such conditionals can get extremely complex, even for relatively simple creatures. They require specifying a virtually infinite number of conditions for when beliefs should and should

[1] Unrevisability can be achieved by taking the prediction to be so precise that no amount of prediction error will suffice to revise it. For an account of motivation in terms of unrevisable predictions, see (Miller Tate [2021]).

not be updated, for every scenario the organism might find itself in. It is highly implausible that our beliefs have this kind of complexity. Moreover, it rests on the dubious assumption that we come pre-wired with complex expectations fit for nearly every possible scenario, instead of adapting to environmental changes by learning about changes in value. Klein ([2020]) points out that Humean theories have an easier time explaining such cases. When we learn that the water is contaminated, we can leave other beliefs untouched and simply update our desires (for example, 'store-bought beverages are now strongly preferable to tap water'). Beliefs can thus remain responsive to the evidence via Bayesian updating, while desire updates ensure that we can still adapt and pursue the best course of action when circumstances change.

2.2. The revisionist response

Some proponents of PP have argued that PP can in fact explain the phenomena targeted by Humean critics without the need to posit distinct motivational states. Clark ([2020]) argues that motivational states can be cast as counterfactual predictions, the content of which is what we would observe if we acted in a certain way. We counterfactually predict that we are already in the desired state. This initially gives rise to prediction error, which is minimized by bringing the agent into the desired state. According to Clark ([2020], p. 12), PP should treat desires as 'varying forms and time-scales of prediction', the motivational force of which is dictated by the relative precision-weighting of those predictions. These counterfactual predictions simultaneously have belief-like features (they predict what will occur as a consequence of the action) and desire-like features (they are poised to bring about the predicted consequences of acting). This interpretation of PP essentially reduces all mental state types to a single primitive: precision-weighted predictions.

To summarize, critics have raised a number of challenges to PP based on its anti-Humean commitments, while its proponents have argued that PP can meet these challenges by revising our Humean intuitions. What matters for our purposes is the shared assumption that PP is inescapably anti-Humean. The assumption that PP does not allow for distinct types of mental states is not just held by philosophers working on PP, but is repeated in the more technical literature (Friston et al. [2009]; Friston et al. [2012]; Friston et al. [2015]; Friston [2019]; Parr et al. [2022]). Specifically, it is sometimes claimed that PP entails a kind of desert landscape ontology, where desires, goals, reward signals, and the like do not exist and are replaced by more parsimonious models of purely prediction-driven embodied exchanges of creatures with their environment (Friston [2019]).

3. Predictive Processing, the Devil, and the Details

Thus far, we have discussed PP in broad outline as a framework according to which everything is precision-weighted prediction error minimization. We would now like to argue that the current discussions have failed to recognize that PP architectures come in different guises, which account for motivation and action in different ways. To make this more explicit, we introduce a distinction between two theories: optimistic PP and preference PP.

3.1. Optimistic predictive processing

PP accounts vary in what they take the main purpose of the generative model to be. Initially, PP accounts were seen as a continuation of a Helmholtzian view of perception (Friston et al. [2012]; Clark [2013]; Hohwy [2013]). On such an account, the main purpose of PP is to reconstruct the hidden state of the world based on proximal sensory input. The content of perception is then determined by the set of predictions that manage to explain the sensory signals, or, equivalently, that manage to explain away the prediction errors. Following the Helmholtzian approach, action is seen as a kind of experiment that disambiguates between competing hypotheses and increases the evidence for one's current hypothesis (Friston et al. [2012]). As illustrated by the dark room problem and the other Humean challenges discussed in the previous section, it is difficult to see how a purely prediction-driven version of PP can explain all aspects of agency and deliver a unified theory of the mind.

One way to respond is to argue that agents are endowed with so-called optimistic priors, which dispose the agent to predict favourable outcomes, consistent with having their bodily needs met (for example, a full stomach, stable blood-glucose and hydration levels, and a body temperature of around 37°C). These priors are 'optimistic' in that the agent expects to observe outcomes that are 'good' for the agent, in the sense that they are compatible with its continued existence. They are also sometimes referred to as 'stubborn predictions' (Yon et al. [2019]), as they resist revision and can only be satisfied by making the agent's observations conform to its expectations. There is a second sense in which the agent's priors are 'optimistic': the agent expects a beneficial environment in which all its bodily needs can be met. Typical environments are not like this; they might be cold or lack food and water. To act adaptively, the stubborn agent needs to keep believing that it will find beneficial environments. Hence, adaptive action involves what Wiese ([2017]) calls 'systematic misrepresentations of the environment': the agent's expectations are (and need to be) systematically skewed towards the kinds of environments in which it thrives.

Some simulation studies of PP principles, where agents learn the structure of an environment, simply assume that the environment is beneficial. In Friston et al. ([2009)]), for example, the agent is first shown the correct sequence of actions without being able to intervene. After having learned the optimal solution in a controlled environment, the agent is then endowed with the capacity to act, and to make its observations congruent with the previously learned sequence of observations. The idea is that an agent equipped with optimistic priors will not discover the causal regularities in its environment, and update its model accordingly, but will instead make the environment conform to its expectations (Bruineberg et al. [2018a]; Yon et al. [2019]). To get out of the dark room, a PP agent needs to stubbornly predict that the world is different from how it is currently observed to be. If the agent is equipped with optimistic and stubborn expectations of a full stomach, then, as the agent in the dark room grows hungry, this leads to prediction errors that can only be minimized by leaving the room to eat. Thus, optimistic priors serve as self-fulfilling prophecies that compel the agent into action, even if this means facing less predictable environments. Since the defining construct of this version of PP is optimistic priors, we will refer to it as optimistic PP.[2]

Optimistic PP runs into the problems raised by Humean critiques. An account that needs priors to be both unrevisable and systematically skewed to account for adaptive action will have trouble providing an empirically adequate account of value learning: to learn values, the agent's priors need to be revisable. There is, however, a different version of PP on the market with a different architecture.

3.2. Preference predictive processing

To our knowledge, Friston et al. ([2015]) provides the first articulation of a PP theory that involves the minimization of prediction error in the future; that is, minimization of expected

[2] In some discussions of active inference, the explanation of optimistic priors is delegated to the FEP. The FEP holds that any system that maintains its organization over time will engage (or appear to engage) in a form of model-based inference in which the generative model embodies the optimal state of being for that system. We discuss the implications of the FEP in section 6.

free energy. At first glance, the introduction of expected free energy involves more of the same: 'Our basic approach is to cast optimal behavior in terms of inference, where actions are selected from posterior beliefs about behavior. This allows one to frame goals and preferences in terms of prior beliefs, such that goals are subsequently fulfilled by action' (Friston et al. [2015], p. 188). The authors seem to argue for an anti-Humean position: pro-attitudes, including goals and preferences, reduce to doxastic states, particularly prior beliefs. However, the devil is in the details. To see this, let us unpack the commitments of expected free energy minimization.

An agent has a finite number of policies or strategies available. To rank the policies, the expected free energy of each policy is evaluated. Heuristically, it amounts to the evaluation of the following counterfactual: 'What is the free energy I expect to receive if I were to pursue this policy?'. Calculating the result for each policy gives a policy-specific expected free energy. The probability of pursuing a policy is then proportional to the relative expected free energy of the policies: probability of policy $\propto$ expected free energy of policy. In other words: 'I will pursue those policies most often that I expect will minimize free energy'. But what exactly is expected free energy? One way to decompose expected free energy is as follows: expected free energy = expected ambiguity + risk.

The expected ambiguity term roughly captures: 'How much uncertainty will be reduced by pursuing this policy?'. A policy that brings the agent to a location where the agent expects it can gain new information (which would reduce uncertainty) will have lower expected ambiguity than a policy that brings the agent to a location where the agent does not expect to learn anything new. The risk term roughly captures: 'How close will following a particular policy bring me to a preferred outcome?'. Here, a lot of the explanatory work is done by the notion of preferred outcomes. So, what is a preferred outcome? In their introduction of expected free energy minimization, Friston et al. ([2015], p. 188) define preferred outcomes as follows: 'In active inference, constructs like reward, utility, epistemic value, etc. are described in terms of prior beliefs or preferences. In other words, preferred outcomes are simply outcomes one expects, a priori, to be realized through behavior (e.g., arriving at one's destination or maintaining physiological states within some homoeostatic range)'.

A preferred outcome is thus an outcome the agent expects given the kind of agent it is. The agent is set up to bring about expected outcomes by selecting policies that it expects will lead to those outcomes. If the agent's preferred outcome is 'tasting coffee', then a policy that involves pouring oneself a cup of coffee will involve less risk than a policy that doesn't. Preferred outcomes are defined in terms of a probability distribution over observations some time into the future. For this reason, they are also frequently referred to as 'preferred observations'. The two terms are used interchangeably in the literature and refer to the same formal construct: a probability distribution over observations at some time point in the future. For simplicity, we will stick to the term 'preferred outcomes'.[3]

If you think preferred outcomes sound suspiciously Humean, you would be correct. In introducing the notion of preferred outcomes, Friston et al. ([2015]) contrast it with pro-attitudinal constructs like reward and utility, but also likens it to pro-attitudes like preferences. We will return to this question shortly, but for now, let us point out some advantages of the version of PP that operates with expected free energy minimization.

First, minimizing expected free energy involves selecting policies that are expected to lead to new information and preferred outcomes. Another way to put this is that an agent trying to minimize expected free energy is trying to strike an optimal balance between explorative behaviours (reducing expected ambiguity) and exploitative behaviours (reducing risk). In the absence of ambiguity, the agent will simply select the policy that leads to preferred outcomes. In the absence of preferred outcomes, the agent will select the policy that reduces most uncertainty. Taken together, the minimization of expected free energy provides a relatively simple account of action selection in uncertain environments.

A second major advantage of these models is that they add a capacity for planning and decision-making. Expected free energy allows the agent to score different policies about how to act in the future. It allows the agent to consider possible future observations, as well as how possible future observations are conditioned on the policies that the agent could pursue. It adds an internal loop that considers possible future observations and evaluates them relative to preferred outcomes. The inferred policies that strike the optimal balance between bringing the agent towards preferred outcomes and reducing uncertainty will be considered most probable, and, therefore, be selected for action. Since the defining construct of this version of PP is preferred outcomes (at least in the context of explaining motivation), we will refer to it as preference PP.

[3] Other terms used include 'preferences', 'preferred states', and 'preferred sensations'. We take it that these are all used interchangeably.

4. Preference Predictive Processing: Going Humean

Let us return to our main question: did PP grow a Humean branch with preference PP? Consider a chess player who selects a move because its consequences are close to the kinds of consequences the player would like to see from a move (that is, seeing her opponent's position crumble rather than her own). The comparison of expected consequences of following an action with preferred consequences is an indispensable element of policy selection using expected free energy minimization. If there is a Humean, desire-like, pro-attitudinal element in preference PP, it is to be found in preferred outcomes.

4.1. Preferred outcomes

How are we to understand the notion of preferred outcomes? Can desires be replaced by beliefs about observations? Opinions on this seem to be mixed. In a recent treatment, Parr et al. ([2022] p. 53) offer the following proposal:

> […] using the notion of expected free energy amounts to endowing the agent with an implicit prior belief that it will realize its preferences. Hence, *the agent's preference for a course of action becomes simply a belief about what it expects to do, and to encounter, in the future*—or a belief about future trajectories of states that it will visit. This *replaces the notion of value with the notion of (prior) belief*. This is an apparently strange move, if one has a background in reinforcement learning (where value and belief are separated) or Bayesian statistics (where belief does not entail any value). (emphasis added)

These remarks suggest that talk of 'preference' and 'value' is somewhat deceptive or derivative. If preferences are fundamentally just beliefs about what the agent expects to do, then preference PP does not seem able to draw a genuine distinction between beliefs and desires, in anything but name. If so, preference PP seems, like optimistic PP, to be firmly anti-Humean.

However, one should be careful not to take such discourse at face value. Generally speaking, the word 'belief' in PP does not stand for any standard folk-psychological concept. Instead, a belief, in the technical sense employed in PP and Bayesian inference, is a probability distribution over a set of states. Importantly, the Bayesian notion of belief does not entail the mind-to-world direction of fit.[4] Indeed, this could not be the case, because the direction of fit between

[4] This is especially true of Bayesian beliefs over policies, which depend upon preferences over counterfactual outcomes.

the world and Bayesian beliefs can be decided by the relative precision of predictions and observations. This means that whether a Bayesian belief has a mind-to-world or a world-to-mind direction of fit depends on the decision architecture in which it is embedded.

So, what is the decision architecture in which preferred outcomes are embedded? Here it is worth taking a closer look at the details of expected free energy minimization. The canonical implementations of expected free energy minimization make use of partially observed Markov decision processes (Friston et al. [2015]; Friston et al. [2017]). Partially observed Markov decision processes make a number of assumptions about the conditional dependencies involved in the decision process. Most notably, they assume that observations at time t (o_t) are only dependent on the current hidden state (s_t), and that the probability of a hidden state s_{t+1} is dependent only on the previous hidden state s_t and the policy $\pi(t)$. By exploring its environment, an agent can learn the conditional dependencies of its environment: 'Given that I am in this location, I expect to observe this', or 'given that I am in this location, and plan to walk in this direction, I expect to end up in this other location'. In the terminology of the PP literature, the transition probabilities from hidden states to observations are provided by matrix A, while the transition probabilities between hidden states, conditioned under each policy, are provided by matrix B (Friston et al. [2015], [2017]). Preferred outcomes are provided in a separate matrix C.[5] Standard implementations of expected free energy minimization using partially observable Markov decision processes, provide the update equations for how the free energy minimizing agent should change its beliefs about the statistical structure of its environment (for a derivation of those equations, see, for example, Bruineberg et al. [2018b], table 2, appendix B).

A few points are worth emphasizing. First, and crucially, the beliefs about the structure of the environment are kept apart from the preferred outcomes that drive the agent's policy selection. The former are stored in matrices A and B, while the latter are stored in matrix C. Second, whereas canonical implementations provide update equations for matrices A and B (that is,

[5] Although this labelling of the transition probabilities seems specific to the PP literature, the general idea of a partially observed Markov decision process is well established. A standard machine learning textbook introduces them as 'basically a hidden Markov model augmented with action and reward nodes' (Murphy [2015], p. 331). Transitions between hidden states are conditioned on the agent's actions, and these actions themselves are selected in accordance with the agent's rewards. In preference PP terminology, transitions are conditioned on policies, and policies are selected in accordance with preferred outcomes.

equations that capture how the agent updates its beliefs about the environment after observation), far less has been written on implementations of expected free energy minimization that provide update equations for matrix *C*.[6]

In summary, preference PP can explain motivated behaviour through a set of states, specifically, preferred outcomes, with the following properties:

(1) Preferred outcomes are used as a benchmark in policy selection: how close will following a particular policy bring the agent to its preferred outcomes? The probability of selecting a policy is proportional to its proximity to preferred outcomes.

(2) Preferred outcomes are independent of the beliefs the agent has about the causal and statistical structure of its environment. To the extent that preferred outcomes have updating rules, these are independent of the rules for updating beliefs about the environment.

In other words, these mental states guide action selection, have a world-to-mind direction of fit, and are updated independently of the updating of beliefs about the structure of the environment. They thus have all the functional characteristics of desires. By contrast, the beliefs about the structure of the environment have all the functional characteristics of beliefs. Note that reflecting the nature of the creature, preferred outcomes could take the form of a wide range of pro-attitudes from basic drives and urges (for example, for food and shelter) to higher-level evaluative judgements (for example, about the favourability of chess positions or moral actions). However one fills in preferred outcomes, preference PP allows for the existence of distinct desires (broadly construed) that motivate action over and above belief-like states. Hence, preference PP is consistent with Humeanism.

To illustrate the basic points, consider an agent engaging a very simple blue or red environment. Let us assume that the values stored in its *C*-matrix are such that the agent has 'perceiving red' as a preferred outcome and 'perceiving blue' as an undesirable outcome. The agent starts out not knowing which parts of the environment are blue and which are red. As it explores its environment, it only ever encounters blue. The agent correctly infers that it inhabits a blue environment. This knowledge gets stored as a prior belief in its *A*-matrix (which stores the probabilities of observations given one's current state): 'Given that I am in this location, I expect to observe blue'. But throughout this exploration, its preferred outcomes (stored in the

6 One exception is the work on preference learning by Sajid et al. ([2021]). We return to this in section 4.2.

C-matrix) remain unchanged. It still has 'perceiving red' as a preferred outcome and 'perceiving blue' as an undesirable outcome. The simple fact that the agent can maintain these preferences while learning that its environment only contains blue implies that the agent is a Humean creature. It has distinct beliefs and desires that are updated independently.

4.2. Preference learning

In our discussion so far, we have presupposed that the updating rules for the *C*-matrix will need to be substantially different from the updating rules for the *A* and *B* matrices. After all, the *A* and *B* matrices try to approximate the structure of the agent's environment, while the *C*-matrix captures the agent's preferred outcomes. As detailed above, these are very different functional states. A crucial point is that to benefit from the resources offered by the distinction between belief-like and desire-like states, a Humean agent needs to have independent updating rules for both types of states. Consequently, in line with Klein ([2020]), value learning needs a non-Bayesian updating rule.

Recently, Sajid et al. ([2021]) have developed an account of preference learning that provides updating rules for the *C*-matrix and demonstrates how PP agents can learn preferred outcomes to guide policy selection, even where pre-specified preferred outcomes are absent.[7] On this account, the agent learns its preferred outcomes by engaging with its environment in much the same way that it learns about the structure of its environment: preferred outcomes are learned via Bayesian updating. To better understand the behaviour of such an agent, let us examine the details of the simulations by Sajid et al. ([2021]) of this learning process.

When placed into an unfamiliar environment, the agent is initially uncertain about the structure of the environment and its own preferred outcomes. At first, the agent engages in purely exploratory behaviour and learns the structure of the environment (that is, what to expect where). Next, the agent is equipped with the ability to learn preferred outcomes, and the outcomes observed more often are the outcomes the agent learns that it prefers. As the agent becomes less uncertain about what its preferences are, it will move away from exploring various outcomes and start to seek out its learned preferences.

What are we to make of such an updating rule for preferences? Let's transpose our PP agent to Italy where it spends considerable time in a village with only one restaurant. Not knowing what it wants, or what the items on the menu mean, the agent randomly picks a different item

[7] We would like to thank two reviewers for bringing this literature to our attention.

from the menu every night (careful not to order the same thing twice). This is not a bad strategy: the agent now knows what is on offer and how it tastes. In a separate second phase, the agent starts to learn its preferences. It does so by picking a random item from the menu each evening and by keeping a tally of its choices. Over time, it observes itself choosing spaghetti slightly more often than other dishes and starts to bias its ordering towards spaghetti, leading to even more observations of eating spaghetti. At some point, having eaten spaghetti consistently for several days in a row, the creature comes to the inevitable conclusion (paraphrasing Sajid et al. [2021], p. 30): 'I am the sort of creature that enjoys eating spaghetti', and happily eats its favourite food for the rest of its life.

Such an account of value learning seems deeply unsatisfactory. First, it is implausible that preferences are simply determined by the relative frequencies of outcomes encountered. Might the spaghetti-eater not eventually grow tired of spaghetti? The issue runs deeper: Of course, it is true that we learn what we like by trying out different things. If we are presented with two items from a menu, we can like one and dislike the other. Why? Because one tastes good, and the other one does not. But this is a distinction that a purely Bayesian account of value learning cannot make. In reality, repeated exposure is not necessary for acquiring preferences. Sometimes we immediately learn to prefer or disprefer certain outcomes (we can confidently judge good and bad taste after a single bite), and sometimes preferences are hard-wired (for example, aversion to pain). Bayesian updating over multiple rounds of observations will, therefore, often be the wrong place to look for preferences. Sajid et al. ([2021]) acknowledge the counterintuitive consequences, or potentially suboptimal strategies, implied by their Bayesian updating framework. As demonstrated by their simulations, in an environment where the agent encounters obstacles more frequently than the goal state, the PP agent will learn to prefer and seek out the obstacles rather than the goal state. Hence, this updating rule risks teaching the agent plainly counterproductive strategies.

On their account, preference learning is essentially a form of Bayesian updating applied to the preferred outcomes stored in the C-matrix. Preferred outcomes are cast as prior beliefs that the agent will encounter those outcomes, and the updating rules for likelihood and preferred outcomes are the same (they are both the experience-dependent updating of concentration parameters of a Dirichlet distribution). This suggests that preferred outcomes are not relevantly distinct from the beliefs about statistical regularities in the environment stored in matrices A and B. Consequently, the anti-Humean's problem resurfaces: we need an explanation of our ability to keep beliefs about statistical regularities fixed while independently revising our desires. The fact that preferred outcomes are stored in a separate matrix does not seem to help by itself. If prior beliefs across matrices are sensitive to the same evidence and are all revised via Bayesian updating, then preferred outcomes and other prior beliefs should be revised in tandem and converge over time. Indeed, the spaghetti-enjoying creature will both like and expect spaghetti every evening. Given the updating rules it is subject to, it seems difficult to explain how it can end up liking one thing and expecting another.

To conclude, a view on which preferred outcomes are updated via Bayesian updating faces a dilemma. If preferred outcomes are simply a form of prior beliefs subject to Bayesian updating, it becomes hard to provide a mechanism for value learning that is independent of general belief updating, and thereby addresses the problems associated with anti-Humeanism. By contrast, if preferred outcomes are not updated in tandem with beliefs about the structure of environment, the agent is somehow able to weigh evidence differently and arbitrate between updating its beliefs and preferred outcomes. Consequently, either PP insists on the parsimony of updating rules and is unable to explain all aspects of agency, or it insists on a broad explanatory scope, which requires independent value learning. In the latter case, we relax the strict Bayesian updating requirement and allow that preference PP is consistent with a Humean account of motivation.

The discussion above shows that the question of preference PP's Humeanism requires answering two questions. First, is there a desire-like element in preference PP that is separate from belief-like states? We have argued in section 4.1 that preference PP does contain a separate desire-like construct. Second, even if there is such a distinction, are the updating rules for both types of states sufficiently different to make use of that distinction? There are two options here: The first is to follow Sajid et al. ([2021]) in trying to subsume value learning under Bayesian inference. We have seen that this leads to an account of value learning that must confront the problems associated with anti-Humeanism. The other option is to treat preferred outcomes as a placeholder for some to-be-specified value learning mechanism. With the exception of the purely Bayesian account of preference learning discussed above, the current literature on preference PP models seems to leave open this approach. For PP to avoid the problems facing anti-Humeanism, whatever learning mechanism is slotted in ought to make preferred outcomes independently updatable from beliefs. This will likely require distinct value representations that do not reduce to prior beliefs, along with a non-Bayesian learning mechanism thereof. Both are characteristics of other prominent computational frameworks, such as reinforcement learning and Bayesian decision theory. In fact, nothing seems to prohibit preference PP from adopting an account of value learning from alternative frameworks except the aspiration for a theory of the mind that is both universal and based purely on prior beliefs and Bayesian updating.

5. Two Theories of Action

Pro-attitudes and their motivational role are central to standard accounts of agency. Some version of the Humean theory of motivation is assumed by most theories of action. The difference between optimistic PP and preference PP, and their respective explanatory potential with respect to agency, comes clearly into focus by looking at their implications for a theory of action. In this section, we outline the respective theories of action that optimistic PP and preference PP have on offer. Fundamental to both theories of PP is that all aspects of agency, including motivation and action selection, are cast as inference problems (Friston et al. [2012]; Friston et al. [2013]; Clark [2020]). Beyond that, the theories have quite different stories to tell.

5.1. Optimistic predictive processing: A revisionist theory of action

In optimistic PP, all mental state types are reduced to a single construct, namely, precision-weighted predictions.[8] To perform an action, the brain predicts that it is currently receiving the sensory input it would expect to receive if the action had already been performed (for example, the proprioceptive signals associated with having raised one's arm). These prediction errors are minimized by activating processes that move the body towards the predicted state. Descending predictions can thus serve as motor commands. They activate motor processes by triggering reflex arcs (neural pathways that control reflexes), which move the body (for example, by contracting muscles in the arm) towards a state where it receives the (proprioceptive) sensory signals associated with the predicted state (Adams et al. [2013]; Grünbaum and Christensen [2024]). Again, this assumes that the agent's predictions are stubborn in the face of prediction errors, so that the agent is driven to act instead of revising the predictions.

The stubbornness of predictions is accounted for by their precision: when the precision of a prediction is high, it resists revision and causes the agent to pursue the action that makes it come true. Optimistic PP thereby accounts for some aspects of motivation: precise predictions can drive action. If all goes well, the agent is able to prioritize those predictions that help her

[8] Some might prefer the alternative phrasing that the single fundamental construct is the prior beliefs about states and state transitions that generate predictions. We do not think much hinges on the choice of terminology here.

navigate her environment and meet her needs (Pezzulo et al. [2015], [2018]; Clark [2020]). This context-sensitive prioritization of predictions is non-trivial. The challenge is to provide an empirical account without presupposing an agent that can tweak its precision, as it sees fit. One issue with explaining action by a voluntary and context-sensitive tweaking of precision is that it turns a presumably sub-personal mechanism into a capacity governed by the person's will. Not only does this sneak motivation and desire in through the back door, but it also seems to commit the homunculus fallacy of trying to explain agency by positing another agent inside the agent, which itself requires explanation. This puts a heavy explanatory burden on precision.

Assuming that a proponent of optimistic PP can provide a satisfactory anti-Humean account of precision-weighting, such an account of precision needs to stay within the Bayesian commitments of the framework. Given both prior expectations, the prior precision over those expectations, and sensory input, the posterior expectations and precision over those expectations should approximate Bayesian inference. The resulting account would clearly be revisionist, as it compels us to revise standard conceptions of motivation and action. We have already seen how optimistic PP clashes with the Humean theory of motivation, unlike other computational frameworks, such as reinforcement learning and Bayesian decision theory. In these frameworks, probabilities and values are represented and computed independently, and value signals are necessary to motivate action.

There are also other aspects of agency that optimistic PP struggles to explain. Consider Buridan cases, where an agent faces a choice between multiple incompatible options that are equally desirable and probable. The canonical case is of a donkey placed right between two equally attractive bales of hay, having to make a choice so as not to starve. Arguably, we often face choices like this in real life (for example, when faced with the choice between equally desirable holiday destinations or items on a menu). The optimistic PP equivalent of such cases seems to be an agent facing two incompatible courses of action, the information about which is equally precise. How can the optimistic PP agent break the tie? It seems puzzling how a predictive brain that deals solely in predictions, could come to favour one option when the precision of all options is identical.[9] According to some, such cases require distinct intentions to break the tie and motivate action (Bratman [1987]). To overcome the impasse, we arbitrarily form an intention to pursue one option. The intention then motivates us to act in accordance

[9] Ransom et al. ([2017]) raise a similar critique. They argue that PP cannot explain our ability to voluntarily shift attention between two overlapping film-streams when the signals from each stream are equally precise.

with it, and structures further planning and practical reasoning. To some, this suggests that intentions cannot be reduced to belief-desire pairs, since such pairs cannot play the same roles in planning and practical reasoning. However, there appears to be no element in optimistic PP that could play the role of intentions to resolve ties in Buridan cases and in structuring further planning and practical reasoning.

Optimistic PP might respond that if the Buridan agent is able to break the tie, it would merely demonstrate that the agent has deep or evolving priors that somehow make the expected precision of one option relatively higher.[10] This effectively amounts to denying the possibility of Buridan cases, because there would always be some prior to the effect that the options are not truly considered equally desirable and probable. Furthermore, this response would appear to settle optimistic PP with some additional problems. This strategy reads off the agent's preferences from her choices. But if this is how preferences are determined, there seems to be no room for any divergence between motivation and action. Some accounts of action reject the claim that actions are always caused by the strongest desire and argue that agents can act against their strongest desire. That is, agents are able to do something they find truly undesirable or refrain from doing what they find truly desirable (Schueler [1995]). Accepting a distinction between desire and intention enables this type of divergence, because intentions can then play the role of controlling actions, even when they run counter to our strongest desires (Holton [2009]). Yet, with only precision-weighted predictions at its disposal, this is not a divergence optimistic PP is able to accommodate. In sum, optimistic PP is not only incompatible with Humeanism, but also with accounts of action that distinguish intentions from belief-desire pairs.[11]

[10] This is essentially Clark's ([2017]) response to Ransom et al. ([2017]): Your voluntary shift of attention takes the form of a counterfactual prediction that you are currently perceiving one of the film-streams, which, in a self-fulfilling manner, increases the expected precision of inputs from that stream, thereby making you perceive that stream. Clark suggests that such a counterfactual prediction can be understood as a desire to see one film-stream rather than the other. Since Clark attributes both belief-like and desire-like roles to a single primitive (precision-weighted predictions), his explanation is of the optimistic PP and anti-Humean variety. The problems facing anti-Humeanism seem just as pressing for mental actions, such as voluntary shifts of attention, as for any other action type. Moreover, as explained in the main text, since each option can be considered equally desirable, distinct intentions are arguably needed in addition to beliefs and desires to break the tie.

[11] A reviewer suggested another response to Buridan cases: It might be that intrinsic noise disturbs the equilibrium. But this simply amounts to denying the phenomenon. If the equilibrium is disturbed by noise, the agent will not be faced with a choice between equivalent options. Our argument is that if Buridan cases exist, they pose a challenge to optimistic PP. We have some reasons for thinking that such cases exist. Hence, claiming that they do not requires a substantial argument.

5.2. Preference predictive processing: A non-revisionist theory of action

According to preference PP, actions are brought about through the inference and selection of optimal policies. Policy selection is the process of inferring which policy minimizes expected free energy; that is, which policy strikes the optimal balance between reducing uncertainty and leading to highly weighted preferred outcomes (Pezzulo et al. [2018]; Parr et al. [2022]). This process requires a generative model that can model more and more abstract relations between actions and outcomes and score them on how well they minimize expected free energy.

As argued earlier, preference PP is consistent with a distinction between beliefs and desires. To make clearer how preference PP relates to standard conceptions of agency, let us take a closer look at policy selection under preference PP. Expected free energy is composed of both expected ambiguity and risk. Expected ambiguity encodes expectations about how much uncertainty will be reduced by pursuing a certain policy or how much information we stand to gain under a certain policy. Risk encodes expectations about the preferred outcomes a certain policy might bring about.

The evaluation of expected ambiguity is made possible by the fact that an agent has access not just to what it believes, but also to its uncertainty about its beliefs. Expected ambiguity evaluation is therefore dependent on a particular kind of belief: 'Given that I am in this state and execute this policy, I expect to observe a state with this much uncertainty'. The expected ambiguity term promotes policies that lead to observations that reduce uncertainty, while penalizing policies that do not, and from which the agent, therefore, expects to learn little. The evaluation of the risk is made possible by two different kinds of mental states. On the one hand, there are beliefs of the form 'given that I am in this state and execute this policy, I expect to observe this state'. On the other hand, there are desires, or preferred outcomes, of the form 'I want to observe this state'. The risk term promotes policies that lead to observations that match preferred outcomes, while penalizing policies that do not.

Some critics might argue that the best interpretation of preferred outcomes is something like: 'Given that I am this kind of creature, these are the kinds of things I expect to observe' (note that these expectations need not be conscious). Hence, one might argue that preferred outcomes are more akin to beliefs (or predictions) than desires, and therefore, preference PP does not contain distinct desire-like states (we will discuss this objection further in section 6). In section 4, we argued that preferred outcomes are (1) used as a benchmark for action selection, and (2) independent of beliefs the agent has about the causal structure of its environment.

A state with these characteristics is best understood as a desire, not as a kind of belief. The doxastic gloss of preferred outcomes one sometimes encounters in the literature simply mischaracterizes the role preferred outcomes actually play in the preference PP architecture.

By separating out policy selection from state estimation, preference PP has the theoretical tools for both the modulation of precision of sensory signals and the modulation of precision of policies (Parr and Friston [2017], [2019]). The distinction between the two forms of precision modulation allows for more flexibility. For example, an agent that has a precise high-level policy of following a diet—that is to say, is very motivated to follow a diet—will lower the expected precision of low-level gustatory outcomes related to consuming high-calorie foods and increase the expected precision of signals and policies related to eating healthier alternatives (Pezzulo et al. [2018]). This explains how one action (for example, eating the healthy option) is selected over alternatives (for example, eating cake). Where the optimistic PP agent needs to stubbornly predict that she will not eat the cake, the preference PP agent can infer that eating the cake will meet certain preferences (for example, for high-calorie foods), but nonetheless opt for an alternative diet-congruent policy by making the alternative highly precise, to the point of making it the optimal policy. In short, higher precision of policies translates to higher motivational force.

How about intentions? Recently, some have developed a notion of intentions within a preference PP framework. Friston et al. ([unpublished]) argue that in intentional behaviour, the agent tries to bring about so-called preferred latent states when selecting policies. Latent states are the presumed but not themselves observable causes of sensory input. A preferred latent state is, they suggest, simply a prior belief that the agent will bring that state about. In action theory, intentions are often distinguished from beliefs and desires by their distinctive functional and normative roles in practical reasoning and planning. If intentions are simply prior beliefs over latent states, it is not clear that they are sufficiently distinct from beliefs and desires to play the distinctive roles often attributed to intentions.

Another potential interpretation is to identify intentions with selected policies. Since selected policies result from optimal belief-desire pairs (that is, those that minimize expected free energy), this seems to imply that intentions are reducible to belief-desire pairs. As argued above, this clashes with accounts that attribute distinctive and irreducible roles to intentions. Others maintain that the belief-desire account can explain all aspects of intentions (Sinhababu [2013]). We do not intend to settle this complex issue. Our aim is simply to clarify what theories of action are available within a preference PP framework.

Revisionist aspirations are considered important in some PP circles. However, such aspirations are optional within a preference PP framework. Preference PP does not require any major revisions to standard conceptions of motivation and action. In this respect, preference PP deviates little from other prominent computational frameworks, such as reinforcement learning and Bayesian decision theory, which also contain distinct representations and computations of value. Thus, preference PP is potentially much less revisionist than its predecessor, optimistic PP, and does not entail the radically revisionist programme advocated by some theorists. Preference PP might aspire to offer a universal account of agency, and, perhaps, the mind in general, but this universality is traded off against the supposed simplicity of the framework and its formalisms.

6. What about the Free Energy Principle?

6.1. The low road and the high road

So far, we have been pursuing what some PP theorists have called the low road to PP (Friston [2019]; Parr et al. [2022], chap. 2). The low road starts from the assumption that the brain is a Bayesian inference engine trying to optimize its model of the causes of its sensory input. PP is then proposed as an explanation of how the brain can solve this inferential problem and the neurocognitive mechanisms involved in this process. By enriching PP models with whatever constructs necessary to explain the empirical data, PP might gradually come to explain more and more aspects of cognition and behaviour, including action, motivation, planning, and decision-making.

The low road is sometimes contrasted with the high road to PP. The high road takes as its starting point fundamental questions about what properties systems that manage to persist must have. According to Friston ([2019], p. 177), 'any system that exists will appear to model and predict its exchange with the environment'. More specifically, any self-organizing system will necessarily engage in (or necessarily appear to engage in) the minimization of free energy. This idea is known as the free energy principle (FEP). According to proponents of the high road, PP turns out to be a necessary feature of self-organizing systems. For this reason, the high road has been described as a 'top-down journey from near existential nihilism to the riches of predictive processing' (Friston [2019], p. 175).

Let us unpack the FEP. A living organism must resist a tendency to disintegrate; that is, it must keep its internal states within a viable range as reflected by their homeostatic properties.

To do so, it must conserve a boundary that distinguishes it from its environment. Under the FEP, this boundary is formalized as a Markov blanket (Friston [2013]). Markov blankets are supposed to partition states into those internal to the system, external to the system, and the states of the boundary itself. Some boundary states are influenced by external states (namely, sensory states) and some by internal states (namely, active states). States that, according to the organism's model of the world, are expected to be incompatible with its continued existence are deemed surprising (in a technical sense). These states are deemed unlikely to occur when the organism inhabits a hospitable environment. Since calculating surprise directly would require knowing all the hidden states of the world that cause the sensory input, it is impossible for any organism to calculate this directly. Instead, it must minimize variational free energy, which is an upper bound on surprise. The organism thus effectively minimizes surprise in the only tractable manner. In other words, living systems expect to occupy states compatible with their continued existence. By minimizing variational free energy, the system keeps its internal states within a range consistent with its survival.[12]

In short, the high road involves developing so-called process theories that align with the FEP. These theories account for the structure and functions of neurocognitive mechanisms, which are consistent with the imperative of minimizing free energy. Under this approach, PP is essentially the suite of such process theories. It is important to emphasize that preference PP models, which explain policy selection in terms of expected free energy minimization, are not committed to the broader claims of the FEP. Variational free energy and the expected free energy of policies are not the same. Variational free energy minimization serves to model and predict the environment based on past and present observations. Expected free energy minimization, by contrast, pertains to action selection based on expectations about the consequences of future actions (see Parr et al. [2022], pp. 31–39).

6.2. The desert landscape

Why discuss the FEP? Because some claim that the FEP entails a desert landscape view of the mind; that is, a minimalist ontology in which 'there are neither goals nor reward signals as such' (Clark [2013], p. 200). All that really exists is self-organizing systems that appear to

[12] As some authors have argued (Seth [2015]; Pezzulo and Cisek [2016]), this makes the FEP a modern version of cybernetics, according to which control consists in using feedback signals to keep essential internal variables within an expected range.

model and predict their environment via free energy minimization (Friston [2019]). Although resisted by Clark (2013), this view is (at least sometimes) endorsed by Friston ([2019]) and other proponents of the FEP (Ramstead et al. [2019]) and presented as an inevitable consequence of the FEP. The desert landscape interpretation of the FEP denies the existence of pro-attitudinal constructs, such as value, reward, goals, drives, and desires. Therefore, if this radical interpretation is true, then PP (in any guise) is necessarily anti-Humean under the FEP.

There has been much disagreement about the high road that takes the FEP as its theoretical starting point, and what it entails exactly (Clark [2013]; Friston [2019]; Williams [2022]). For our purposes, the relevant question is whether the FEP entails a desert landscape ontology, which would restrict process theories to anti-Humean varieties with no pro-attitudinal constructs. As we will see, this depends on how the FEP is interpreted.

On one reading, the FEP strives to explain how self-organizing systems manage to persist over time by means of mechanisms that implement free energy minimization. If free energy minimization is a necessary condition on self-organizing systems, and any mechanism implementing it precludes pro-attitudinal states, then this would entail a desert landscape view of the mind.

On another reading, the FEP merely posits that all self-organizing system can be redescribed as if they minimize free energy. Under such a reading, the FEP places no constraints on how such systems minimize free energy, and process theories are free to include pro-attitudinal constructs. The specific mechanisms involved could be very different for rocks, oil drops, and humans. Under this interpretation, PP process theories could include pro-attitudinal constructs, so long as it remains true that the target system can be described as it if it minimizes free energy—even if this is not strictly the objective of the mechanisms underlying the system's behaviour. This interpretation does not entail a desert landscape ontology and is consistent with Humeanism.[13]

Others have argued that the FEP is consistent with a folk-psychological distinction between beliefs and desires. Smith et al. ([2022]) argue that there are terms within the expected free energy formalism (that is, within preference PP), which can be functionally identified with

[13] For an argument that the FEP places no necessary constraints on explanations of how self-organizing systems manage to maintain their existence, see (Williams [2022]).

desire-like constructs with a world-to-mind direction of fit as described by folk psychology.[14] Even when described by the FEP, they argue, the organism can still be described as having desires. This illustrates that there is no clear consensus that the FEP entails a desert landscape ontology. For those who deny this implication of the FEP, there need be no conflict between the FEP and a Humean interpretation of preference PP.

7. Conclusion

We have explored the intricacies of the predictive processing framework by uncovering two distinct theories within it and their distinct implications for motivation and action. The difference between these has not been properly appreciated. Optimistic PP bases all processing on optimistic priors and entails a revision of standard accounts of motivation and action. This gives rise to significant explanatory challenges. Sticking to its simplistic formalism, optimistic PP must relinquish its ambition to provide a universal account of the mind. There are aspects of motivation and action that seem beyond its explanatory scope. By contrast, preference PP posits that actions are selected to minimize expected free energy and aligns more closely with standard accounts of motivation and action in philosophy and cognitive science.

Contrary to some radical interpretations, the FEP does not necessitate a fundamental overhaul of standard desire-like or pro-attitudinal constructs. Preference PP better explains more aspects of motivation and action. However, the broader explanatory scope requires moving beyond attempts to reduce all mental phenomena to a single process of Bayesian belief updating. In its preference PP incarnation, PP has instead come to resemble other prominent computational frameworks implementing a distinction between beliefs and desires, such as reinforcement learning and Bayesian decision theory. The general lesson is that a tension exists between the parsimony often aspired to in PP theories and accepting enough primitives to give a complete account of agency and the mind.

[14] Though similar in some respects, our analysis is different in others. One difference is that we focus on the role of desire-like constructs in philosophical and scientific theories of motivation and action, not simply on consistency with how they are described in folk psychology. Another is our focus on value learning and the need for a non-Bayesian account thereof.

Acknowledgements

Thanks to the referees for this journal for helpful comments. This material was presented to audiences at the Rome–Copenhagen Workshop on Cognitive Motor Control and Active Inference and at the 1st Annual Web Conference of the International Society for the Philosophy of the Sciences of the Mind. We are grateful for the comments we received on these occasions. Special thanks to Giovanni Pezzulo and the members of the CoInAct Research Group for helpful discussions. Support for this project was provided by the Carlsberg Foundation (Semper Ardens Advance Grant CF22-1111), awarded to Thor Grünbaum.

Frederik T. Junker
Section for Philosophy, Department of Communication
University of Copenhagen
Copenhagen, Denmark
and
CoInAct Research Group
University of Copenhagen
Copenhagen, Denmark
ftj@hum.ku.dk

Jelle Bruineberg
Center for Subjectivity Research, Department of Communication
University of Copenhagen
Copenhagen, Denmark
and
CoInAct Research Group
University of Copenhagen
Copenhagen, Denmark
bruineberg@hum.ku.dk

Thor Grünbaum
Section for Philosophy, Department of Communication
University of Copenhagen
Copenhagen, Denmark
and
CoInAct Research Group
University of Copenhagen
Copenhagen, Denmark
and

Department of Psychology
University of Copenhagen
Copenhagen, Denmark
tgr@hum.ku.dk

References

Adams, R. A., Shipp, S. and Friston, K. J. [2013]: 'Predictions Not Commands: Active Inference in the Motor System', *Brain Structure and Function*, **218**, pp. 611–43.

Bratman, M. [1987]: *Intention, Plans, and Practical Reason*, Cambridge, MA: Harvard University Press.

Bruineberg, J., Kiverstein, J. and Rietveld, E. [2018a]: 'The Anticipating Brain Is Not a Scientist the Free-Energy Principle from an Ecological-Enactive Perspective', *Synthese*, **195**, pp. 2417–44.

Bruineberg, J., Rietveld, E., Parr, T., van Maanen, L. and Friston, K. J. [2018b]: 'Free-Energy Minimization in Joint Agent–Environment Systems: A Niche Construction Perspective', *Journal of Theoretical Biology*, **455**, pp. 161–78.

Clark, A. [2013]: 'Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science', *Behavioral and Brain Sciences*, **36**, pp. 181–204.

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

Clark, A. [2017]: 'Predictions, Precision, and Agentive Attention', *Consciousness and Cognition*, **56**, pp. 115–19.

Clark, A. [2020]: 'Beyond Desire? Agency, Choice, and the Predictive Mind', *Australasian Journal of Philosophy*, **98**, pp. 1–15.

Colombo, M. [2017]: 'Social Motivation in Computational Neuroscience', in J. Kiverstein (*ed.*), *The Routledge Handbook of Philosophy of the Social Mind*, New York: Routledge, pp. 320–40.

Davidson, D. [1963]: 'Actions, Reasons, and Causes', *Journal of Philosophy*, **60**, pp. 685–700.

Friston K. J. [2009]: 'The Free-Energy Principle: A Rough Guide to the Brain?', *Trends in Cognitive Sciences*, **13**, pp. 293–301.

Friston, K. J. [2010]: 'The Free-Energy Principle: A Unified Brain Theory?', *Nature Reviews Neuroscience*, **11**, pp. 127–38.

Friston, K. J. [2013]: 'Life as We Know It', *Journal of the Royal Society Interface*, **10**, available at <doi.org/10.1098/rsif.2013.0475>.

Friston, K. J. [2013]: 'Active Inference and Free Energy', *Behavioral and Brain Sciences*, **36**, pp. 212–13.

Friston, K. J. [2019]: 'Beyond the Desert Landscape', in M. Colombo, E. Irvine and M. Stapleton (*eds*), *Andy Clark and His Critics*, New York: Oxford University Press, pp. 174–90.

Friston, K. J., Daunizeau, J. and Kiebel, S. J. [2009]: 'Reinforcement Learning or Active Inference?', *PLOS One*, **4**, available at <doi.org/10.1371/journal.pone.0006421>.

Friston, K. J, Samothrakis, S. and Montague, R. [2012]: 'Active Inference and Agency: Optimal Control without Cost Functions', *Biological Cybernetics*, **106**, pp. 523–41.

Friston, K. J., Thornton, C. and Clark, A. [2012]: 'Free-Energy Minimization and the Dark-Room Problem', *Frontiers in Psychology*, **3**, available at <doi.org/10.3389/fpsyg.2012.00130>.

Friston, K. J., Schwartenbeck, P., Fitzgerald, T., Moutoussis, M., Behrens, T. and Dolan, R. J. [2013]: 'The Anatomy of Choice: Active Inference and Agency', *Frontiers in Human Neuroscience*, **7**, available at <hdoi.org/10.3389/fnhum.2013.00598>.

Friston, K. J., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T. and Pezzulo, G. [2015]: 'Active Inference and Epistemic Value', *Cognitive Neuroscience*, **6**, pp. 187–214.

Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P. and Pezzulo, G. [2017]: 'Active Inference: A Process Theory', *Neural Computation*, **29**, pp. 1–49.

Friston, K. J., Salvatori, T., Isomura, T., Tschantz, A., Kiefer, A., Verbelen, T., Koudahl, M. et al. [unpublished]: 'Active Inference and Intentional Behaviour', available at <arxiv.org/html/2312.07547v2>.

Grünbaum, T. and Christensen, M. S. [2024]: 'The Functional Role of Conscious Sensation of Movement', *Neuroscience and Biobehavioral Reviews*, **164**, available at <doi.org/10.1016/j.neubiorev.2024.105813>.

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

Hohwy, J. [2016]: 'The Self-Evidencing Brain', *Noûs*, **50**, pp. 259–85.

Holton, R. [2009]: *Willing, Wanting, Waiting*, Oxford: Oxford University Press.

Klein, C. [2018]: 'What Do Predictive Coders Want?', *Synthese*, **195**, pp. 2541–57.

Klein, C. [2020]: 'A Humean Challenge to Predictive Coding', in D. Mendonça, M. Curado and S. S. Gouveia (*eds*), *The Philosophy and Science of Predictive Processing*, London: Bloomsbury, pp. 25–38

Lewis, D. K. [1988]: 'Desire as Belief', *Mind*, **97**, pp. 323–32.

Lewis, D. K. [1996]: 'Desire as Belief II', *Mind*, **105**, pp. 303–13.

Miller Tate, A. J. [2021]: 'A Predictive Processing Theory of Motivation', *Synthese*, **198**, pp. 4493–521.

Murphy, K. P. [2015]: *Probabilistic Machine Learning: An Introduction*, Cambridge, MA: MIT press.

Parr, T. and Friston, K. J. [2017]: 'Working Memory, Attention, and Salience in Active Inference', *Scientific Reports*, **7**, available at <doi.org/10.1038/s41598-017-15249-0>.

Parr, T. and Friston, K. J. [2019]: 'Attention or Salience?', *Current Opinion in Psychology*, **29**, available at <doi.org/10.1016/j.copsyc.2018.10.006>.

Parr, T., Pezzulo, G. and Friston, K. J. [2022]: *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*, Cambridge, MA: MIT Press.

Pezzulo, G. and Cisek, P. [2016]: 'Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition', *Trends in Cognitive Sciences*, **20**, pp. 414–24.

Pezzulo, G.; Rigoli, F. and Friston, K. J. [2015]: 'Active Inference, Homeostatic Regulation, and Adaptive Behavioural Control', *Progress in Neurobiology*, **134**, pp. 17–35.

Pezzulo, G.; Rigoli, F. and Friston, K. J. [2018]: 'Hierarchical Active Inference: A Theory of Motivated Control', *Trends in Cognitive Sciences*, **22**, pp. 294–306.

Ramstead, M. J., Constant, A., Badcock, P. B. and Friston, K. J. [2019]: 'Variational Ecology and the Physics of Sentient Systems', *Physics of Life Reviews*, **31**, pp. 188–205.

Ransom, M., Fazelpour, S. and Mole, C. [2017]: 'Attention in the Predictive Mind', *Consciousness and Cognition*, **47**, pp. 99–112.

Sajid, N., Ball, P. J., Parr, T. and Friston, K. J. [2021]: 'Active Inference: Demystified and Compared', *Neural Computation*, **33**, pp. 674–712.

Schueler, G. F. [1995]: *Desire: Its Role in Practical Reason and the Explanation of Action*, Cambridge, MA: MIT Press.

Schroeder, T. [2004]: *Three Faces of Desire*, New York: Oxford University Press.

Searle, J. R. [1983]: *Intentionality: An Essay in the Philosophy of Mind*, Cambridge: Cambridge University Press.

Seth, A. K. [2015]: 'The Cybernetic Bayesian Brain: From Interoceptive Inference to Sensorimotor Contingencies', in T. Metzinger and W. Wiese (*eds*), *Philosophy and Predictive Processing*, Frankfurt am Main: MIND Group.

Seth, A. K., Millidge, B., Buckley, C. L. and Tschantz, A. [2020]: 'Curious Inferences: Reply to Sun and Firestone on the Dark Room Problem', *Trends in Cognitive Sciences*, **24**, pp. 681–83.

Sinhababu, N. [2013]: 'The Desire–Belief Account of Intention Explains Everything', *Noûs*, **47**, pp. 680–96.

Smith, M. [1987]: 'The Humean Theory of Motivation', *Mind*, **96**, pp. 36–61.

Smith, R., Ramstead, M. J. D. and Kiefer, A. [2022]: 'Active Inference Models Do Not Contradict Folk Psychology', *Synthese*, **200**, available at <doi.org/10.1007/s11229-022-03480-w>.

Sun, Z. and Firestone, C. [2020]: 'The Dark Room Problem', *Trends in Cognitive Sciences*, **24**, pp. 346–48.

Van de Cruys, S., Friston, K. J. and Clark, A. [2020]: 'Controlled Optimism: Reply to Sun and Firestone on the Dark Room Problem', *Trends in Cognitive Sciences*, **24**, pp. 680–81.

Wiese, W. [2017]: 'Action Is Enabled by Systematic Misrepresentations', *Erkenntnis*, **82**, pp. 1233–52.

Williams, B. [1979]: 'Internal and External Reasons', in R. Harrison (*ed.*), *Rational Action*, Cambridge: Cambridge University Press, pp. 101–13.

Williams, D. [2022]: 'Is the Brain an Organ for Free Energy Minimisation?', *Philosophical Studies*, **179**, pp. 1693–714.

Yon, D., de Lange, F. P. and Press, C. [2019]: 'The Predictive Brain as a Stubborn Scientist', *Trends in Cognitive Sciences*, **23**, pp. 6–8.

Yon, D., Heyes, C. and Press, C. [2020]: 'Beliefs and Desires in the Predictive Brain', *Nature Communications*, **11**, available at <doi.org/10.1038/s41467-020-18332-9>.