# Consciousness, attention, and the motivation-affect system

**Tom Cochrane**

Forthcoming in the *Journal of Consciousness Studies*

**Abstract**

It is an important feature of creatures like us that our various motivations compete for control over our behaviour, including mental behaviour such as imagining and attending. In large part, this competition is adjudicated by the stimulation of affect- the intrinsically pleasant or unpleasant aspects of experience. In this paper I argue that the motivation-affect system controls a sub-type of attention called 'alerting attention' to bring various goals and stimuli to consciousness and thereby prioritise those contents for action. This view allows me to flesh out the global workspace theory of consciousness, as well as some of the phenomenal characteristics of conscious experience.

---

Strong links are widely noted between attention and consciousness. For instance, the prominent global workspace theory of consciousness (e.g. Baars 2002) makes attention necessary for consciousness and Jesse Prinz and Filipe de Brigard (2010) have even claimed that attention is sufficient for consciousness. Yet various observations indicate that attention and consciousness can be disassociated. In this essay I propose that we can reconcile these disassociations by focusing on 'alerting attention'. This is a sub-type of attention that is controlled by the motivation-affect system.

The motivation-affect system is the name I give to the collection of our affective states and drives. It includes every process that has a motivational aspect, ranging from our fundamental homeostatic systems, to pains, pleasures, fatigue, nausea, itchiness, emotions, moods, and even character traits. At every moment of our lives, these various motivational processes jostle against each other, competing for control over the organism's activities, including its mental activities. Understanding how this competition works is, I think, crucial for understanding consciousness.

Working within the framework of the global workspace theory of consciousness, I will develop an account of how the motivational-affect system directs an organism's alerting attention. I will explain how attention to some content has to achieve motivational priority before it can trigger the wider broadcast of that content to other content-consuming mental capacities such as verbal report and memory. When this happens, according to the global workspace theory, that content becomes conscious. Connecting consciousness to motivational priority also has the strong attraction of explaining the function of consciousness. Finally, connecting consciousness to motivational priority helps to make sense of the subjective character of consciousness; how it fluctuates in its structural, temporal, and spatial characteristics depending on our motivational engagement. Overall, this paper will show how the motivation-affect system can fill out a rich and plausible account of consciousness.

**1. Prima facie connections**

My strategy in this paper is to start with some very broad connections between consciousness and the motivation-affect system, which I will explain and elaborate as the paper progresses. To begin, we can point to both a conceptual or intuitive connection and a clinical connection.

First the conceptual/intuitive connection: Our focus here is on phenomenal consciousness, where the core characteristic of phenomenally conscious states is that there is something that it is like, subjectively speaking, to undergo them. One of the main arguments against physicalist-functionalist accounts of phenomenal consciousness is that the functional characteristics of a state like seeing red could conceivably be maintained while its intrinsic phenomenal character (its qualia) could be entirely absent, or else spectrum-inverted such that what it's like to see red is replaced by what it's like to see green. If such differences are even possible, it indicates that phenomenal character and functional properties cannot be identical, since it is not possible for something to be distinct from itself.

It may be noted that whatever intuitive appeal this argument has is at least partly due to focusing on the functional characteristics of colour perception; behaviours like verbally reporting colour, or distinguishing red from blue. If instead, one were to consider the functional characteristics of pain or intense despair—such as the motivation to scream and shout or curl up into a ball—it is far less plausible that these could remain unchanged while their phenomenal character was either entirely absent or inverted. For instance, it is harder to imagine that the intrinsic feeling of intense joy or pleasure could just as easily replace the intrinsic feeling of pain or despair.

Presenting this argument, Craig DeLancey claims that "our concept of the experience of an emotion is both richly phenomenal and functional; and the particular phenomenal properties—the emotional qualia—are specific to their function" (1996: 496).[1]

It seems that the intrinsic conscious character of states like pain, pleasure, despair and joy have a definite role to play. They help to supply the urgency or *oomph* that make us respond with alacrity to whatever we discern as responsible.[2] Specifically I want to point to the *affect* of these mental states, which is their intrinsically pleasant or unpleasant aspect; the painfulness of pain, or the pleasantness of pleasure. This aspect is distinct from sensory components such as noticing bodily damage or the taste of ice-cream. It is also arguably the aspect that various affective states have in common. That is, it is intrinsic pleasantness that is shared by the otherwise varying experiences of drinking a cool beer, winning a game of tennis, feeling pride in one's children, or enjoying an interesting book. While on the negative side it is intrinsic unpleasantness that is shared by physical pain, grief, nausea, or being bored by a dull book.[3]

Overall, the intuitive argument points out that pleasurable and painful affect is a paradigm of conscious experience. Affect is sufficient for phenomenal consciousness, and its phenomenal characteristics link up directly with its functional characteristics.

Let us now turn to a clinical case which suggests that affective states and consciousness stand or fall together. Akinetic mutism is a condition caused by damage in frontal areas of the brain in which patients maintain wakefulness and some behavioural responsiveness, but have a marked deficit in consciousness. There has been some discussion regarding whether akinetic patients completely lack consciousness (e.g. Klein 2017; Davies & Levy 2016). The condition can be more or less severe such that less severe cases may retain fleeting experiences. Furthermore, some patients retrospectively report their blankness, which might be interpreted

[1] See also Deonna & Teroni (2020: 119) for a similar argument that an emotional zombie could not understand our evaluative practices.
[2] Jaak Panksepp makes a somewhat similar point (2003: 122).
[3] There is some controversy concerning what unifies our various pleasant and painful experiences. Pleasure in particular seems heterogeneous (see e.g. Bramble 2013 for a review). However, some unifying factor is demanded to explain why, in general, we avoid painful experiences and embrace pleasurable ones. In fact, certain accounts of felt affect analyse it precisely as a motivational imperative (e.g. Martínez 2011; Barlassina & Hayward 2019). This is not my approach, however. I defend an 'evaluativist' account according to which felt affect is essentially a higher-order representation of the goodness or badness of the sensory input. This higher-order representation then reinforces the motivational tendency of the underlying affective system. See my (2018: Ch.2) for further elaboration and defence, as well as Carruthers (forthcoming).

either as an attempt to describe the loss of consciousness or as a phenomenally conscious experience of blankness. However, as Davies and Levy note, patients only ever report absences of various kinds and offer no positive characterisations that we could only attribute to phenomenal experience. Thus I judge that at least in severe cases, akinetic mutism involves a genuine loss of consciousness.

For our purposes, what is notable about such patients is that their lack of consciousness accompanies a lack of self-driven (endogenous) agency and affective states, including emotion, pain, pleasure, hunger, and thirst. For instance, Antonio Damasio describes the case of 'L', a woman who remained in an akinetic state for several months:

> She might make a normal movement with arm and hand, for instance, to pull her bed covers, but in general her limbs were in repose. Together, body and face never expressed any emotion of any kind, background, primary, or secondary, although there were plenty of inducers offered, day to day, in the attempts at focused conversations or just plain bedside banter of physicians, nurses, medical students, friends, and relatives… Months later, as she emerged from this state of narrowed existence and gradually began to answer some questions, she would clarify the enigma of her state of mind... It appeared that there had not been much mind at all, and nothing that would resemble core consciousness, let alone extended consciousness. The passivity in her face and body was the appropriate reflection of her lack of mental animation. She had no recall of any particular experience during her long period of silence; she had never felt fear; had never been anxious; had never wished to communicate. (Damasio 1999: 102-103)

Affect is not present in individuals suffering akinetic mutism. More than this, motivational states generally seem to be lacking. Damasio also describes cases of absence seizures (1999: 96-100) and late stage Alzheimer's disease (1999: 103-104) in which loss of consciousness is similarly accompanied by loss of emotion and affect. Thus there is evidence that we do not see affective states where consciousness is absent. That is, these conditions indicate that affective states might be necessary for consciousness.

How seriously should we take the preceding intuitive and clinical connections between affective states and consciousness? It may be argued that affective states cannot be sufficient

for consciousness because they can sometimes be unconscious. For instance, in an experiment by Winkielman and Berridge (2015) participants were briefly presented pictures of happy and angry faces which were backwards masked such that they could not be consciously discerned. Following this, Winkielman and Berridge observed differences in the participants' preference judgements and consumption behaviour, indicating at least a subtle emotional impact. Yet even when participants were explicitly directed to attend to how they were feeling, they reported no emotional experience. Meanwhile, Prinz (2004: 203) reports an experiment by Fischman and Foltin (1992) in which cocaine addicts were hooked up to two intravenous drips; one infused with cocaine, the other with saline. At very low concentrations of cocaine, the addicts reported no discernible difference between the drips, yet statistically showed a definite preference for the cocaine drip. Their preference for the cocaine drip might accordingly be interpreted as an unconscious pleasure.

My response to these sorts of experiments is to agree that an affective state can be unconscious, but only so long as it is mild enough not to win priority over other ongoing affective states. In general, affective states should primarily be conceived as representational states that automatically trigger responses to approach or avoid their targets.[4] Affective states can vary considerably in their representational and behavioural complexity. Thus hunger may represent the simple lack of nutrition where fear represents that a harm is approaching (a temporal contrast), and jealousy represents that someone else possesses a good to which one is entitled (a social contrast). Yet even complex representations and responses are achievable in the absence of consciousness.

Where I draw a stronger link to consciousness is with affect, as described above. It seems clear that the generation of affect is sufficient for consciousness (though a specific source of affect can get crowded out by another source, as I discuss below). All affective states are capable of stimulating affect (hence the name) but there are conditions on this.

## 2. The generation of affect

To understand why affect is sufficient for consciousness, we need to discuss how affect is generated. As I have described elsewhere (Cochrane 2018) affect is both a felt intrinsic quality

[4] In my 2018 book I analyse these states as 'valent representations'; representations that automatically trigger some valent (i.e. positive or negative) response.

and a system for managing our various affective states. The causal trigger of affect is the success or failure of affective states to deliver their functions. A negatively valent affective system like the one responsible for monitoring bodily damage will only trigger affect (i.e. the painfulness of pain) if it *fails* to regulate signals of bodily damage.[5] Merely registering bodily damage by itself is not enough. We know this because individuals with pain asymbolia can register bodily damage without experiencing unpleasant affect (Grahek 2007). These individuals even preserve some (though not all) reflexive responses to bodily damage, and of course their bodies still repair themselves. Thus I have argued that painful affect must be a higher-order representation, dysfunctional in asymbolia, which combines a representation of the organism's attempts to avoid or repair damage with a representation that this damage is still present; i.e. a kind of failure. I moreover predict that the overall intensity of affect will be proportional to the detected intensity of bodily damage multiplied by the detected intensity of the organism's response. If either level is zero, then affect will not be triggered.

Symmetrically, a positively valent affective system like the emotion of excitement will trigger pleasurable affect when the creature both represents the presence of something good and responds in such a way that the (anticipated) presence of this good is increased. Pleasurable affect is thus a representation of *success* in acquiring something to which one is positively drawn. Again, I predict that its intensity will be proportional to the detected level of the good multiplied by the detected level of the response. These symmetrical models are illustrated side by side below.

[5] The system is negatively valent in the sense that it seeks to avoid or reduce the presence of its target (in this case bodily damage).

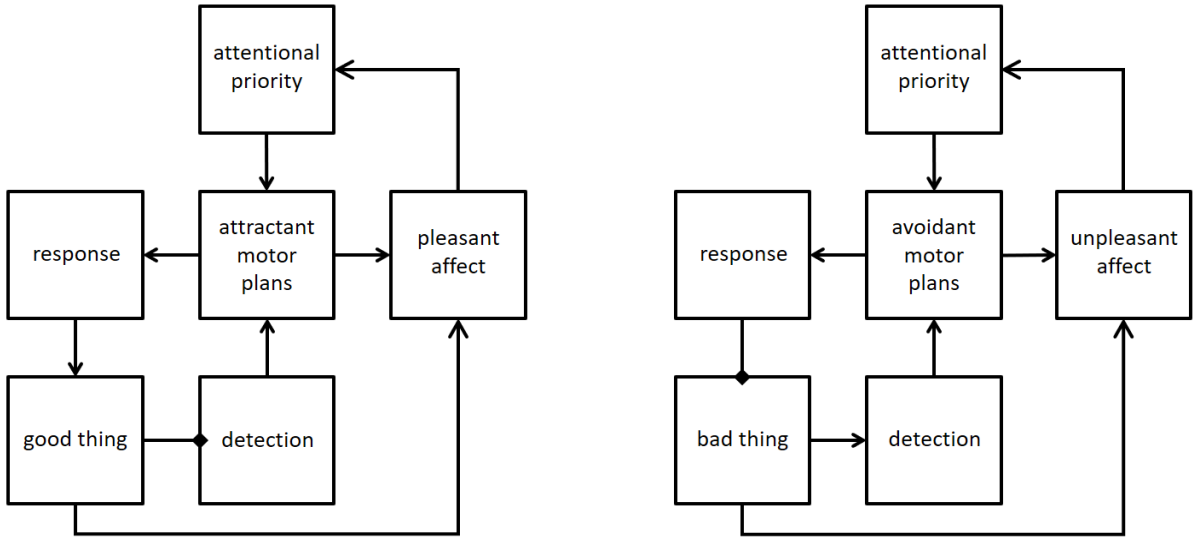


**Functional models of pleasurable and painful affect (taken from Cochrane 2018: 61)**

The other crucial feature of this model of affect is that its output is the boosting of attention towards the underlying positively or negatively valent system. The above diagrams link the arrow from attention to the response plans in the sense that the call to respond is intensified. However, so far as conscious experience is concerned, the model allows that attention brings the sensory objects of those response plans into sharper focus (see Cochrane 2018: 60 for discussion).[6]

As I conceive it, the essential regulative role of affect is to manage the conflict between different affective systems that simultaneously call for the organism to respond in some way. So long as a response system draws on resources that no other response system is drawing on (be it a certain physiological mechanism, or metabolic energies) then there is no need for management (and no affect). But once these systems compete for resources it is necessary to have some way to decide which one takes priority. The intensity of failure or success represented by affect is precisely how I suppose that an organism ascribes relative importance to the functioning of a mechanism. Thus failure or success to any degree will call for attention, but this can be overridden by a stronger call from another system. This is why it is important that affect is both distinct and unified across all kinds of pleasant or unpleasant states. Affect is a way of collating urgency across different motivational systems.

[6] I am persuaded by Carruthers (2019: 57-61) that, in general, the dorsal processing pathways that directly guide motor behaviour operate non-consciously, while conscious experience applies to ventral perceptual pathways.

For example, consider the phenomenon of battlefield analgesia in which a soldier may be severely injured during the heat of battle and yet not feel the pain from that injury for several minutes (Melzack et al. 1982). As soon as injury occurs, the system for monitoring that body part will certainly be trying and failing to repair the injury, and accordingly calling for attentional priority; to get the entire person to take care of that body part. Yet other affective systems within the soldier are simultaneously in operation- i.e. fear of being killed, or anger towards the enemy. These emotional systems also demand immediate whole-person responses and attention towards the sensory qualities most relevant to guiding those responses. Thus it is plausible that, in extreme circumstances, such emotional systems can more intensely arouse affect and win attentional priority. I also anticipate that the scope of attention will be tied to the intensity of affect. The more intense the affect, the more it will fill the subject's attentional capacity, crowding out calls from other affective states.

So far then, my claim is that not all affective states will arouse affect, and of those that do arouse affect, not all capture attention. An alternative position is to simply equate affect with a certain kind of conscious feeling, and so if a negatively valent system's failure does not manage to capture attention, and is not experienced, then it does not count as arousing affect. Other than intuitions about the consciousness of affect, I don't see any way to decide which alternative is preferable. But let us at least have on the table the simpler view that affect may be (functionally speaking) the attended representation of a failing negatively valent system or the attended representation of a succeeding positively valent system. Either way, attention will be dominated by whichever affect-generating system is failing or succeeding the most.

I can now offer a refinement of my earlier claim that affect is sufficient for consciousness: Affect-winning affective states dominate or prioritise consciousness. They need to, in order to spur us to act on the most urgent successes or failures of their affective systems. Specifically, we are conscious of whatever content feeds into a winning affective state and is available to be broadcast across the mind. This content will be experienced with the gloss of pleasurable or painful affect, which supplies motivational and attentive urgency. Thus we will tend to experience both the call and the response of the affective state, though absent self-conscious reflection, we are unlikely to distinguish the two.

**3. Motivational relevance**

It is already fairly obvious that we will experience whatever is most intensely pleasurable or painful, though I think it is helpful to clarify its functional characteristics. The more difficult issue is to pin down what is necessary for consciousness. Below I will endorse the claim that a kind of attention is necessary for consciousness. But before that, let us consider what is necessary for attention. Intuitively, affect does not seem necessary. When I am calmly looking out the window, my experience does not seem to have any noticeable pleasurable or painful character. It should be noted that affect can be subtle. For instance, sensory scenes are subject to aesthetic satisfactions or dissatisfactions. It is also easy to note how comfortable or content one's body feels at any given moment. Thus, affect may be more pervasive than is commonly recognized. Yet rather than insisting on the presence of subtle affect in all cases of attention, I believe we should consider in more detail the motivational side of the motivation-affect system.

What should be recognized is that our motivational systems are not only concerned with moving our bodies around in the service of our various goals, but also with bringing forth representational contents. Curiosity and surprise are obvious examples of motivational states that seek representational content. But all motivational states require input and guidance from representational content; they all *call for content*. The way that affect grabs attention merely intensifies a call that is already there.[7]

To give an example, the presence of another animal is information of major significance for several different motivational systems, such as the predator avoidance system, and the mate-seeking system. Amongst the undergrowth, the initial appearance of a face (i.e. two eyes and a mouth) may not be sufficient to clarify which of these motivational systems should be activated. Thus the first response of an animal is likely to be a re-orientation to gather more information, disambiguate the stimulus, and thereby trigger the appropriate response.

Besides such overt bodily movements aimed at gathering information, we also bring forth representational content in our imaginations, understood here in a very broad sense to encompass visualising, planning, empathic projects, inner speech, and the voluntary aspects of

[7] In Cochrane (2023) I offer an extended description of the mechanism of calling for content. I specify that basic or direct mental actions are ones in which a currently prioritised goal brings content to working memory. Where attention is aroused exogenously by stimuli that is highly relevant to a background or dormant drive, this does not count as mental agency, though it is still goal relevant and still falls under alerting attention (more on which below).

episodic memory. In all cases, there is some goal that we aim to satisfy. We might explicitly articulate this goal to ourselves—'where shall I go on holiday this year?' or it can be driven by motivational systems such as hunger (one starts daydreaming about food). Note in particular, that such mental acts need involve no stimulation of affect; no special success or failure of one's motivational projects. Rather such mental acts will just be what is motivationally *most relevant* for the agent at that time.

How is motivational relevance graded? Again, assuming that there is really no affect in operation (so no distinctively unpleasant hunger, thirst, discomfort and so on), then I suppose it to be largely opportunistic or stimulated by cues. Motivational systems are always seeking their dedicated inputs and so long as we are awake we gather sensory data that will be relevant to at least one of them. For instance, if I walk into my living room and see my piano, then my long-standing goal to play the piano is stimulated to some a degree. If nothing else is particularly pressing at that moment, I am liable to sit down and start playing. At the very least, I'll attend more closely to the piano and a brief image of playing may come to mind. Similarly, if I'm watching television and a French film comes on, I may be reminded of my goal to holiday in France, and start to imagine what this would involve. It will take interference from my other goals to resist the switch to such trains of thought.

We also structure our environments and our daily routines to provide these sorts of motivational cues. My office is filled with cues that remind me of my various work tasks; my to-do list, the books I'm supposed to read; the email icon on my computer desktop. Similarly, I have established habits such that as soon as I enter my office, I will switch on my computer and check my email, which will quickly inform me of various things I'm supposed to be doing.

In fact, I claim that the moment we wake up, we are motivationally engaged in something or other, even if the goal is just to identify what's going on over there. We may wake up because of some affective state such as the pressure to urinate, or we may wake up because the sound of an alarm stimulates the goal to be aware of the auditory environment. Either way, wakefulness involves constantly shifting one's attention towards anything that is important or novel. Sometimes the level of novelty is sufficient to stimulate affect, as in the emotion of surprise, but much of the time, we are simply attentive to changes.

In accordance with this, predictive coding theories of the mind emphasise that the perceptual-cognitive system constantly seeks the reduction of predictive error (e.g. Hohwy 2013; Clark 2013). We anticipate that our perceptual inputs will be a certain way, and if our expectations are not matched by reality, the mismatch or error signal is fed up the perceptual hierarchy. The greater the signal of error, the more intensely attention is called for. Predictive coding thus serves a motivation to keep track of what is going on around us, whereby change or novelty is motivationally relevant as a failure of one's predictive goals. I emphasise, however, that mere novelty will only be attended when there is no more active goal calling for content. Much more frequently, we only attend to what is novel relative to a focused motivational system.

Overall, I claim that attention is captured by whatever representational contents are both available for wider processing and most relevant to a motivational system. Motivational systems call for contents. Affect is always relevant to a degree and thus attention will be dominated by the inputs to whichever affective state is arousing (the most) affect. But where affect is not stimulated, relevance is graded by whichever motivational system happens to be most active and we will attend to the contents serving that motivation. At a minimum, we are motivated to attend to anything novel within our perceptual fields, in so far as this serves a goal to identify what is happening around us. However, this motivation can be overridden by various affective drives. Stimuli that has greater affective salience can crowd out our awareness of the background environment. The affective state of fatigue can also override our awareness goals.

Thus it seems to me that relevance is a core feature of attention, and this is dictated by the competition between one's motivational systems. How can this be experimentally verified? It should in principle be possible to objectively establish which of a person's motivational systems is most active. I also expect there to be bandwidth considerations which we can tie to empirically established observations of the scope of working memory.[8] If the scope of working memory can be firmly established, then one way to test my claim would be to see if filling the working memory with motivationally important tasks inhibits conscious awareness of stimuli that are not relevant to those tasks. More generally, I predict that we will not be conscious of

[8] Cowan (2001) reports this to be 4 distinct items of content at any given moment (where an item can include chunked data, enabling quite complex information to be captured at once). For discussion see also Carruthers (2019: 82-83).

stimuli (even those lying squarely within our perceptual fields) that are either not novel or not relevant for some motivationally prioritised goal.[9]

I can now make my more general claim about phenomenal consciousness. I propose that phenomenal consciousness is the product of the motivation-affect system calling for representational content. Our consciousness awareness is filled by both the affective states that call for content (e.g. hunger, fear, painful affect) and the content that gets called for (e.g. food, danger, bodily damage). These affective calls make use of attention to draw on perceptually available contents, or contents stored in memory, and then broadcast this content to various content-consuming faculties of the mind, making the content available for subjective report, reasoning, and a flexible range of responses.[10] The limits of working memory constrain how many items of sensory data can be broadcast at any given moment. Only the most relevant calls for content and their inputs will be broadcast, where relevance is graded by the degree to which a motivational system is active, and especially whether it is generating any affect. Thus we might put it that 'relevance-winning content' is necessary and sufficient for consciousness.

**4. Assessing the claim**

To the extent that I appeal to the mechanism of attention triggering the broadcast of information, the view I am defending operates within the framework of the global workspace theory of consciousness. The main distinctive feature of my approach is to specify the trigger or 'ignition' point of broadcast. Furthermore, I do not regard 'attention' conceived merely as a mechanism of selection for further processing or for guiding behaviour to be sufficient to trigger consciousness. I constrain the type of attention that concerns me with respect to both a determination of relevance and the availability of content to be broadcast to other content consuming faculties.

[9] This will be challenging to empirically verify, since as soon as one asks a subject to report on the contents of their awareness, those contents become motivationally relevant to the reporting task, thus trivially satisfying the conditions I have outlined. Instead, it would make more sense to check on what subjects spontaneously report. If subjects report contents that are neither novel nor motivationally salient, this would significantly undermine my claims. Alternatively, we could test whether subjects fail to notice non-novel, motivationally inert items being surreptitiously removed from the perceptually available field, although removing stimuli risks making background stimuli novel.

[10] I also should note here that 'broadcast' does not mean that the content is literally shuffled over to some other part of the brain. Rather what is enabled is the mutual interaction or reinforcement between the area(s) processing the relevant content and other content consuming faculties (cf. Prinz 2012: 101-102; 321-322). See also Cochrane (2021) for a defence of this view.

To test and further elaborate this model, my task is now is to see if it can accommodate empirical observations made about the relationship between attention and consciousness. In particular, I want to know if it can handle the supposed disassociations between attention and consciousness.

First a quick review of the positive correlations between attention and consciousness: The main evidence that tends to get mentioned here is inattentional blindness and related phenomena. In the classic demonstration of inattentional blindness (Simons & Chabris 1999) participants are engaged in an attentionally demanding task such as counting how many basketball passes are made between players wearing white. This task causes viewers to fail to notice a man dressed in a gorilla suit who walks into the middle of the viewing screen, waves, and then walks off. Inattention seems to make participants literally unconscious of this stimulus. Similarly, in attentional blink (e.g. Raymond et al. 1992) when two stimuli are presented, one immediately after the other, attention to the first inhibits consciousness of the second. Third, in change blindness experiments (e.g. Simons & Levin 1997) participants fail to notice quite major changes in a visual scene when their attention is directed elsewhere, even though these features can be present to the participant's foveal region. Yet words and shapes that capture attention (such as the subject's name or a smiley face) can be reported. As de Brigard and Prinz put it "when attention is captured invisible stimuli become visible" (2010: 54).

Note in particular, the operation of relevance in the first and third of these examples. The task to count basketball passes makes that stimuli relevant for subjects. Participants are motivated to detect basketball passes and so they are constantly calling for this information. Change blindness similarly involves consciousness to stimuli coming or going with the emotional or personal relevance of that stimuli. The phenomenon of attentional blink does not seem to have much to do with relevance (other than the participant being motivated to pay attention to the screen). The unconsciousness of the second stimuli is merely a limitation of how fast attention can shift.

Let us turn now to the disassociations between attention and consciousness. Evidence against the sufficiency of attention for consciousness generally involves experiments in which an individual's performance on a task is facilitated by the selective processing of information, even though the subject is not consciously aware of this information. Since attention is

conceived as the selective uptake of information for further processing, these are taken to be cases of attention without consciousness.

One such piece of evidence comes from studies with the blindsight patient GY (Kentridge et al. 2004). GY can more quickly and accurately report whether or not an item has been presented to his blind field if a cue is first presented- again within the blind field. This indicates that this area of screen has been selected for increased processing prior to the presentation of the item.[11]

The global workspace view is already well-placed to accommodate this apparent case of attention without consciousness. If, due to cortical damage, content is not available for global broadcast, then it does not matter whether attention is given to it or not, it can never become conscious. No special appeal to motivation or affect is required to make sense of this case. It is, however, the reason why this theory of consciousness specifies that content must be available for global broadcasting.[12]

More evidence is where participants can pick up on subtle informational properties independent of conscious awareness.[13] In one experiment (Mulckhuyse et al. 2007) a grey circle was flashed up on a screen 16ms earlier than two others—too brief a window to be consciously discerned. Yet participants were then more quickly able to respond to following targets displayed in the location of the earlier stimulus. In a similar, yet more cognitively challenging task (Zhao et al. 2013) sequences of symbols were simultaneously displayed at 4 different locations on a screen. One of these sequences was more regular or repetitive than the other 3. Participants were not aware of this and yet again performance at that location was facilitated.

Another relevant experiment makes use of binocular rivalry. This is where incompatible images are presented one to each eye, and instead of experiencing a blend of the two images, participants fluctuate between consciousness of one and consciousness of the other. In an experiment by Jiang et al (2006) the task given to participants was to judge in what direction a Gabor patch was tilted. This was a challenging task since the Gabor patch was only presented

[11] It is, incidentally, worth bearing in mind that GY has 'Type 2' blindsight. He sometimes reports experiences correlated with his successful performance on visual tasks. These experiences however need not be comparable to normal visual experience (Kentridge 2015). Kentridge suggests they are often only 'a feeling of knowing'.
[12] Bayne also points out that blind field content cannot be used for the selection of goals, even if it can enable the execution of intentions (2013: 171). This supports the claim that the activation of our goal dispositions by perceptual cues is a conscious event.
[13] The following three experiments are all discussed in detail by Mole (2020).

for 100ms, and only tilted by 1 degree. Participants were better able to judge the tilt if the image was shown on a side in which previously an erotic image was displayed, indicating that attention was already being given to that side. Yet this erotic image was never consciously perceived because a dynamically changing, high contrast image was shown to the other eye, which causes that image to dominate in binocular rivalry scenarios.

Overall these experiments, alongside the blindsight case, are quite compelling evidence that selective information processing can occur in the absence of conscious awareness. One response would be to dispute whether such cases count as attention. It is quite plausible to interpret the above experiments as cases in which selective processing enables participants to *then* more easily attend to the tasks that follow. However, there is not much point engaging in such semantic disputes over what counts as attention proper, which would mostly generate confusion over what we are actually talking about. Better, in my opinion, to specify a sub-type of attention that is aligned with consciousness.[14]

Indeed, Taylor (2018) convincingly argues that the experimental paradigms of psychologists and neuroscientists make use of three distinct notions of attention. One of these is the selective processing of information. The other two are 'alerting' in which individuals maintain a high sensitivity to incoming stimuli, and 'executive attention', which is a mechanism for managing conflicts between thoughts and responses.

It is not very plausible to align consciousness with executive attention. Executive attention is the sort of process that occurs when we resist distractions from our goals, for instance, if we try to concentrate on a lecture instead of daydreaming about getting ice cream. This level of control seems sufficient for consciousness and is moreover highly motivationally related. It seems to require that our goals are consciously salient. Yet it is too demanding for this to apply to all cases of phenomenal consciousness, many of which may involve no resistance of distractors or sophisticated self-awareness.

[14] De Brigard and Prinz suggest that these sorts of experiments show 'spatial attention' but this is distinct from the 'attentional modulation of perceptual representations' which they regard as necessary and sufficient for consciousness. Taylor (2013) argues that this definition ties attention too closely to working memory, itself defined in terms of subjective report, for it to be falsifiable. Since we can only determine consciousness via subjective report, any time an experiment purports to show attention without consciousness, the mere lack of subjective report means that attention hasn't delivered the representation to working memory, which means that de Brigard and Prinz can exclude it as attention properly so-called.

Alerting however, is much more congenial to our interests. It is goal driven, and the affective state driving our goal may well be obvious to us, yet it need not involve the reflective consciousness of one's goals. It also seems able to incorporate both cases in which the stimuli is not yet present but one calls it to consciousness (i.e. by searching around or completing an imaginative task) and cases in which a stimuli that is currently present is sufficiently motivationally relevant that one desires to keep it present. That is, alerting can incorporate both endogenous and exogenously stimulated attention.[15]

So what I propose is that we specify the alerting sub-type of attention as responsible for triggering consciousness, given the understanding that this involves the interaction of (perhaps already selectively processed) information with one's motivational systems. I suppose that the cue and the motivational call reinforce each other in a reciprocal loop. In a burst of positive feedback, calls from other motivational systems are overwhelmed, and a system-wide broadcast is triggered, thereby bringing the cue and the motivational call to consciousness.

So what should this view say about the above experimental evidence? It seems clear enough that whether a circle is presented slightly earlier than others, or whether a sequence of symbols is subtly more regular than others, is not motivationally relevant enough to trigger a definite goal to focus on a certain portion of a screen. Selective processing attention merely facilitates alerting attention for the following task. The third reported experiment in which an erotic picture is displayed does seem more motivationally relevant. Yet this is an experiment in which consciousness is dominated by an alternate image within the binocular rivalry paradigm. My response then is that the dynamic high-contrast image must be of *greater* motivational relevance for the individual than the briefly presented erotic image, even if that erotic image would most likely become conscious given enough time.

**5. More on binocular rivalry**

More generally, the binocular rivalry paradigm is worth examining in some detail to see if we can give an explanation for why consciousness fluctuates between the two images in the ways that have been observed. One of the best explanations for this phenomenon comes from the predictive processing theory (Hohwy et al. 2008). Indeed, it is one of strongest lines of evidence

[15] This can also be linked with the 'saliency network' identified by Corbetta and colleagues (Corbetta & Shulman 2002; Corbetta et al. 2008) which is supposed to continually monitor unconscious percepts and memories and check them against one's current goals.

offered in support of this general approach to the mind (cf. Clark 2013: 184-185). The basic claim of predictive processing is that the image we are conscious of is a high-level hypothesis or 'prior' about the external world. Since our hypotheses are gained from experiences of ordinary objects, certain combinations, such as a house and a face in the same spatial location, are not hypothesised and become incompatible. Thus in rivalry set-ups, one high-level hypothesis crowds out competitors, so that we are generally only conscious of one image at a time. However, bottom-up error signals from perceptual processes responsive to the alternate image are fed up the perceptual hierarchy. Eventually, this error signal is strong enough to trigger the replacement of the high-level hypothesis. Yet now bottom-up error from the first image will start to accumulate. Therefore, so long as the two images cannot be reconciled, we can expect our experience to fluctuate.

Meanwhile, the lengths of time we are conscious of the two images in the rivalry scenario are not always equal. As noted above, one of the main ways in which one image can come to dominate is if it displays higher visual contrast. The predictive processing model can account for this by claiming that higher contrast generates a stronger error signal, thereby making that image more likely to overwhelm error from the other image. Another cause of domination is if the background context is congruent with one of the images (Hohwy et al. 2008: 697). Again, the predictive processing model can suppose that congruency influences the priors that we bring to the scenario. Related to this, the individual's deliberate attention can influence which image comes first to consciousness (though overall fluctuation is not voluntary). Indeed, Paffen & Alias (2011) propose that where attention can make a difference to binocular rivalry it is due to attention affecting the stimulus contrast.

Yet there are other causes of domination in binocular rivalry that are less easy to fit with the predictive processing model. Emotional imagery (both positive and negative) has been found to dominate over neutral images. For instance, the image of a fearful face will more often be seen first and for longer (e.g. Alpers & Gerdes 2007). Similarly, Müller, Gerdes and Alpers (2022) found that images of spiders will dominate for spider-phobic individuals, particularly those who prefer to cope with that fear by avoiding spiders. Meanwhile, on the positive side, stimuli that are associated with a monetary reward (or negatively associated with a monetary cost) will lead those stimuli to dominate. This occurs even when the stimuli are otherwise neutral (such as letters and numbers) and have only temporarily been linked to a reward (Balcetis et al. 2012; Wilbertz et al., 2014; Marx & Einhäuser 2015).

On the basis of some of these experiments, Haas (2021) argues that the purely epistemic account of predictive processing—that perceptual systems produce the representations that they do purely for the sake of getting the world right or explaining what's going on out there—must be revised. If epistemic accuracy was all that was happening we should not expect motivationally salient images to dominate so effectively in rivalry scenarios.

Haas' own proposal is that predictive processing should include an 'optimism bias' in the generation of priors, such that the creature is more apt to be aware of things that offer a reward. An optimism bias is, however, in tension with the dominance of fearful faces, or spiders for people who strongly desire to avoid spiders. Instead I propose that the predictive processing model of binocular rivalry should be incorporated into a more general account of motivationally-relevant attention. That is, we see things that feed into affective states, whether they be positively or negatively valent. In section 3, I argued that the motivation to be aware of one's surroundings aligns predictive processing with motivational views of attention, and moreover that perceptual error can be a source of affect, as in the emotion of surprise (cf. Cochrane 2018: 80-82). We are generally motivated to get things right, even when this brings unwelcome news. Yet this motivation operates within a larger context of avoiding things are bad for us and approaching things that are good for us.

The alerting sub-type of attention should make us more likely to attend to rewarding or scary images in rivalry scenarios. However, it is not the case that the rewarding or scary image is always conscious (despite being available). Thus it is necessary to include these motivational cues alongside the motivation to get things right. Both must be balanced, proportional to their importance. Note in particular that it is plausible that when only still images are presented (and contextually, one knows that one is safe) individuals can quickly habituate such that after a short period of time they do not generate affect so intensely and other motivationally relevant information can come to the fore.

**6. Function and phenomenology**

Overall, I believe that linking consciousness specifically to motivationally-driven alerting attention can accommodate the various experimental evidence. Of course, as a version of the global workspace theory, the view also aligns with the considerable experimental evidence in favour of that view which I have not reviewed here (I recommend instead Carruthers 2019:

Ch.4 and Ch.5 for a detailed review and comparison with other theories). My view predicts that so long as information is potentially available for global broadcast, and there is a motivational system in operation that calls more intensely for that information than other motivational systems, then the individual will be conscious of that information. Accordingly, the view will be falsified if we can find cases where the individual has a clear motivational priority for some perceptually available information (i.e. it is more closely associated with a reward or cost than anything else available) and yet, when probed, they report either no consciousness of the affective state or no consciousness of information relevant to that affective state. Alternatively, my view will be falsified if subjects are conscious of stimuli that lack motivational relevance (as described at the end of section 3). Given that I've allowed that merely knowing what is going on has some motivational relevance for individuals, it will be challenging to design decisive experimental paradigms, especially since subjects are motivated to perform experimental tasks. Yet I believe it should be possible.

Meanwhile the view has the definite attraction that it helps to make sense of why we are conscious. We know from various experimental evidence that information can be processed unconsciously and still successfully guide behaviour. And though I'm willing to bet that consciousness is widely spread throughout the animal kingdom, there are still likely to be species that are capable of quite complex cognition and yet unconscious. So what does consciousness add? The answer that we can draw from the above discussion is that consciousness tracks motivational relevance, enabling the flexible or spontaneous redirection of our activities (including mental activities) where opportunities arise, and keeping fixed the priority of those states arousing affect. This is a view that has been proposed before. Here, for instance, is Jaak Panksepp:

> I assume that the most primitive function of consciousness is to facilitate adaptive response selection from alternative courses of action: It allows organisms to cope with complex environmental situations in which several behavioral alternatives are competing, with comparable urgency, for a common output channel in the brain. (1982: 451)

Panksepp is known for putting emotions and affect at the centre of the mind, and in this regard I am in full agreement.

Further support for the claim that the motivation-affect system guides and dominates consciousness can be found by noting some of the broad phenomenal characteristics of experience. For instance, we can expect that if consciousness is always dictated by motivational relevance, then relevance will structure the field of consciousness. Thus, the things that matter to us stand out more clearly and centrally, with less important contents receding into the background.

Of particular note is that the way that our motivational engagement can affect the basic spatial and temporal characteristics of conscious experience. For instance, time can seem to slow down when we are under extreme or life-threatening stress (Hancock and Weaver 2005). To some extent, this is likely to be explained by changes in bodily arousal, such as a raised heart rate, or heart rate variability. Yet given that these situations involve a major sense of actual or anticipated failure, we can expect the call of the relevant motivational systems for perceptual and cognitive contents to be very intense. The heightened frequency of such calls may very well help to explain the time distortion effect, i.e. one is experiencing more calls for content per second than usual.

Given that positive affect also generates calls for content, we should also expect the time-slowing effect to accompany positive stimuli, and this is indeed what we find (see Ciu et al. 2022 for a recent meta-review). It is interesting however that positively-valenced stimuli generate less of a time-slowing effect than negatively valenced stimuli. Ciu and colleagues hypothesise that this is due to a human bias to attend more towards negative events (2022: 14). It may be harder to achieve a comparable positive intensity with the sorts of pictures and sounds that researchers are using for stimuli. We should also distinguish between expected positive events and unexpected ones. It is much more plausible that one's sense of time slows down during the experience of unexpectedly receiving a prize in front of a large crowd of people than when, say, going for a pleasant walk in the countryside.

Another important phenomenon to consider is the experience of flow, in which it is commonly reported that time seems to move faster (e.g. Csíkszentmihályi 1990, also Shepherd 2022 for a recent review of studies). Flow states are fairly rare experiences in which we become highly absorbed in a task such as playing sport or performing a piece of music. Csíkszentmihályi claims that flow states occur when the difficulty of a task is perfectly balanced against one's capacity to perform it. As such, we can expect that calls for attention generated by performance

errors are eliminated. Simultaneously, if one's intentions to act are being immediately satisfied then one is less likely to need to hold one's intentions in mind until they are satisfied. Moreover, while successful performance will certainly generate positive affect, this positive affect should concentrate one's attention on the task one is already fixated upon. This is in marked contrast to the unexpected prize example given above, in which positive affect would shift one's attention around the crowd or the need to go up on stage. Thus in flow cases the calls for attention that may affect our sense of temporal duration are plausibly reduced and thereby we can explain why time may seem to move faster.

Another common report in experiences of flow is that the individual's sense of distinctness from their task or the tools they are using is reduced. Thus flow affects the spatial characteristics of conscious experience in addition to its temporal characteristics. Elsewhere I have described how the distinct sense of self can be generated by the sense that the world is resistant to one's goals and that one's intentions are distinct from reality (Cochrane 2017). Where one's intentions are instead immediately matched by performance outcome, the awareness of the outcome is likely to occlude the awareness of one's intention. One's intentions don't disappear, they just become less obvious. Thus again, we have a phenomenon in which the salience of one's calls for content could help to explain the intrinsic characteristics of conscious experience.

## 7. Conclusion

Overall, I have claimed that phenomenal consciousness is the product of the motivation-affect system calling for representational content. This explains the process by which we become conscious, the function of consciousness, and some of its intrinsic phenomenal characteristics. I have outlined how the motivation-affect system commands the alerting sub-type of attention, which calls for perceptual and cognitive contents, both present and non-present, to be widely broadcast throughout the mind. Though a certain kind of attention can occur independently of consciousness, this is merely the selective processing of information, and not the motivated use of information that we find in alerting attention. Without this motivated use, I doubt that we would be conscious at all.[16]

[16] Thanks to the referees from this journal for their comments on this paper. I would also like to thank my student Galo Cheung, discussions with whom helped me to develop the view I have presented here. Though in various respects our views on consciousness differed, Galo and I agreed that conscious states are ones that we motivationally call upon.

**References**

Alpers, G. W. & Gerdes, A. (2007). Here is looking at you: emotional faces predominate in binocular rivalry. *Emotion*, 7(3), 495.

Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, 6(1), 47-52.

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

Barlassina, L. & Hayward, M. K. (2019). More of me! Less of me!: Reflexive imperativism about affective phenomenal character. *Mind*, 128(512), 1013-1044.

Bayne, T. (2013). Agency as a marker of consciousness. In A. Clark, J. Kiverstein, & T. Vierkant (Eds.). *Decomposing the will*. Oxford: Oxford University Press, 160-181.

Bramble, B. (2013). The distinctive feeling theory of pleasure. *Philosophical Studies*, 162(2), 201-217.

Carruthers, P. (2019). *Human and animal minds: The consciousness questions laid to rest.* Oxford: Oxford University Press.

Carruthers, P. (forthcoming). On valence: imperative or representation of value? *The British Journal for the Philosophy of Science*.

Cui, X., Tian, Y., Zhang, L., Chen, Y., Bai, Y., Li, D., Liu, J., Gable, P. & Yin, H. (2022). The role of valence, arousal, stimulus type, and temporal paradigm in the effect of emotion on time perception: A meta-analysis. *Psychonomic Bulletin & Review*, https://doi.org/10.3758/s13423-022-02148-3

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

Cochrane, T. (2017). Group Flow. In M. Lesaffre, M Leman & P. J. Maes (eds.) *The Routledge Companion of Embodied Music Interaction*, London: Routledge, 133-140.

Cochrane, T. (2018). *The Emotional Mind: A control theory of affective states*. Cambridge: Cambridge University Press.

Cochrane, T. (2021). A case of shared consciousness. *Synthese*, 199(1), 1019-1037.

Cochrane, T. (2023). Imagination, Endogenous Attention, and Mental Agency. *Phenomenology and the Cognitive Sciences* https://doi.org/10.1007/s11097-023-09909-y

Corbetta, M. & Shulman, G. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3), 201–15.

Corbetta, M., Patel, G., & Shulman, G. L. (2008). The reorienting system of the human brain: from environment to theory of mind. *Neuron*, 58(3), 306-324.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1), 87-114.

Csíkszentmihályi, M. (1990). *Flow: The Psychology of Optimal Experience*, New York: Harper and Row.

Damasio, A. R. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. London: Vintage Books.

Davies, W. & Levy, N. (2016). Persistent vegetative state, akinetic mutism, and consciousness. In Walter Sinnott-Armstrong (ed.) *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, New York: Oxford University Press, 122-34.

De Brigard, F. & Prinz, J. (2010) Attention and consciousness. *Wiley interdisciplinary reviews: Cognitive science* 1 (1), 51-59.

DeLancey, C. (1996). Emotion and the function of consciousness. *Journal of Consciousness Studies* 3(5-6): 492-499.

Fischman, M. W. & Foltin, R. W. (1992). Self-administration of cocaine by humans: A laboratory perspective. In G. R. Bock & J. Whelan (Eds.), *Cocaine: Scientific and social dimensions*, Chichester, UK: John Wiley and Sons, 165-180.

Grahek, N. (2007). *Feeling pain and being in pain*. Cambridge, MA: Bradford Books/MIT Press.

Haas, J. (2021). Can hierarchical predictive coding explain binocular rivalry? *Philosophical Psychology*, 34(3), 424-444.

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.

Jiang, Y., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender-and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences*, 103(45), 17048-52.

Kentridge, R. W., Heywood, C. A., & Weiskrantz, L. (2004). Spatial attention speeds discrimination without awareness in blindsight. *Neuropsychologia*, 42(6), 831-835.

Kentridge, R. W. (2015). What is it like to have type-2 blindsight? Drawing inferences from residual function in type-1 blindsight. *Consciousness and Cognition*, 32, 41-44.

Klein, C. (2017). Consciousness, intention, and command-following in the vegetative state. *The British Journal for the Philosophy of Science*, 68(1), 27-54.

Martínez, M. (2011). Imperative content and the painfulness of pain. *Phenomenology and the Cognitive Sciences*, 10(1), 67–90.

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

Melzack, R., Wall, P. D. & Ty, T. C. (1982). Acute pain in an emergency clinic: latency of onset and descriptor patterns related to different injuries. *Pain*, 14(1), 33-43.

Mole, C. (2020). Consciousness and Attention. In U. Kriegel (ed.) *The Oxford Handbook of the Philosophy of Consciousness*, Oxford: Oxford University Press, 498–519.

Mulckhuyse, M., Talsma, D. & Theeuwes, J. (2007). Grabbing attention without knowing: Automatic capture of attention by subliminal spatial cues. *Visual Cognition*, 15(7), 779-788.

Müller, U. W., Gerdes, A. B., & Alpers, G. W. (2022). You see what you avoid: Fear of spiders and avoidance are associated with predominance of spiders in binocular rivalry. *Journal of Anxiety Disorders*, 86, 102513.

Paffen, C. L., & Alais, D. (2011). Attentional modulation of binocular rivalry. *Frontiers in Human Neuroscience*, 5, no. 105, 1-10.

Panksepp, J. (1982). Toward a general psychobiological theory of emotions. *Behavioral and Brain sciences*, 5(3), 407-422.

Panksepp, J. (2003). Damasio's error? *Consciousness and Emotion*, 4(1), 111-134.

Prinz, J. (2004). *Gut Reactions: A perceptual theory of emotions*. Oxford: Oxford University Press.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human perception and performance*, 18(3), 849.

Shepherd, J. (2022). Flow and the dynamics of conscious thought. *Phenomenology and the Cognitive Sciences*, 21, 969-988.

Simons, D. J. & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1(7), 261-267.

Simons, D. J. & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28(9), 1059-1074.

Taylor, J. (2013). Is attention necessary and sufficient for phenomenal consciousness? *Journal of Consciousness Studies*, 20(11-12), 173-194.

Taylor, H. (2018). Attention, psychology, and pluralism. *The British Journal for the Philosophy of Science*, 69(4), 935-956.

Deonna, J. & Teroni, F. (2020). Emotional Experience: Affective Consciousness and its Role in Emotion Theory. In U. Kriegel (ed.) *Oxford Handbook of the Philosophy of Consciousness*. New York: Oxford University Press, 102-123.

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.

Winkielman, P. & Berridge, K. (2015). Unconscious emotion. *Current Directions in Psychological Science*, 13(3), 120-123.

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention is spontaneously biased toward regularities. *Psychological Science*, 24(5), 667-677.