Can Hierarchical Predictive Coding Explain Binocular Rivalry?
Julia Haas

Abstract: Hohwy et al.'s (2008) explanation of binocular rivalry is taken as a classic illustration of hierarchical predictive coding's explanatory power. I revisit the account and show that it cannot explain the role of reward in binocular rivalry. I show that a modified version of the predictive processing approach may account for the role of reward by recasting it as a form of optimism bias. Accepting this account, however, is at odds with the epistemic commitments favored by proponents of hierarchical predictive coding.

Keywords: Hierarchical predictive coding, binocular rivalry, reward

1. Introduction

The visual system must estimate the shapes and sizes of objects from retinal stimulation. However, the distribution of light on the retina is consistent with an indefinitely large combination of stimulus objects and patterns of object illumination. This is known as the *inverse optics* or *underdetermination problem* (Helmholtz, 1867).

There are two general approaches to characterizing and explaining the underdetermination problem. The first is *pragmatic* in emphasis, characterizing perceptual systems as generating representations of distal stimuli for the purposes of action (for a review, see Jacob, 2015). That is, on this type of view, how an individual perceives an object is thought to depend in part on what she is going to do with it. For example, *enactivist* theories propose that perceptual experiences depend on dispositional motor responses (Noe, 2004), and *efferent readiness* theories propose that perceptual states prepare an observer to move and act in relation to her environment (Coren, 1986; Vishton et al. 2007). On broadly pragmatic approaches to perception, our perceptual systems can misinform us about what is out there in the world, but still be thought to perform their function if they enable us

to act appropriately, as when an individual perceives the reconstructed shape of an otherwise occluded object.

The second general approach to the underdetermination problem characterizes perception in much more robustly *epistemic* terms (for a discussion and review, see Siegel & Silins, 2015). On this type of view, perceptual systems produce veridical representations of distal stimuli principally for the purposes of *knowledge*; in many cases, perception is simply defined as "a way of acquiring information, beliefs, or knowledge about the world by means of the senses [...] 'perceive' and its derivatives 'see', 'hear', and the like, are usually taken to be success verbs" (Macpherson 2009, 502). On this type of view, then, how an individual perceives an object depends on what is really out there in the world, and whether and how her perceptual system's is able to represent it accurately. Notably, these representations are analyzed independently of actions undertaken by the organism, and they are taken to be better and worse to the degree that they 'get it right' about the external world. Hence, if an individual claims to see an object in front her when in fact it isn't there, her perceptual system is taken to have failed.

Proponents of hierarchical predictive coding widely adopt an epistemic approach to perception, emphasizing the twin roles of knowledge and inference. For example, Clark (2013, p. 182) describes perception as "probabilistic, knowledge-driven inference;" Hohwy (2013, p. 2) states that, "By testing hypotheses[,] we get the world right;" and Friston (2018, p. 1019) observes that perception constructs explanations for "what's going on out there." Even on views according to which the main function of knowledge is to enable action, such as on the view put forward by Wiese and Metzinger (2017, p.3) the authors maintain that "one's brain constantly forms statistical estimates, which function as representations of 'what is currently out there in the world.'"[1]

---

[1] Clark (2015, Section 6.7) offers a nuanced consideration of the relationship between the inferential nature of perception and the role of perception in guiding action. Describing the tension between what he calls the 'intellectualist' and 'non-reconstructive' approaches to perception – where the latter denies that perception must involve rich internal

These epistemic considerations inform predictive coding's engagement with phenomena such as *binocular rivalry*. Binocular rivalry occurs when one stimulus is shown to one eye at the same time as a different stimulus is shown to the other. The resulting experience is of the two images alternating back and forth. For example, if one eye is shown an image of a face and the other eye is shown an image of a house, then, rather than seeing the face and the house superimposed over one another, the experience is of seeing a face, then a house, and so on. Since binocular rivalry is a case of perception clearly not representing 'what's going on out there,' it constitutes an important test case for theories that aim to explain why we have particular perceptual states at particular times.

To try and address this challenge, Jakob Hohwy and colleagues' (2008) self-described 'epistemological' explanation recasts the problem of binocular rivalry in terms of Bayesian inference. The authors propose to explain the phenomenon by providing a rigorously parsimonious explanation of why the two stimuli appear alternate in visual perception (p. 700). The approach is widely cited across the cognitive sciences as an illustration of predictive coding's ubiquity in the brain, as well as its explanatory power as a general theory of the mind (e.g., see Clark (2013) for a prominent discussion; see also Metzinger & Wiese, 2017).

In this paper, I argue that Hohwy et al.'s (2008) epistemological account fails to explain the role of reward in perceptual dominance, afeature of binocular rivalry. There are two types of perceptual dominance: where one percept is seen first (dominance at *onset*), and where one is seen for a longer period of time (dominance over duration). The epistemological solution fails to explain

---

models - he argues that "'inference', as it functions in the PP story, is not compelled to deliver internal states that bear richly reconstructive contents. It is not there to construct an inner realm able to stand in for the full richness of the external world. Instead, it may deliver efficient, low-cost strategies whose unfolding and success depend delicately and continuously upon the structure and ongoing contributions of the external realm itself, as exploited by various forms of action and intervention" (p. 191). Note, however, that Clark's point here is still largely one about the relative *richness* or sparsity of a given representation, rather than discussion of its relative *veracity* relative to utility for guiding action. In other words, Clark acknowledges that perception needn't reconstruct a comprehensive picture of the external world, but he does not deny that the job of perception remains, to whatever degree of detail, to tell us about the world. So, Clark is open to a non-reconstructive view of perception; but he remains committed to an *epistemic* account of perception. Thanks to an anonymous reviewer for encouraging me to clarify this point.

both types of dominance and thus fails to provide a comprehensive explanation of the perceptual data.

Next, I show that a modified version of the predictive coding approach can account for the role of reward in both types of perceptual dominance. However, this model no longer offers an epistemic picture of perception: instead, it characterizes perceptual experience in pragmatic terms, making perceptual inferences conditional on inferences about policies and, specifically, control states. Proponents of predictive coding thus cannot remain committed to the strictly epistemic picture of perception and at the same time explain the phenomenon of binocular rivalry. I consider the implications of this impasse for predictive coding accounts of perception, as well as for rival computational approaches to the mind, most notably, for reinforcement learning views.

I begin by sketching the epistemological view and its proposed explanation of binocular rivalry.

2. Hohwy et al.'s (2008) epistemological solution

2.1 The general view

Hierarchical predictive coding is one answer among many to the question of how perceptual systems solve the underdetermination problem.[2] As its name suggests, the view proposes that perceptual systems actively make predictions about what their sensory input should be, rather than passively using inputs to draw inferences about the world, as had been previously thought. Perceptual systems compare these predictions with actual sensory input to calculate *prediction errors*, which reflect the difference between a given predicted and actual input. In a process known as *prediction error minimization*, the systems aim to minimize such prediction errors by continually revising their predictions in light of incoming sensory inputs.

---

[2] For an excellent review and discussion other approaches, see Rescorla (2015).

The view is *hierarchical* because it proposes that higher levels send predictions about the input to lower levels, and lower levels send bottom-up prediction errors which signal any discrepancies between the top-down predictions and actual input. This structure is thought to help explain how the brain is able to implement what are otherwise intractable inferences.

The view is *epistemological* in the sense described above, that is, in that the perceptual systems are in the business of 'getting it right' about the external world. As Howhy et al. (2008, p. 688) observe, "The motivation behind this approach is the idea that binocular rivalry is an epistemic response to a seemingly incompatible stimulus condition where two distinct objects occupy the same spatiotemporal location."

The view is also epistemological in a second sense. The brain's inferences - based on predictions, prediction errors, and prediction error-minimization - are described in terms of Bayesian inference. That is, the view proposes that the brain updates its predictions by approximating *Bayes rule*, which specifies how we should update an existing probability in light of a new piece of evidence. Specifically, the rule dictates that we update probabilities according to Bayes theorem:

$$P(H|E) = P(E|H)P(H)/P(E) \qquad (1)$$

That is, our posterior probability for a given hypothesis—the probability of the hypothesis conditional on our evidence ($P(H|E)$)—should be equal to the product of the probability of the evidence conditional on our hypothesis ($P(E|H)$) and the probability of the evidence, divided by the unconditional probability of the evidence. Applied accordingly, the posterior probability $P(H|E)$ replaces the prior probability $P(H)$ until new evidence is introduced, and so on. The brain thus also applies a distinctly epistemic type of inference to solve the underdetermination problem.

Perceptual inference is also determined by the *precision* or likelihood of the sensory data. Precision accounts for the uncertainty or noise associated with sensory experience (Hohwy 2013, p. 59). For example, visual inputs are thought to be broadly reliable, but may be degraded underwater

or in the dark. Precision measures act as weightings on bottom-up precision errors in the aforementioned inferences.

Finally, some defenders of predictive processing hold that action results from 'active inferences', which use the same framework of prediction error-minimization. In perceptual inference, prediction error-minimization results in updating percepts. By contrast, in active inference, prediction error-minimization is used to drive behavior. To illustrate the idea, take the idea of someone wanting a glass of water. On this framework, the mind *predicts* that it is reaching for a glass of water. Since this is initially not the case, this prediction results in substantial prediction error. The mind minimizes this error, not by updating its prediction, but by causing the arm to move toward the cup. Error is minimized via active inference when the prediction is realized (i.e. the person drinks). A central virtue of this approach is that the same framework aims to account for both perception and action, the resulting picture is of an organism or agent that alternates back and forth between perception and action.

2.2 The problem

Hohwy et al. (2008, p. 690) recast the problem of binocular rivalry in correspondingly epistemological terms. They propose that a successful view must explain three features of the phenomenon.[3]

The first explanandum is called the *individual vs. aggregate selection* problemThe question is this: Why does a participant only see one of the two stimuli at a time, rather than seeing some combination of the two? For example, why does the participant perceive a *house* or a *face*, rather than

---

[3] Hohwy et al. divide the problem into two: a selection problem and an alternation problem. Since the selection problem itself consists of two components, however, and one of these components plays an important role in what follows, I have preserved their general analysis but split the classification into three.

a house superimposed over a face (a *houseface*)? In Bayesian terms, it raises the question of why the hypothesis *face* is selected over the conjunctive hypothesis *house-face*.

The second explanandum is called the *individual vs. individual selection problem*. How does the perceptual decision select a given stimulus for perception? In other words, given that the participant only sees one stimulus, what procedure determines which one the participant sees first, e.g., a house rather than a face, or vice versa? In Bayesian terms, it asks why the hypothesis *face* is favored over the hypothesis *house*.

Finally, the third explanandum is called the *alternation problem*. The alternation problem asks why perceptual experience alternates between the two stimuli rather than simply staying with one stimulus or the other. That is, if the participant first sees a house followed by a face, why do these two images continue to alternate? Again, in Bayesian terms, what it is about the nature of the hypothesis of either *face* or *house*, and of the process of revising such a hypothesis, that results in the phenomenon's characteristic alternation?

2.3 The solution

The epistemological account explains the explananda as follows. *Individual vs. aggregate selection* occurs because *face and house* combined has a much lower prior than do either face or house: it is a priori improbable that what is being seen is really a *house-face*, and interactions with the environment that could have induced a prior for such a hypothesis are unlikely. Thus, as long as the low prior offsets the likelihood advantage for *face* and *house* over *face* or *house*, *face* and *house* will not be selected over either *face* or *house*.

Second, *individual vs. individual* selection occurs because, assuming the contents of the stimuli are independent, *face* and *house* explain the evidence equally well, even though they each are unable to account for a large part of it – in other words, each explains the evidence equally poorly, and so they

are roughly equally likely. Given equal likelihood, the perceptual inference will tend to depend on the prior probability of the hypotheses. If, for some reason, say, *face* has a higher prior than does *house*, then *face* will be selected for perceptual dominance.

Finally, *alternation* occurs because the hypothesis for either *face* or *house* only explains half of the stimuli. As a result, *face*, say, results in a strong prediction error signaling *house*, and so on. Since no single hypothesis combines a high prior and high likelihood, alternation between the two hypotheses results.

3. A challenge for the view

Hierarchical predictive coding has faced a number of objections. Some have criticized efforts to present predictive coding as a 'grand unifying theory of the mind,' (Anderson & Chemero, 2013; Ransom, Fazelpour, & Moll, 2017), notably on the grounds that it fails to account for the nature of motivation and desire (Huebner, 2012; Colombo & Wright, 2017, Klein 2018, Klein *forthcoming*). Others, in contrast, have offered more specific arguments that the predictive coding account fails to provide a satisfactory explanation of binocular rivalry in particular (Gershman, Vul, & Tenenbaum, 2009; Gershman, Vul, & Tenenbaum, 2012; Rescorla, 2015). The challenge I present here extends both of these lines of reasoning. Specifically, I argue that since hierarchical predictive coding struggles to explain the nature of reward, the epistemological solution to binocular rivalry will also struggle to account for the role of reward in perceptual dominance, both at onset and in duration

The role of reward in binocular rivalry can be examined in three general paradigms conducted across a number of experiments (Balcetis, Dunning, & Granot, 2012; Wilbertz, van Slooten, & Sterzer, 2014; Marx & Einhauser, 2015; though see also Wilbertz, Kamenade, Schmack, & Sterzer, 2017 for less conclusive results). The first paradigm pertains to rewarded stimuli; the

second to rewarded percepts; and the third to punished percepts. All three of these can be taken as pertaining to individual vs. individual selection. Let's look at each of these in turn.

3.1 Rewarded stimuli

The first paradigm involves three experiments using a rewarded stimulus and focused on onset dominance (Balcetis, Dunning, & Granot, 2012).

In these experiments, participants were trained to associate letters and numbers with, say, positive or negative point values, respectively (these associations were varied randomly in the actual experiments). Participants were also taught to use glasses with a blue lens in one eye and a red lens in the other, such that the stimulus were in fact overlain over one another, but only allowing one eye to see one stimulus at a time, resulting in binocular rivalry, and to report what they saw (see Figure 1).

Over a series of trials, participants were then presented with numbers and letters simultaneously under binocular rivalry and asked to report which stimulus they perceived first. Based on their responses, they were awarded points using the aforementioned reward structure – say, positive points for letters and negative points for numbers – with their total positive point values earning them raffle tickets for a monetary prize.

Broadly, the results of this study suggest that the stimulus associated with reward results in increased perceptual dominance of that stimulus. More specifically, they found that when letters were the rewarded stimulus, they were perceptually dominant at onset; when numbers were rewarded, they were perceptually dominant at onset; and, importantly, this effect only occurred when the participant himself or herself was rewarded. The effect disappeared when the reward points went to a disliked other (Table 1).

Notably, Balcetis *et al.* were careful to control for four confounds. It is useful to understand them in detail, since they isolate the role of reward in perceptual dominance. First, the authors controlled for salience and exposure by using neutral, equally common stimuli, i.e., numbers and letters. Previous studies had used domestic animals and sea creatures, and it was plausible that the former may have been more familiar than the latter to most participants. They also controlled for implicit learning by asking participants to report the dominance percept at onset, i.e., to report which stimulus they saw first, and so limiting how much they could learn over the course of the perceptual experience.

The authors controlled for response selection, i.e., motivated responding to maximize rewards, by adding a dot to each image that lay on the surface of one image but not the other. For example, in Figure 1, the dot is lying on the number 5 but not on the letter G. Participants were asked whether the dot was on or off the reported stimulus, and since there was no predetermined association between stimulus type and dot placement, participants could only correctly describe the location of the dot if accurately reported the percept which they actually saw. This feature of the experiment prevented participants from falsely reporting that they saw the rewarded stimulus first or for a longer duration.

Finally, the authors also controlled for the possibility of task availability. This might have occurred as a result of either reading the instructions, and so priming the participants to perceive letters-or, conversely- because the task involved rewards, such that participants might be primed to perceive numbers. Although the random distribution of both types should have controlled for this, the experimenters also shifted the reward structure over the course of the experiment. In one block, rewards were to go to the participants; in another, rewards were to go to a disliked other. The principle behind this structure was that if task availability played a role in what stimulus was perceived, this effect should be present across both reward structures. By contrast, if it was reward

driving perceptual dominance, the effect should be seen in the first type of structure, rewarding the participant, but not in the second, reward the disliked other.

And indeed, this latter result is in fact what the experimenters found. We thus have our first piece of evidence that reward plays a role in perceptual dominance.

3.2 Rewarded percepts

Despite the dot task controlling for response selectivity, the role that learning the associations between the letters, numbers, reward and costs plays remains open to an alternative interpretation, namely, that participants have had a chance to learn about the experiment's reward structure, and so are responding in a way that maximizes their monetary reward, rather than faithfully reporting what they see. Consequently, the second paradigm aims to control for the role of learning by examining the role of rewarded *percepts* in perceptual dominance. That is, rather than introducing the participants to the reward structure in advance, this paradigm simply asks participants to report which percept they experience and introduces latent rather than explicit rewards.

Developed by Wilbertz, van Slooten, & Sterzer (2014), in this experiment, participants were first presented with red and blue rotating grating stimuli rather than with numbers and letters (Figure 2). As above, the experiment controlled for reporting bias by asking participants to indicate the location of a target stimulus (a black dot) located in one of four positions on the grating. Unlike in the first paradigm, however, the participants did not know anything about the reward structure in advance. Instead, they trained participants to report what they saw in binocular rivalry, and then gave them the following instructions: "From now on, you will sometimes hear the sound of a falling coin during one of the colors (red or blue) and this means that € 0.10 have been added to [...] your balance. Your task is still to respond to every target you see, just as before'" (2014, p. 3). Because it

would be impossible to measure perceptual dominance at onset in this structure, experimenters instead measured perceptual dominance over duration.

The results support the role of reward in perceptual dominance. The authors found that the rewarded percept resulted in increased perceptual dominance over the duration of that percept (see Figure 3). Analogous results were reported by Marx and Einhauser (2015), who found that rewarded percepts resulted in increased perceptual dominance. Here, Marx and Einhauser described the effects of reward as qualitatively similar to the effects of attention on perceptual dominance (p. 8).

3.3 Punished percepts

In the previous paradigm, the rewarded percept resulted in the increased perceptual dominance over the duration of that percept. The third paradigm performs the same experiment, but rather than latently rewarding a given percept, the experimenter introduced small punishments. Specifically, the experimenters trained participants to report what they saw in binocular rivalry, and then gave them the following instructions: "From now on, you will sometimes hear the sound of a falling coin during one of the colors (red or blue) and this means that € 0.10 have been [...] subtracted from your balance. Your task is still to respond to every target you see, just as before'" (2014, p. 3).

Notably, in this paradigm, the punished percept resulted in the perceptual dominance of the *non-punished* percept (see Figure 3). As above, similar effects were also observed by Marx and Einhauser (2015), who reported that "non-punishing showed similar effects as reward, suggesting that the observed effects are specific to positive value and not a mere consequence of general stimulus relevance" (pp 8-9).

3.4 Precision: a possible explanation?

Hierarchical predictive coding will *prima facie* struggle to accommodate the role of reward in binocular rivalry, because the framework strictly reduces perception to prediction error minimization, and so implicitly denies that reward plays any discrete or useful explanatory role.

However, proponents of the framework could potentially argue that reward is reducible to *precision* and try to accommodate the foregoing results that way. Since this possibility is not explicitly discussed in the literature, I briefly sketch this potential line of argument here.[4] However, I then show why even such a potential response will fail to provide a convincing explanation of reward in binocular rivalry.

Various predictive coding accounts of reward can be reconstructed as follows (; Friston, Samothrakis, & Montague, 2012; Hohwy, 2013;; Clark, 2015, 2019; see also Feldman & Friston, 2010 for a discussion of the relation between precision and attention):[5]

1. Rewarded states are expected states;

2. High reward states are then expected with high precision;

3. Therefore, reward can be explained by precision.

If reward could be explained by precision, then hierarchical predictive coding could explain phenomena involving reward without positing any additional theoretical machinery, including the role of reward in binocular rivalry.

However, premise 1 is problematic. As Klein (2018, *forthcoming*) has pointed out, there is little reason to think that an organism should expect rewarded states. After all, the world can be a fairly terrible place: we should expect the environment to present us rewards and punishments alike, and

---

[4] Thanks to Friston (personal correspondence, February 2018), Schwartenbeck (personal correspondence, February-March 2018), and Hohwy (personal correspondence, February 2019) for helpful discussions of this possibility.
[5] This specific reconstruction is indebted to Hohwy (personal correspondence, February 2019): It it [sic] common to consider 'reward' as the absence of prediction error, under active inference. This means that rewarding states are expected states. This makes for an easy transition from predictive processing type views to views that include reward. High reward states are then expected with high precision."

our survival depends as much on anticipating the latter as it does on anticipating the former. Moreover, as Gershman and Daw (2012, p. 10) point out, expectation and value come apart in cases where rare circumstances are unusually good or bad. For example, if a wolf eats a deer, not because it is rewarding but because this is what its evolutionary history leads it to expect it will do, then it should only continue to eat the same amount, even in the face of a substantial caloric windfall.[6]

Proponents of predictive coding present two responses to this objection. The first response appeals to what are sometimes called *deep expectations* (Friston, 2013; Hohwy, 2013; Seth, 2014). Clark (2019, pp. 5-6) characterizes deep expectations as follows:

> Predictions are made on the basis of a generative model, and the generative model that we (considered as whole embodied organisms) instantiate will have been shaped by both evolution and lifetime learning so as to be one that ensures we are deeply disposed to predict, with high action-entraining precision, the kinds of sensory state that help to keep us alive and viable. Among such deep-set predictions we will find, for example, ones that mandate keeping key features of the bodily plant within tolerable limits.

On this view, organisms have evolved in such a way that they expect to find themselves in homeostatic equilibrium. When these expectations are violated, they adjust via active inference to return to it.

One problem with the deep expectations view is that it is very difficult to make these expectations more precise; the resulting predictions are either empirically inadequate or computationally intractable (Klein, 2018). Another problem is that, as explanations go, it simply kicks the explanatory can down the road. Presumably, these organisms' ancestors must have learned what kinds of things to expect; but how did they do so if, they too, were guided exclusively by prediction error minimization? Appealing to deep expectations thus fails to provide a substantive answer to the question of how such expectations of reward could have come about. Clark acknowledges this point, noting (2019, p. 6),

---

[6] See also Ransom *et al.* (2017) on this issue.

But, of course, such deep-set predictions get us only so far. At some point, the PP- theorist needs to accommodate the ordinary shifting webs of (as we would ordinarily say) desire—the ebbs and flows of intention that sometimes lead us to play the piano, then to work on a paper, then to order a Chinese rather than Indian takeaway, watch a certain movie, and so on.

Enter the second response to the objection. The second response appeals to what Clark more technically calls "shifting webs of precision assignments" which, together with multiple time-scales of precision, combine to produce the ins and outs of motivated behavior (2019, p. 6).

These webs reverse the traditional direction of causation in motivation. On a traditional approach, for an organism to desire p is for that organism to be disposed to act so as to bring about p, i.e., to be disposed *towards p* (e.g., see Smith, 1994). By contrast, webs of precision produce a certain disposition or action, i.e., an agent is disposed *from* the webs of precision. They are *drivers* of action from behind, as it were. For example, Clark (2019, p. 6) explains, "As both our inner states (hunger, thirst, etc.) and outer contexts ebb and flow, some predictions enjoy increased precision, becoming positioned to drive immediate actions, while others remain in the background, awaiting the right opportunity to arise." Or again, he writes,

Right now, for example, it is my high-precision prediction that I am exploring Klein's argument that is selecting my actions—both at the level of looking up various papers to check my claims, and then making specific key-strokes (cashing out precise proprioceptive predictions) on my computer (2019, p. 6).

The problem with this explanation, of course, is that it can only describe motivated behavior in hindsight. As Clark himself recognizes (2019, p. 6), the skeptic will inevitably ask "why a given agent predicts the very things she does, with their various weightings. Perhaps she chooses to order tofu rather than chicken for the take-away. Why did her lifetime learning position the tofu prediction so as to trump her colleague's suggestion of chicken?" He concludes that "The PP-mechanism itself offers no concrete story here."

This is deeply dissatisfying. It leaves us with little reason to think of rewarded states as expected states, other than perhaps by stipulation. But if we reject the first premise of the

reconstruction regarding reward and precision, then we ought to reject the general analysis of reward *as* precision.

Where does this leave us with respect to an epistemological explanation of the role of reward in binocular rivalry? As noted at the outset of this section, the epistemological approach *prima facie* struggles to accommodate the role of reward in binocular rivalry because it reduces perception to prediction error minimization, and so does not appear to have the theoretical resources to explain the foregoing experimental results.[7]

Taking it one step further, namely, to try and anticipate what the proponent of the epistemological explanation might say, I then sketched a potential line of argument based on the reduction of reward to precision. However, multiple version of this line of argument were found lacking. And if *this* is true, I content, then we ought to reject even this potential response from epistemological hierarchical predictive coding to explain the role of reward in binocular rivalry.

4. A modified approach

If the foregoing analysis is correct, then even a potential response from the epistemological version of the hierarchical predictive coding framework struggles to provide a satisfactory account of the phenomenon of reward, and so struggles to explain the specific effect of reward on perceptual dominance in binocular rivalry. By extension, and contra Hohwy et al. (2008), it follows, I argue, that the hierarchical predictive coding account cannot explain oneaspect of binocular rivalry, and so fails as a comprehensive explanation of the phenomenon.

---

[7] Another way to put the same idea is that, *prima facie,* the epistemological approach does not *predict* that reward will play a role in binocular rivalry. By contrast, an alternative theory, such as a reward-based theory of visual fixation and attention, does predict such a role (Hayhoe and Ballard 2005).

Nonetheless, the challenge that emerges from the role of reward in binocular rivalry does not mean that all is lost for the proponents of predictive processing in general. On the contrary, those seeking to defend PP have a different theoretical version of the view available to them if they want to account for reward. That is, they can deploy a modified version of Bayesian model averaging.[8]

In this section, I will show that this modified version may be able to account for the role of reward in binocular rivalry by recasting it as a special form of optimism bias.[9] Notably, my aim is not to defend this approach, but rather to show that, independently of whether it can account for the role of reward, it is fundamentally at odds with the favored, strictly epistemic picture of perception committed to by most proponents of predictive coding, as outlined at the outset of the paper.

4.1 Integrating perception and action

On the hierarchical predictive coding view, an organism *alternates* between perceptual and active inference (Hohwy 2013, p. 91). [10] That is, the organism makes inferences about perception and action using the same Bayesian principles. Nonetheless, the organism calculates these inferences independently of one another.On the modified Bayesian averaging approach, by contrast, an organism's inferences about perception and action are *integrated* in a way that may help explain the role of reward in binocular rivalry (Friston, Schwartenbeck, FitzGerald, Moutoussis, Behrens, &

---

[8] Thanks to [xxxx] and [xxxx] for helpful discussions of this issue. Thanks also to an anonymous reviewer for encouraging me to see the two versions as part of a single, overarching framework.

[9] The authors of the modified approach called it a "Variational Bayes Approach." This sometimes causes confusion, as variational methods are separable from active inference. I thus use the more minimal "modified approach."

[10] Specifically, Hohwy writes of the system, "We should therefore aim to alternate between perceptual and active inference. This alternation of inferential activity seems to me a very fundamental element of who we are and what we do. Getting the weighting of these inferential processes right is crucial to us: if the bound on surprise is not minimized enough by perceptual inference, then action suffers. If we persist with minimizing the bound for too long before we act, then we become inactive and end up spending too much time in states that are too surprising in the long run. If we persist with active inference for too long without pausing to revisit perceptual inference, then inaccuracy mounts and action becomes inefficient. If we react too soon to mounting prediction error during active inference then we get lost in overly complex and detailed models of the world" (2013, 91).

Dolan, 2013; Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2016; see also Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2017). Let's look at the modified version in more detail before seeing how this version of the explanation might work.

The central feature of the Bayesian averaging version of the approach is that it recasts traditional questions about goal-directed behavior, decision-making, and agency in terms of prior beliefs about what it would be best to do. More specifically, it recasts beliefs about external *actions* into beliefs about internal *policies*.[11] Notably, much like perceptual states, these policies are *inferred*.

Several further features follow from this. First, because policies are inferred, they are themselves associated with *precision* or *confidence*. Here, precision or confidence refers to the obtainability of a goal under a given policy (Friston et al. 2013, p. 11). For example, if the policy of drinking a cup of coffee is conducive to the goal of attaining alertness, then that policy is associated with a high measure of precision or confidence. Conversely, if the policy of drinking a cup of coffee is at odds with the goal of getting to sleep, then that policy is associated with a low measure of precision or confidence.

Second, Friston et al. (2013, p. 2) argue that "because policy optimization is absorbed into the more general problem of inferring hidden states of the world, [i.e., perceptual states,] inferences about policies depend upon inferences about hidden states and *vice versa*." That is, we can expect an organism's inferences about what it will *perceive* to depend on what it will do or, more specifically in this case, to depend on its *inferences* about what it will do, and vice versa. Hence, on the modified approach, its inferences about perception and 'action' are fundamentally *integrated* in a way that was not true on the hierarchical predictive coding view.

---

[11] Another way of describing the transformation is from *action states to control states;* I stick to policies for the sake of consistency.

Following from the foregoing two features, third, inferences about *perceptual states* are informed not only by inferences about policies, but also by the precision or confidence measures associated with those policies. That is, an organism may make an inference about what it is likely to be seeing based not only on inferences about what it is doing, but also based on how likely it is that what it is doing will enable it to achieve its goal. For instance, a brushtail possum may infer that it is likely to be perceiving a eucalyptus leaf based not only on its inference that it is climbing a tree, but also based on a relatively high precision measure associated with that strategy of climbing a tree, which is likely to help it reach food. Conversely, it may infer that it is not likely to be perceiving a eucalyptus leaf based not only on its inference that it is clambering around the base of a tree rather than climbing it, but also on the relatively low precision measure associated with clambering as a means of reaching food. This means that the framework incorporates a bias between what one infers it would be good to do - a policy associated with high precision - and what one infers about what one is actually perceiving. Since the possum's policy of climbing is associated with a high measure of precision, the critter also becomes more likely to perceive the very thing that it wants to see, namely, the eucalyptus leaf. The authors describe this bias as a kind of "optimism bias" (Sharot, Guitart-Masip, Korn, Chowdhury, & Dolan, 2012).

How *much* more likely is the possum to see the leaf? In other words, how strongly optimistic is the bias on this view? The authors only specify the issue by saying that the status quo is roughly 'optimal.' They characterize the concern in terms of precision, observing (Friston et al., 2013, pp. 10-11):

> One of the key insights, afforded by the [modified view], is that precision has to be optimized. So what would happen if (estimated) precision was too high or too low? If precision was zero, then perception would be unbiased and represent a veridical representation of worldly states. However, there would be a failure of action selection in the sense that the value of all choices would be the same. One might plausibly associate this with the pathophysiology of Parkinson's disease—that involves a loss of dopaminergic cells and a poverty of action selection. Conversely, if precision was too high, precise choices are made but there would be a predisposition to false perceptual inference—through the augmentation

of optimism bias. This might be a metaphor for the positive symptoms of schizophrenia, putatively associated with hyper-dopaminergic states (Fletcher and Frith, 2009). In short, there is an optimal precision for any context and the expected precision has to be evaluated carefully on the basis of current beliefs about the state of the world.

Hence, on this version of the framework, degree of bias is context sensitive and, in healthy organisms, optimizes a trade-off between action selection and veridical perception – but without systematically misrepresenting the world in overly optimistic terms.

4.2 The modified explanation of the role of reward in binocular rivalry

With these assumptions in place, the modified version of the view can explain the role of reward in perceptual dominance in binocular rivalry. Specifically, on this version, the role of reward can be understood as a specific instance of the foregoing phenomenon of optimism bias.

Recall the paradigms using the red and blue gratings in Wilburtz et al.'s (2014), pertaining to individual vs. individual selection. Supposed reporting red percepts is associated with a reward. As in the possum climbing example, a participant may infer that she is likely to be perceiving a red grating based not only on her inference that she is participating in a binocular rivalry task, but also based on a relatively high precision measure associated with the policy of reporting this grating, given that it is very conducive to her goal of receiving a monetary reward. She is thus slightly more likely to see the red rather than the blue grating.

An analogous argument can be made for the effect of punished percepts. Suppose in this case that reporting a blue grating is associated with a small cost. As in the possum clambering case, a participant may infer that she is less likely to be perceiving a blue grating based not only on her inference that she is participating in a binocular rivalry task, but also based on the very low precision measure associated with the policy of perceiving this grating and her goal of receiving a reward. Rather, the best policy remains seeing a red percept, and the probability of the red percept will

remain slightly higher than seeing its blue counterpart. Consequently, she is more likely to perceptually infer rewarded percepts, as is suggested by the findings described in Section 5.1. And overall, the modified version can explain the role of reward in perceptual dominance in a way that Hohwy et al.'s (2008) epistemological account of binocular rivalry could not.

5. Problem solved?

We can set aside for now the discussion of whether we accept the modified solution to the role of reward in perceptual dominance or, indeed, if we accept the broader, modified version of the framework as applied to the problem of inferring policies. Let's assume that the modified version of the approach has it right. What are we to make of this modified solution to the phenomenon of binocular rivalry? And more broadly, what are we to make of the modified approach to the underdetermination problem?

The modified solution comes at the cost of abandoning a fully *epistemic* characterization of the underdetermination problem. Recall that on an epistemic approach to the problem, perceptual systems are thought to produce veridical representations of distal stimuli for the purposes of knowledge as well as action. Recall also that many proponents of hierarchical predictive coding are quite explicitly committed to this approach, suggesting that the function of perception is to "get the world right" (Hohwy 2013, p. 2) and capture "what's going on out there" (Friston 2018, p. 1019).

My analysis of the role of reward in binocular rivalry suggests that they can't have it both ways. Hohwy *et al.*'s (2008) epistemological solution to binocular rivalry fails to account for the role of reward in perceptual dominance, and potential appeals to precision struggle to get the job done. But the modified version of the view integrates perception and policies in such a way that perception can no longer be described as being only in the business of telling us about the world as it really is.

Instead, the modified version of the approach is profoundly *pragmatic* in nature. Not only does it not characterize perceptual systems as generating veridical representations of distal stimuli for the purposes of knowledge, it rejects the commitment to exclusive veridicality, and characterizes perceptual systems as generating *biased* representations of distal stimuli for the purposes of action.

If this is right, however, then we are no longer in the business of understanding perception as consisting entirely of beliefs. Rather, we must recognize the importance of what folk psychological frameworks call desires and other approaches call reward (Sutton and Barto 2018). This is an interesting result and, in my view, by no means a disqualifying feature of the modified account. On the contrary, I take it to be a positive outcome of the preceding analysis. Nonetheless, if we want to accept the modified approach to binocular rivalry, we must also be prepared to take up this fuller picture and consider its implications for our understanding of perception.

6. Conclusion

Proponents of hierarchical predictive processing are committed to an epistemic approach to perception, according to which perceptual systems produce veridical representations of distal stimuli for the purposes of knowledge as well as action. However, I have argued that a prominent, epistemological version of this framework fails to account for the role of reward in binocular rivalry. I have further sought to show that a modified version of the general approach may be able to account for the role of reward in binocular rivalry, but that it is for its part fundamentally at odds with the epistemic commitments favored by its proponents.

References

Anderson, M. L., & Chemero, T. (2013). The problem with brain GUTs: Conflation of different senses of 'prediction' threatens metaphysical disaster. *Behavioral and Brain Sciences*, *36*(3), 204.

Balcetis, E., Dunning, D., & Granot, Y. (2012). Subjective value determines initial dominance in binocular rivalry. *Journal of Experimental Social Psychology*, *48*(1), 122-129.

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

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

Clark, A. (2019): Beyond Desire? Agency, Choice, and the Predictive Mind. *Australasian Journal of Philosophy.* DOI: 10.1080/00048402.2019.1602661.

Colombo, M., & Wright, C. (2017). Explanatory pluralism: An unrewarding prediction error for free energy theorists. *Brain and Cognition*, *112*, 3-12.

Coren, S. (1986). An Efferent Component in the Visual Perception of Direction and Extent. *Psychological Review*, 93(4), 391–410.

Feldman, H., & Friston, K. (2010). Attention, Uncertainty, and Free-Energy. *Frontiers in Human Neuroscience* 4(215), 1–23.

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

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

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions: Biological Sciences*, 369(1456), 815–836.

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

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

Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2013). The anatomy of choice: active inference and agency. *Frontiers in human neuroscience*, *7*, 598.

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68, 862-879.

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

Gershman, S., Vul, E., & Tenenbaum, J. B. (2009). Perceptual multistability as Markov chain Monte Carlo inference. In *Advances in Neural Information Processing Systems* (pp. 611-619).

Gershman, S. J., Vul, E., & Tenenbaum, J. B. (2012). Multistability and perceptual inference. *Neural computation*, *24*(1), 1-24.

Gershman, S. J., & Daw, N. D. (2012). Perception, action and utility: The tangled skein. *Principles of brain dynamics: Global state interactions*, 293-312.

Hayhoe, M., & Ballard, D. (2005). Eye movements in natural behavior. *Trends in cognitive sciences*, 9(4), 188-194.

Helmholtz, H. von. (1867). Handbuch der Physiologischen Optik. Leipzig: Voss.

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

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

Huebner, B. (2012). Surprisal and valuation in the predictive brain. *Frontiers in psychology*, *3*, 415.

Jacob, P. (2015). Action-based accounts of perception. In M. Matthen (Ed.), *The Oxford handbook of the philosophy of perception*. Oxford: Oxford University Press, 217-236.

Klein, C. (2018). What do predictive coders want?. *Synthese*, 195(6), 2541-2557.

Klein, C. (*Forthcoming*). A Humean challenge to predictive coding. In *The Philosophy and Science of Predictive Processing* eds Steven Gouveia, Dina Mendonça, and Manuel Curado. Bloomsbury Press.

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

Macpherson, F. (2009). Perception, Philosophical Perspectives. *The Oxford Companion to Consciousness,* T. Bayne, A. Cleeremans and P. Wilken (eds), Oxford University Press, 502-508.

Marx, S., & Einhäuser, W. (2015). Reward modulates perception in binocular rivalry. *Journal of vision*, *15*(1), 11-11.

Metzinger, T., & Wiese, W. (2017). Philosophy and predictive processing. *Frankfurt am Main: MIND Group.*

Noë, A. (2004). *Action in perception.* MIT press.

Ransom, M., Fazelpour, S., & Mole, C. (2017). Attention in the predictive mind. *Consciousness and cognition*, *47*, 99-112.

Rescorla, M. (2015). Bayesian perceptual psychology. In M. Matthen (Ed.), *The Oxford handbook of the philosophy of perception.* Oxford: Oxford University Press, 694-716.

Sharot, T., Guitart-Masip, M., Korn, C. W., Chowdhury, R., & Dolan, R. J. (2012). How dopamine enhances an optimism bias in humans. *Current Biology*, *22*(16), 1477-1481.

Siegel, S., & Silins, N. (2015). The epistemology of perception. In M. Matthen (Ed.), *The Oxford handbook of the philosophy of perception.* Oxford: Oxford University Press, 781-811.

Smith, M. (1994). *The Moral Problem.* Wiley.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction.* MIT press.
Vishton, P.M., N.J. Stephens, L.A. Nelson, S.E. Morra, K.L. Brunick, and J.A. Stevens. (2007). Planning to Reach for an Object Changes How the Reacher Perceives It. *Psychological Science*, 18, 713–719.

Wilbertz, G., van Slooten, J., & Sterzer, P. (2014). Reinforcement of perceptual inference: Reward and punishment alter conscious visual perception during binocular rivalry. *Frontiers in psychology*, *5*, 1377.

Wilbertz, G., van Kemenade, B. M., Schmack, K., & Sterzer, P. (2017). fMRI-based decoding of reward effects in binocular rivalry. *Neuroscience of consciousness*, *2017*(1), nix013.