



Perceptual Motivation for Action

Tom McClelland¹ · Marta Jorba²

Accepted: 8 December 2021 / Published online: 19 January 2022
© The Author(s) 2022, corrected publication 2022

Abstract

In this paper we focus on a kind of perceptual states that we call perceptual motivations, that is, perceptual experiences that plausibly motivate us to act, such as itching, perceptual salience and pain. Itching seems to motivate you to scratch, perceiving a stimulus as salient seems to motivate you to attend to it and feeling a pain in your hand seems to motivate actions such as withdrawing from the painful stimulus. Five main accounts of perceptual motivation are available: Descriptive, Conative, Imperativist, Reflexive and Dual Content views. Some of these have been developed in detail and others are natural extensions of existing views. However, we argue that none provides a satisfactory account of the target perceptual states. Driven by the failings of each account, we identify three conditions that any satisfactory theory of perceptual motivation must meet: Direct Motivational Force, Non-Obligatory Force and Intention-Independent Influence. Then we offer an account that attempts to fulfil these conditions: the Affordance Model. This model of perceptual motivation is premised on the claim that each of the target states involves the perception of a specific affordance for action and the triggering of Automatic Motor Initiations (AMIs). We suggest that AMIs can explain the motivational force of these perceptual states.

1 Perceptual Motivations

Actions have motivations. In ordinary cases, the mental states that motivate our actions are *desires*. The desire to eat a cake, for example, might motivate your act of eating a cake. Although perceptual states are likely to play an important role in such an action, their role is not ordinarily *motivational*. Your eating of the cake might depend on you seeing it, but the perceptual state of seeing the cake is not what drives your action. It is rather the desire for cake that carries the motivational force. In many cases seeing the cake may be the state that initiates an action, but here the perceptual state tells you how to act on a motivation that you have rather than the

✉ Marta Jorba
marta.jorba@upf.edu

¹ University of Cambridge, Cambridge, UK

² Pompeu Fabra University, Barcelona, Spain

perceptual state itself constituting a motivation. Put another way, perception tells you how things are but doesn't tell you what to do. It is tempting to think that this picture holds across the board – that *no* perceptual states are motivational. However, there are some perceptual experiences that plausibly do motivate us to act.

Consider the experience of *itching*. Being in this perceptual state seems to motivate us to scratch. Consider also our experience of *perceptual salience*. Perceiving a stimulus as salient seems to motivate you to attend to it. And consider the experience of *pain*. A pain in your hand seems to motivate actions such as withdrawing from a painful stimulus. These three perceptual experiences are quite different in character, but what they seem to share is a *motivational force*. They do not merely tell us what's going on in our body or in the world. Rather, they tell us *how to act* – they somehow point to the action required. In fact, it is hard to describe these mental states *without* mentioning the actions that they motivate you to perform. Hall goes as far as claiming that it is 'unintelligible' for someone to say that they have a terrible itch yet feel no urge to scratch (2008, p. 525). Perhaps it is similarly unintelligible to perceive something as highly salient without feeling an urge to attend to it, or to feel pain from contact with something without feeling an urge to withdraw from it. We will call these types of mental states *perceptual motivations*.¹

Perceptual motivation presents us with the challenge of explaining how (if at all) these perceptual states, with their intimate connections to action, have motivational force. This explanatory challenge raises a number of more specific questions that need to be confronted. Do the perceptual states motivate a *specific* action, such as scratching, or do they motivate any action that helps achieve some goal? Do these motivations for action qualify as *reasons* to act, or do they fall outside the space of reasons? When we respond to these motivational states, do our responses always qualify as actions or might they be non-agential movements? These questions, and others, must be answered if we are to understand how perceptual states can be motivational. This will have implications for our understanding of the specific kinds of mental state cited: itches, pains and perceptions of salience. More broadly, it will have implications for the character of perceptual experience, the place of perception in action and the nature of motivation.

Five main accounts of perceptual motivation seem to be available to account for perceptual motivation. Some of these have been developed in detail and others are natural extensions of existing views. However, we argue that none provides a satisfactory account of the target perceptual states. Each of the five accounts will be considered in Section 2. Driven by failings of each account, we will identify three conditions that any satisfactory theory of perceptual motivation must meet. In Section 3 we will offer an account that attempts to fulfil these conditions: the Affordance Model. This model of perceptual motivation is premised on the claim that

¹ We take a liberal view on what counts as *perceptual* experience which includes cases of visual/auditory perception and cases of bodily sensations such as pains and itches. While we acknowledge that more restrictive views on perception would leave proprioceptive states aside such restrictions are not relevant to our project. We are interested in how the three cases presented share important features and present a similar profile regarding their motivational force.

each of the target states involves the perception of a specific affordance for action. Itches involve perceiving an affordance to scratch, salience involves perceiving an affordance to attend and pain involves perceiving an affordance to withdraw. Once we appreciate the interface between affordance perception and action, we can understand how these perceptual states can have motivational force.

2 Five Accounts of Perceptual Motivation

2.1 Descriptive Accounts

One possible way of explaining perceptual motivation is with a Descriptive account, according to which the motivational force of a perceptual state is explained entirely by its descriptive content. Descriptive content has a mind-to-world direction of fit: it represents how things are and it is 'successful' if the world is indeed as it represents it to be. If the descriptive content does not match the world, then the representation must be adjusted to bring it into alignment with reality. Although the descriptive contents of perceptual states might resist being fully captured linguistically, we can approximate such content quite easily. The visual experience of a tomato might have content along the lines of <this object is round and red>. According to descriptive accounts, if a perceptual state has the right kind of descriptive content then it will have motivational force. The idea is that there is something about perceiving the presence of an itch, a pain or a salient perceptual stimulus that thereby motivates you to scratch, withdraw from it or turn your attention to it, respectively.

Although there are various candidate descriptive contents here, they can never motivate action in the way that perceptual motivations do. Let's say that the itch experience has the content <unpleasant sensation here>. For this state to motivate you to scratch, you would need to have additional suitable beliefs and desires. You would need to have a *desire* to alleviate unpleasant sensations and a *belief* that scratching will help achieve this desire. But what we actually find with itches is that they motivate us to scratch *independently* of our wider beliefs and desires.

An itch still motivates you to scratch even if you lack the belief that scratching will alleviate the unpleasant sensation. In fact, itches still motivate scratching even when you believe that scratching will make it worse. Similarly, itches motivate you to scratch independently of your wider desires. Itches can motivate you to scratch even in the absence of a desire to alleviate the itch. Even if you had some strange desire to be itchy (or just lacked a desire to alleviate the itch), the itch would still motivate you to scratch. Itches have what we will call *direct motivational force*. They motivate one to act, and they do so independently of one's wider beliefs and desires.

To understand what it means to have direct motivational force, consider cases where a perceptual state motivates us *indirectly*. An itch might motivate us to apply an ointment. But here the motivation depends on your wider mental states: your desire to get rid of the itch and your belief that applying the ointment is an effective means of satisfying this desire. If this belief changed – e.g. if you came to believe that you are allergic to the ointment and that it makes the itching

worse – you would no longer have the motivation to apply the ointment. Thus, in this case the motivational force of the itch is indirect because it depends on the additional belief, so it does not come directly from the itch. Similarly, if your desire changed – e.g. if you came to desire itchiness – the itch would again no longer motivate applying ointment. So on the Descriptive account, the motivational force of the itch is *indirect*: it depends on one's wider beliefs and desires. But the way that itching motivates scratching is quite different. When we have an itch we are motivated to scratch *directly*. Similarly, a pain might motivate us to continue suffering it on the basis of one's belief that suffering pain makes one stronger, for instance. Notice again that if such an additional belief changes or is no longer endorsed, then the motivation to continue having that pain also disappears. According to the descriptive view, pain's motivational force is also indirect because it depends on these additional beliefs and desires. The same applies to perceptual salience: a strong light in your visual field might indirectly motivate you to go in the direction of the strong light because you believe that something relevant is happening there. But once this belief changes or disappears, one is no longer motivated to do that particular action. Thus, direct motivational force is plausibly an essential feature of the class of perceptual motivations that we're targeting.

Direct motivational force captures the apparent *recalcitrance* of perceptual motivation. Directly motivational perceptual states are recalcitrant in that they are not responsive to reasons. No matter how much you try to reason yourself out of scratching, your itch will carry on motivating you to scratch. No matter how much you reason about the need to keep your hand on the stimulus, your pain still motivates you to withdraw from it. And no matter how much you reason about avoiding paying attention to a strong light that appears in your visual field, the perceptual stimulus will motivate you to attend to it. In contrast, where a perceptual state *indirectly* motivates you to φ , your motivation is responsive to reason. For instance, if I reason in the right way then my itch will no longer motivate me to apply ointment. This is because through reasoning you can alter the belief that ointment will diminish the itch sensation and come to believe that the ointment won't help. This reasoning influences the additional beliefs and desires that might be associated with the perceptual state, but not the motivational force of the perceptual state per se. Again similarly, if I reason in the adequate way then my pain will no longer motivate me to continue having it, or I might alter my belief that it is indeed a good idea to follow a strong flash upon perceiving it.

It is worth pausing to consider our reasons for claiming that these perceptual states have direct motivational force. One source of evidence is introspective: reflecting on the phenomenology of itches, pains and perceptual salience we find that there is something about them that motivates us directly. However, introspective disagreement is common and others might find no such reason to posit direct motivational force. But a second source of evidence is theoretical: there is a clear contrast between the case of the itch motivating scratching and the case of the itch motivating applying ointment and the *best explanation* of this difference is that only the former motivates us to act directly (and so the latter motivates us indirectly through the desire). Similar contrasts can be found among the actions that pains motivate and

the actions that perceptual salience motivates, and these are best explained in the same way.

The problem of being unable to directly motivate action generalises to any other descriptive account. Even if we replace the descriptive content <unpleasant sensation here> with some other content, any such content can only *indirectly* motivate us to act. Content that only presents how things are is not the kind of content that can tell us what to do.² One option, for instance, is to claim that the perceptual states represent things as *good to φ* . Thus the itch is perceived as good to scratch, the salient stimulus as good to attend to and the painful stimulus as good to avoid. But even here the motivational force of such states would depend on a wider desire to do good things. The lesson is that representations of how things are can only tell us what to do with outside assistance, yet perceptual motivations tell us to act without any such assistance. They motivate us to act ‘by themselves’. Indeed, if the motivational force of a perceptual state was indirect then the urge to act upon it wouldn’t be recalcitrant in the sense explained.

As such, a descriptive account cannot accommodate the direct motivational force of these perceptual states. An itch motivates scratching independently of your wider beliefs and desires about it. A salient stimulus motivates you to attend to it independently of any desire to attend to it or any belief about what attending to it would achieve. In fact, salience motivates you to attend even when our beliefs and desires push *against* us doing so. And a pain motivates you to withdraw from it independently of your wider beliefs and desires, and continues to do so even if someone convinces you that suffering pain will make you stronger. This gives us our first criterion for judging accounts of perceptual motivation:

- I) DIRECT MOTIVATIONAL FORCE: In perceptual motivation, the subject’s perceptual state motivates them to φ directly.

2.2 Conative Accounts

Conative accounts claim that the target states are constituted (in whole or in part) by desires. Unlike descriptive states desires are paradigm cases of motivational states. Desires have a world-to-mind direction of fit. A desire for cake is satisfied just in case you get cake. If the content of a desire does not match the world, then it is the world that the agent must adjust rather than the mental state. Unlike the descriptive account, advocates of the conative account cannot just build conative content into our perceptual state. We can make sense of a perceptual state describing the world as being some way, but how could a perceptual state constitute a desire for the world

² This general criticism resembles the naturalistic fallacy, or the is-ought problem, that Hume (1739/2000) described for moral reasoning: there seems to be a gap between descriptive statements (about what is the case) and normative or prescriptive ones (about what ought to be the case). Applied to our case, the fallacy consists in believing that descriptive contents of perception can prescribe how to act.

to be some way? The conative account must posit additional states that are *triggered* by the motivational perceptual state.

What is the content of these conative states? One option is to posit a desire that sets some general goal for the subject to achieve. An itch, for example, might trigger a desire to remove the itchy state. Scratching is then motivated as a means of achieving this desired end. The difficulty with this view is that it violates DIRECT MOTIVATIONAL FORCE. The desire to remove the itchy state would only motivate scratching if you also had the belief that scratching is an effective strategy for satisfying this desire. As discussed though, itching, having a pain or perceiving a salient perceptual stimulus motivate their respective specific actions independently of one having such a belief. Our motivation for scratching itches, for instance, lacks the proposed means-end structure.

A related problem for this conative approach is that it cannot capture the *specificity* of perceptual motivation. A desire for some target state is neutral on how the subject achieves that state. A desire for cake is satisfied only by the acquisition of cake but is neutral on how the cake is acquired: the desire itself does not tell us whether to walk to the bakers, drive to the supermarket or bake in the kitchen. Perceptual motivation, on the other hand, is not neutral about how to act. The itch tells us to scratch, the salient stimulus tells us to attend and the pain tells us to withdraw. The action prescribed is thus specific.

The conative account can be refined to make sense of the specificity of perceptual motivation. Some conative states target specific actions, and are satisfied only by the performance of that action. A desire to bake does not merely set some target at which our actions must aim. Instead, it motivates you *to bake*. The conative account can thus adjust its characterisation of perceptual motivation as follows: the relevant perceptual state triggers a desire *for a specific action*. Itches trigger desires to scratch the itch, salient stimuli trigger desires to attend to that stimulus and painful body parts trigger desires to withdraw from the painful stimulus.³

Something to note about this revised version of the conative account is that it takes beliefs out of the picture. An itch triggers the desire to scratch and does so independently of one's wider beliefs. Even if one believed that scratching is a bad idea, the automatic link between the itch and the specific desire would persist. The specificity of the desire means that there is no need for background beliefs to mediate between what we want and the actions we are motivated to perform. As such, the revised conative account avoids a major shortcoming of the Descriptive account.

But despite these advantages over the Descriptive account, the conative theory still has trouble dealing with the DIRECT MOTIVATIONAL FORCE condition. DIRECT MOTIVATIONAL FORCE proposes that in perceptual motivation we are motivated to act independently of our wider beliefs *and desires*. This is certainly not something the

³ Actually, the class of actions that pain motivates may be more diverse than this (Klein 2015). Besides motivating withdrawal, a pain might motivate us to protect a body part, to avoid putting weight on it, etc. But for each token pain, the motivational force seems to have the kind of specificity discussed, i.e., it motivates us to do specific things.

conative account can accommodate. It claims that itches motivate scratching *precisely because* of some triggered desire state distinct from the itch.

Why think that these perceptual states motivate independently of one's desires, though? Because it seems that having an itch, for example, would motivate one to scratch even in the absence of a desire to scratch. Even if you firmly desired *not* to scratch, the motivation to do so would persist (and similarly for pains and perceptual salience). A complication here is that itches do typically come along with desires to scratch. Even when you desire *not* to scratch because you believe it will make things worse, it is plausible that you *also* have a desire to scratch. Desires can, after all, be inconsistent with each other. However, even when an itch comes with a desire to scratch, we can still identify the real locus of motivation. It is the *itch itself* that seems to motivate us to scratch and not just the desire. Let's stipulate that itches automatically trigger a desire to scratch. Now imagine that you have some long-running trouble with itching and get in the habit of applying a soothing ointment to your itches. And let's stipulate that whenever you itch, you desire to apply the ointment, and that the triggering of this desire is automatic. So whenever you itch, you automatically have a desire to scratch *and* a desire to apply ointment. Even with these parallel desires, the *way* that the itching motivates you to scratch would be quite different to the way the itch motivates you to apply ointment. There is something about the itch experience that directly motivates scratching – the itch itself calls out to be scratched in a way that it doesn't call out to be soothed with ointment. But in this scenario the scratching desire and the ointment desire are exactly alike. So the motivational force of the itch cannot be exhausted by any correlated desire.

Here we are probing deeper into DIRECT MOTIVATIONAL FORCE. Even where our wider beliefs and desires *do* motivate us to φ , our perceptual state seems to give us some independent motivation. Itching is a motivation to scratch in its own right. The same goes for perceptual salience and pain. Even where our desires are in alignment with the motivational force of our perceptual state, the perceptual state itself gives us a direct motivation to φ . Overall, the Conative account cannot capture the distinctive way in which perceptual motivations motivate us to act.⁴

2.3 Reflexive Accounts

Reflexive accounts claim that our responses are a reflex-like response triggered by our perceptual state. Itch states trigger a scratching reflex, salient states trigger an attention reflex and pain states trigger a withdrawal reflex. This account promises to satisfy DIRECT MOTIVATIONAL FORCE straightforwardly. Reflex responses do not depend on one's wider beliefs and desires. In fact, a φ -ing reflex will persist even when your beliefs and desires motivate you *not* to φ . The Reflexive account also accommodates the specificity of perceptual motivation. Reflexes are bound to one specific action.

⁴ A version of the Conative Account could be a broader Affective Account according to which the relevant perceptual states motivate us in virtue of their affective aspect, that is, their involving emotional components. But such a view would not satisfy DIRECT MOTIVATIONAL FORCE either unless it characterises the emotional component as *part* of the perceptual state itself, which we regard as an implausible possibility.

Consider the blinking reflex: the reflex is not to do *something* to stop an incoming object going in your eye – the reflex is specifically to blink.

Although the Reflexive account appears to improve on the previous accounts, the difficulty with such an account is that it predicts that when we are in the relevant perceptual state we *always* perform the associated action. If an object comes rapidly towards your eye, your blinking response is obligatory. Yet if you feel an itch, a scratching response is *voluntary*. Itches don't have to be scratched (Hall 2008, p. 527). Similarly, salient stimuli do not always have to be attended and painful stimuli do not always have to be withdrawn from. In fact, the whole notion of being *motivated* to φ relies on the possibility of one *not* φ -ing. It would be odd to say that the doctor applying their hammer to your knee motivated you to move your leg. Rather, your automatic movement would be an unmotivated response. It shouldn't even be classified as an action. Yet scratching in response to an itch *is* an action, so the reflexive account must be getting the relationship between itching and scratching wrong. This gives us a second criterion:

- II) NON-OBLIGATORY FORCE: In perceptual motivation, φ -ing is not an obligatory response.

2.4 Imperativist Accounts

Imperativist views claim that perceptual states have imperative content such as <scratch here!>, <attend here!> or <withdraw your hand!>. Imperative states are like desires in that they have a world-to-mind direction of fit. Unlike desires though, they are always targeted at some specific action rather than some general goal. An imperative is an imperative *to* φ . Crucially, Imperativist accounts build this content into the perceptual state itself. They do not posit some distinct imperative mental state that is triggered by one's perception. Rather, the perceptual state itself commands the subject to act. And, like any perceptual state, it has its content independently of one's wider beliefs and desires.

The account promises to accommodate all three of the conditions introduced so far. It satisfies DIRECT MOTIVATIONAL FORCE because imperatives motivate us to act directly, and retain their motivational force regardless of what one might believe or desire. Your beliefs and desires might tell you that you ought to continue suffering pain but your perceptual state will say <withdraw from it!> regardless and thereby motivate you to withdraw from the pain. The account also promises to satisfy NON-OBLIGATORY FORCE. Imperatives motivate you to act but do not compel you to act. As Hall puts it: '...just as you don't have to do what another person tells you to do, so similarly you don't have to do what your bodily senses tell you to do.' (2008, p. 527) We are thus free not to obey a perceptual command to scratch.

One existing line of objection to Imperativism is that the motivational force of an imperative is contingent on one's background beliefs and desires (Bain 2011). A sergeant's command for you to do push-ups will only motivate you to act if you have a background desire to do what your sergeant commands and a background belief

that they are your sergeant. Where such background mental states are absent, one can be aware of a command without being in any way moved to obey that command. The objection to Imperativism is that the same applies to the proposed imperative content of perception: being aware of a command is one thing but being motivated to act on that command is quite another.

In our terms, this objection to Imperativism casts doubt on its satisfaction of DIRECT MOTIVATIONAL FORCE. However, Imperativists have offered some reasonable responses to this objection. Watzl (2017) defends his Imperativist view of perceptual salience by proposing that if an imperative is part of *your* perceptual state then you are motivated to act on it. Background beliefs and desires are only required in the sergeant case because the command comes from *someone else*. When the imperative content belongs to one's own perceptual state, no such mediation is required for the content to have motivational force. Think of this by analogy with descriptive perceptual content. Someone else's visual representation of rain only gives you reason to believe it is raining if certain background beliefs are in place, such as the belief that the other person doesn't generally hallucinate. But if *you* visually represent rain then your visual state warrants the belief that it is raining without the help of any background beliefs (though this warrant could of course be defeated by wider beliefs).

Klein (2015) offers a different kind of response in defence of his Imperativist account of pain. He concedes that an imperative only has motivational force if certain background conditions are met. But in cases like pain, those background conditions are met *almost all of the time*. In rare pathological cases such as pain asymbolia these conditions might be absent. But in all ordinary cases the conditions are in place for the imperative to motivate you. Whether either of these responses stands up to scrutiny is an open question. However, we argue that Imperativism faces a deeper problem that will be much harder to overcome.

Obedying a command is something you do. A command can only cause you to φ if you yourself elect to φ . There is an *intention* to φ that mediates between the command and the action. For instance, the sergeant's order only results in you doing push-ups if you comply with that order. But what we find in perceptual motivation is that the perceptual state can cause you to φ without any such *intention* to φ on the part of the subject. This is not to say that φ -ing is involuntary. Insofar as you can *stop* yourself from φ -ing your action is under your voluntary control (Levy 2016). The important point is that when you do φ it is something that you *let happen* rather than something that you do. We can capture this with a third criterion:

- III) INTENTION-INDEPENDENT INFLUENCE: a perceptual motivation can lead a subject to φ without the subject intending to φ .

This third condition might seem at odds with the second condition which states that acting on a perceptual motivation is not obligatory. To clarify, our claim is that *sometimes* we are perceptually motivated to φ without in fact φ -ing and *sometimes* perceptual motivation leads us to φ involuntarily without our intending to φ . Reflexive accounts fail to accommodate the former and Imperativist accounts fail to accommodate the latter. But why believe perceptual motivations

can indeed influence our behaviour independently of our intention? We see at least three reasons:

- i) Absent-minded action: sometimes we act on a perceptual motivation without realising. If you concentrate on not scratching you might be able to resist the urge, but then when your mind is occupied by something else you find yourself scratching the itch. Here it is doubtful that you have any *intention* to scratch. Yet the perceptual state of itching can bring about your scratching behaviour without needing any such intention. The same goes for perceptual salience and pain: as you lose concentration you might find your attention falling on a salient item or your arm moving from a painful position.

This is a feature of perceptual motivation that Imperativists cannot accommodate. If our perceptual state merely makes a command then it could only bring about a response if we obey it i.e. if we specifically intend to perform the commanded action. But in absent-minded action we do no such thing. The Imperativist could respond that in each of these cases there *is* an intention to perform the relevant action, albeit one that is formed absent-mindedly. Although this possibility shouldn't be dismissed out of hand, we suggest that the burden of proof would be on the Imperativist to justify such a claim. After all, our claim that these actions are intention-independent fits with cases where one tries to excuse one's scratching, looking or withdrawing by saying "I didn't *intend* to do that!".

- ii) Overpowered behavior: Although perceptual motivation does not make *φ-ing* obligatory, it is continuous with states where *φ-ing* is obligatory. In other words, at the upper end of the continuum of perceptual motivation one is *compelled* rather than motivated to act. Applying one's hand to a moderately painful stimulus will motivate you to withdraw your hand. As the stimulus becomes more painful, this motivation will become stronger. But there reaches a point where you cannot help but withdraw your hand. The withdrawal response becomes obligatory. The same goes for salience. Some stimuli capture our attention in a way that cannot be resisted (Awh et al. 2012). The difference between obligatory and non-obligatory responses is a difference of *degree* rather than a difference of kind. As the pain gets worse, there is no stage at which your pain is replaced with a different kind of mental state that it compels action rather than motivates it. Rather, the relevant connection between the pain and the response reaches a level where our agency is overwhelmed and we are forced to act. This fits with the existence of borderline cases where it is unclear whether our response was obligatory: when our attention shifts to a highly salient object it can be hard to determine whether we *could* have resisted attending to it or whether the urge to attend was irresistible.

Imperativism allows there to be a scale of perceptual motivation (Martínez 2015), with some imperatives being stronger than others. But what it cannot make sense of is perceptual motivation being so strong that it constitutes perceptual compulsion. A command's influence on action depends on the agent intending to perform the commanded act. So a command, no matter how strong, can never compel an agent to act. Put another way, there is no force of imperative such that

one *must* obey it. Imperativists must then say that these cases of obligatory action are *discontinuous* with perceptual motivation: that a fundamentally different kind of psychological process explains the obligatory and non-obligatory cases. And this seems psychologically implausible.

The Imperativist may respond that when we are overpowered the intention to perform the relevant act is generated *involuntarily*. As such, there is an intention to mediate between the imperative content and the action. Again, we don't wish to rule out such a possibility out of hand. But our proposal again seems to better reflect what's intuitive. In the case of the painful hand it would be more natural to say "I didn't intend to withdraw my hand" than to say "I couldn't help intending to withdraw my hand". This shifts the burden of proof over to the Imperativist to explain why their explanation, in terms of obligatory intentions, is preferable.

iii) **Agential Phenomenology:** The agential phenomenology of perceptual motivation has a very distinctive character. When you are perceptually motivated to φ it does not seem as though your φ -ing depends on you intending to φ . Instead, it seems as though one *will* φ unless one resists doing so. When a salient stimulus pulls on your attention, attending to the item feels like something that will happen if you let it. Resisting the impulse to shift attention is effortful, but if one does shift attention it feels effortless. You can simply let the action unfold without having to initiate it. Put another way, our phenomenology is consistent with perceptual motivation influencing our action independently of our intention.

Imperativism does not sit well with this agential phenomenology. When presented with a command to φ one faces a choice of whether or not to φ . There is no reason for it to feel as if one *will* φ unless one resists. On the contrary, such a phenomenology would misrepresent the role of intention in one's action, for if Imperativism is true one will only φ if one *intends* to φ . You might be strongly motivated to form this intention by your perceptual state, and might have a background desire to perform the commanded action, but the point remains that you will only perform that action if you intend to do so. Commands have no way of side-stepping our agency.⁵

Now we might consider a possible way in which Imperativists could indeed satisfy INTENTION-INDEPENDENT INFLUENCE. In fact, some versions of the view would not be affected by our problems because they do not rely on the analogy of imperatives with public language commands, according to which the command itself can only bring about an action if one intends to obey it. Martínez and Klein (2016)'s view, for instance, is an explicit move against attempts to characterise imperative content in terms of public language imperatives, and rather proposes to gloss them in terms of informational signals in sender-receiver games. We believe this Imperativist view resembles very much the Affordance Model we propose in Section 3, and we

⁵ The third requirement of NON-INTENTIONAL INFLUENCE implies there being non-intentional actions, but if one thinks that action is always intentional, as for instance in the sense that it is something an agent does that is always "intentional under some description" (Davidson 1980) then one can read this requirement as implying that there are non-intentional *behaviours*, which is something the Imperativist needs to account for in any case.

will thus leave its assessment until after we have presented our proposal. That being said, not all Imperativist views are of this sort, and we have some reason to believe that Watzl and Hall's way of thinking about imperatives resembles public language imperatives. In particular, Watzl explicitly states that perceptual states share with public language imperatives the feature of having both informational and imperative content (Watzl 2017, p. 127, footnote 34). He later refers to a difference between perceptual imperatives and public language imperatives regarding what the commands tell the subject to do (Watzl 2017, p.129), but this difference doesn't count against the relevant analogous feature that occupies us here, namely, intentions to obey the commands. In the case of Hall, although he denies imperatives are to be understood as how they are expressed in *English* (Hall 2008, p. 526), his overall proposal treats imperatives very closely to the functioning of linguistic entities. So it seems that some forms of imperativism rely on a certain (although arguably weak) analogy with imperatives in public language. Regardless of whether they share other features with linguistic entities or not, one relevant common feature is the fact that imperatives construed this way seem to rely on intentions to bring about action. To the extent that this is so, we believe that our objections regarding INTENTION-INDEPENDENT INFLUENCE apply.

2.5 Dual Content Accounts

Dual Content views posit perceptual states that have both descriptive and desire-like content. This promises to accommodate the thought that it is the experience itself that motivates our action. An itch, for example, could be construed as a perceptual state that is both a representation of a state of your body *and* a command to get rid of that state. Such a view of perceptual motivation comes close to Millikan's (1995) notion of 'pushmi-pullyu' mental states that simultaneously perform a descriptive and a directive function. Thus, dual content states have a mind-to-world direction of fit (descriptive content) and a world-to-mind direction of fit (directive content). Such an account would satisfy DIRECT MOTIVATIONAL FORCE insofar as the motivational force of the itch comes directly from the perceptual state of the subject—specifically from the imperative content—and not from further associated states.⁶

Remember that one of the key problems for Conative Accounts was that they could not satisfy DIRECT MOTIVATIONAL FORCE. The Dual Content account promises to meet this condition. The relevant motivational content is *built into* the perceptual state rather than belonging to some state or states associated with the perceptual experience. Thus, the Dual content view appears to improve on mere Descriptive views by also having an imperative aspect but suffers from the same problems

⁶ Rather than positing one state with dual contents, one might explain perceptual motivation in terms of two *distinct* states. On such a Dual State view, perceptual motivation is constituted by a belief-desire pair (or a descriptive state-imperative state pair) that happens to be instantiated when the relevant perceptual state is. But the Dual State Account cannot accommodate DIRECT MOTIVATIONAL FORCE because the motivational force doesn't come from the perceptual state per se but rather from a package of states that happen to be instantiated together with it. Moreover, doubts can be raised about whether we can really posit such extra states every time there's a case of perceptual motivation.

associated with Imperativist views. The fact that in addition to the imperative content there is a descriptive content doesn't make a difference with respect to the problem that Imperativist views do not satisfy the INTENTION-INDEPENDENT INFLUENCE condition. So with regard to the motivational force of a dual content state, this view collapses into one of the other views already dismissed.

Moreover, the attribution of a complex content to a simple perceptual state like an itch seems an ad hoc move. Should we then think that dual contents of this sort are part of perceptual experiences in general? Perceptual states mainly present the world to us as being a certain way, and so it is difficult to see how a perceptual state could at the same time constitute a command for the world to be some way.

We contend that any view of perceptual motivation that is content-based, eg., descriptive, conative, and imperativist, will have trouble accommodating the sort of connection that seems to exist between motivating perceptual states and action, precisely because content in itself — of whatever kind — cannot per se connect a perceptual state to an action *beyond* merely indicating or commanding which action is to be taken. In contrast, reflexive views, which do not appeal to content to explain perceptual motivation, suggest a closer link by positing an automatic connection between the perceptual state and the required action. But, as already seen, they cannot accommodate the voluntary character of perceptual motivation.

3 An Affordance Model of Perceptual Motivation

3.1 Affordances and Automatic Motor Initiations

As an alternative to the five views presented, we offer an account on which an itch affords scratching, a salient stimulus affords attending and a pain affords withdrawal: the Affordance Model. The model proposes that perceptual motivation be understood in terms of affordances automatically initiating performance of an action and thereby directly motivating us to act in specific, voluntary ways and, in some cases, triggering intention-independent behaviours. The account thus satisfies the requirements set out above for an adequate account of perceptual motivation. The account posits a closer connection between perception and action as it doesn't rely on any kind of content to explain the motivational force of these states.

Affordances are opportunities for action for subjects with certain abilities (Gibson 1979). A tree can afford climbing for an organism capable of climbing, and a teapot can afford grasping for an organism capable of grasping. Any given environment will present a creature with a field of affordances: a set of actions that it's possible for them to perform in the current situation. The notion of affordances has received a great deal of interest and has functioned as a key concept in ecological psychology, cognitive psychology and neuroscience, as well as in music theory and industrial design. It is a theoretical concept that describes perceptual experience and can be captured linguistically by the structure *[verb]-able*: a stone is perceived as throw-able, a tree is perceived as climb-able, or an apple as eat-able. Importantly, these opportunities for action are *perceived* in the environment rather than inferred from other perceptual information (Gibson 1979; Michaels 2003).

Building on Gibson, more recent work in cognitive neuroscience has suggested that perceiving an affordance can *automatically initiate* performance of the afforded action. Affordance perception thus triggers ‘automatic motor initiations’ (AMIs). On this view, itches automatically initiate the action of scratching. Similarly, items perceived as salient automatically initiate the action of attending and the pain of touching a hot surface automatically initiates the action of withdrawing from the stimulus.

What is an AMI? Performance of an action unfolds over a duration of time and the early stages of this process involve covert neural processes that precede any overt bodily movement. Consider the act of catching a ball. Before your hand starts to move towards the ball there is a neural process that results in appropriate motor signals being sent to your muscles. Crucially, these covert early stages of action can occur without actually unfolding into an overt bodily action. Your perception of a catchable ball triggers the initiation of a catching response, but the relevant neural process is suppressed before it has the chance to bring about any bodily movement. This kind of automatic motor preparation is ubiquitous in the cognitive neuroscience literature. Perception of items can automatically initiate actions appropriate to that item: seeing a teapot, for example, results in activation of the motor process required for grasping (Tucker and Ellis 1998). And perception of another agent can automatically initiate imitation of that agent’s action: seeing a ballet dancer leap triggers the motor process involved in actually leaping.

Such responses to affordances in our environment are thought to play an important role in *motivating* particular actions. Ridderinkhof et al. (2011) summarise a range of findings in cognitive neuroscience that indicate an associative mechanism that initiates motor responses independently of intention. They also explore the conscious and unconscious mechanisms involved in the inhibition of inappropriate responses. Regarding the first-person experience of these neural processes, they observe: ‘The phenomenology of action affordances is typically described in terms of attracting or repelling solicitations: intrinsic to the experience of an affordance is that stimuli incite or summon certain actions.’ (2011, p.180).

How does this apply to our target states of itching, pains and salience perception? The neuroscientific evidence suggests that similar AMIs occur for such states. Hall himself notes that itching is associated with motor activation even when subjects do not scratch (2008, p. 533). A salient stimulus triggers the initial stages of an attentional shift even if that shift doesn’t occur (Bisley and Mirpour 2019). Similarly, a painful stimulus triggers a withdrawal motor response even if the agent doesn’t actually withdraw. This response is even present in cases where an agent observes someone else touching something painful (Anelli et al. 2012).

Our proposal is that the motivational force of these perceptual states should be explained in terms of AMIs. When your perception of a catchable ball initiates a catching response, you might thereby be motivated to actually catch the ball. Put another way, the fact that you have covertly *started* the process of catching the ball motivates you to *continue* the process and allows it to unfold into overt behaviour. By the same token, the initiation of scratching by an itch motivates you to scratch, the initiation of an attention shift by a salient stimulus motivates you to shift attention and the initiation of withdrawal (or some other relevant response) by a painful stimulus motivates you to withdraw.

A complication here is that many AMIs occur unconsciously and other neural processes prevent them from unfolding into overt action without the agent ever knowing. The cases relevant to us are those in which the AMI occurs with sufficient strength that the subject notices it and that it will result in overt action unless specifically suppressed by the agent. An itch triggers a scratching response in such a way that you *will* overtly scratch unless you specifically resist the urge to do so. To understand the motivational force of such an AMI, consider the experience of being pushed around in a crowd. You might judge that you shouldn't move forward from your current location. However, if the crowd is pushing you forward you have an immediate motivation to move forward. You feel the *force* of the crowd pushing you forward and have to exert yourself to resist that force and hold your ground. So although the crowd offers no intelligent *reason* for you to move, the force of the crowd pushing you forward *motivates* you to let yourself be pushed forward. Similarly, the force of the AMI unfolding toward overt action *motivates* you to allow the relevant action to occur.

3.2 Requirements for Perceptual Motivation

Now, how does this account capture the three features of target mental states described above? (I) DIRECT MOTIVATIONAL FORCE requires that the subject's perceptual states motivate her to φ directly. AMIs have a *direct* link to action: the itch leads automatically to the initiation of a scratching response without the need for intermediate beliefs or desires, the salient stimulus leads to the automatic initiation of the action of attending without additional beliefs or desires, and the painful stimulus leads to the automatic initiation of withdrawal in the same way. AMIs are motor initiations for actions that endow the perceptual state with DIRECT MOTIVATIONAL FORCE. Thus, in the Affordance Model, the perceptual state itself motivates one to act without the assistance of further beliefs and desires.

There is, however, a possible problem here. As mentioned above, not all instances of affordance perception lead to the automatic motor preparation of the afforded action. The relevant preparation only occurs if certain background conditions are in place. Among these background conditions are the relevant *goals* or *aims* of the agent. If one is in the business of cleaning up the table after dinner, cups can be perceived as graspable and motivate one to grasp them. But if one is just seeing aligned cups on a shelf in a museum, for instance, the graspability of the teacups is still perceived as such but one will probably be totally unmotivated to grasp them. Here the grasping action is not initiated (or not initiated strongly enough to be felt) because the relevant background goals are not in place. In other words, the motivational force of one's perception of a grasping affordance is contingent on one's wider conative states. But now it's less clear that the motivation qualifies as *direct*. Similarly, in other contexts affordances can motivate without any specific goals or intentions on the part of the subject, simply in virtue of the *course of action* that a subject has been following. It is in this sense that Siegel (2014: 2, our emphasis) notes that “[soliciting affordances] are generally structured by how you are *already acting* in a situation—not only by how you can act or are disposed to act in it”. The motivating

affordances in this case depend on your previous steps in the course of action and so the motivational force of one's particular perception would be contingent on one's wider process of acting. Again, it is dubious that the motivation here qualifies as *direct*. It seems that if itches, pains and perceptual salience only trigger AMIs when agents have the right background conditions, then DIRECT MOTIVATIONAL FORCE can't be satisfied after all.

How should we deal with this objection to the Affordance Model? One possible way around this is to maintain that this kind of contingency on background conditions is *compatible* with DIRECT MOTIVATIONAL FORCE. There is a crucial difference between background conditions *enabling* or *causing* perceptual motivation to appear and those background conditions *constituting* perceptual motivation. Background conditions (some or all of them) might be relevant for distinguishing which affordances attract the subject and which do not, and thus which cases of affordance perception may lead to action and which remain at the perceptual stage. But this is significantly different from background conditions by themselves constituting perceptual motivation. The AMI may be contingent on some wider mental states, but the locus of motivation is still the AMI. Just as a desire being causally dependent on wider mental states does not count against that desire being motivational, an AMI being causally dependent on wider mental states does not count against it being motivational.

This defence of the Affordance Model might take us some way, but it doesn't fit very well with our intuitive conception of the target phenomenon. Itches, pains and perceptual salience states seem to motivate by themselves without background conditions playing any role: even if the pain in my arm will hurt more if I want to move the arm in a certain way, it can be the kind of pain that hurts no matter the movement I intend to do or the course of action that I have followed. Similarly, a strong perceptual stimulus motivates one to turn the head towards the stimulus quite independently of one's goals at the moment. Thus, in cases like these the background conditions do not play any role.

The Affordance Model can capture this feature of perceptual motivation and, in the process, offer a more convincing way of accommodating DIRECT MOTIVATIONAL FORCE. Although *some* AMIs are contingent on the subject having suitable background beliefs and desires (or the subject already acting in a certain way), not all AMIs have this contingency. Where the initiation of a grasping response to a teacup only occurs if the right background mental states are in place, the initiation of a scratching response occurs *regardless* of one's wider mental states and previous course of action. Indeed, this captures the distinctive motivational force of perceptual motivations. They motivate us in a recalcitrant manner that we don't even find in other cases of automatic motor initiation.

It is also worth noting here that the Affordance Model captures the *specificity* of perceptual motivation. Affordances invite us to do specific actions, so the required specificity comes built-in to the perceptual state. Motor preparations are preparations to make particular bodily movements. So where a desire can motivate you without motivating any particular action, an AMI motivates a specific response. Itching initiates the particular act of scratching, pain the particular act of withdrawal

and salience the particular act of attending. These responses don't leave open alternative ways of responding to what we are perceiving.

Having dealt with DIRECT MOTIVATIONAL FORCE we can now move on to consider the other two conditions. With respect to (II) NON-OBLIGATORY FORCE, in the Affordance Model perceptual states motivate the action voluntarily. Itches afford scratching but it is not mandatory for me to scratch, and when I feel the pain in my hand I am free to leave my hand on the painful stimulus despite a strong urge not to do so. AMIs are best understood as a kind of *quasi-reflex*. The key difference from reflexes is that performance of the initiated action is not obligatory. Initiation of the action, however, is obligatory. What's voluntary is whether that initiated action is permitted to unfold. Summarising recent work on affordances and motor initiations, Ridderinkhof et al. note: '...the fact that stimuli may attract (or repel) certain actions does not imply that execution of these actions is inevitable. Calls can be resisted, at least in principle.' (2011, p.180).

With respect to (III) INTENTION-INDEPENDENT INFLUENCE, AMIs have the distinctive agential status that we are looking for. If a scratching response is initiated, then the agent need not form a specific intention to scratch. Instead, if the agent simply allows this potentiated response to unfold then they will scratch without needing such an intention. Unlike Imperativist accounts, the Affordance Model cuts out the stage at which the agent *obeys* a perceptual state. In perceptual motivation, the subject lets the action of *φ-ing* happen.

Consider again the three challenges raised against Imperativist accounts. First, we can see how perceptual motivations might trigger actions when we are *absent-minded*. When we aren't concentrating, we often act on affordances without realising (consider playing with your keys, or tapping your feet to some music). This is because perceiving the relevant affordance (under the right circumstances) automatically triggers a motor response. In situations in which we are not absent-minded, the initiated action can be resisted (if the AMI is not too strong). But when we are absent-minded, the action that would otherwise be resisted is *disinhibited* and we make a bodily movement without the need for any intention. For instance, a salient item automatically prepares a shift in attention, and suppressing this automatic response is an effortful task. When we lose concentration, this AMI is left in the driving seat and we find our attention moving to the salient stimulus.

Second, it explains how we can sometimes be *overwhelmed* to act. Ridderinkhof et al. (2011) cite cases in which '...the action-selection system is 'hijacked' by the response activation as triggered by the task-irrelevant action affordances.' (2011, p.177). This kind of involuntary response can also be found in the cases we're considering. Withdrawing a hand from a painful stimulus might be voluntary when the AMI is sufficiently weak, but if the pain is severe and a stronger AMI is triggered we have no choice but to withdraw. This is simply because a sufficiently strong triggering of an action cannot be suppressed. These neural events can generate movement even if such a movement is not intended. Motor-initiation is on a continuum running from easy-to-suppress, to hard-to-suppress and then into *impossible* to suppress. A fairly salient stimulus is such that we can easily resist attending to it. A highly salient stimulus is much harder to resist. And an extremely salient stimulus takes controls

of our actions. In each of these cases, a motor response of attending to the stimulus is prepared. The only difference is in the strength of the preparation.

Third, it explains the *agential phenomenology* of perceptual motivation. It seems, quite correctly, as though one *will* recoil from the painful stimulus unless one resists the urge to do so. It *feels* like you will φ unless you resist the urge to do so because you really *will* φ unless you resist the urge to do so. So when one acts on the affordance it seems that this would happen independently of one *intending* to do so. Put another way, on the Affordance Model our agential phenomenology accurately reflects our situation.

It is now appropriate to recall that not all Imperativist views seemed incapable of satisfying INTENTION-INDEPENDENT INFLUENCE, as Martínez and Klein's (2016) account does not rely on the analogy with public language imperatives. If imperative content is glossed in terms of informational signals in sender-receiver games, then this view of imperativism and the Affordance Model *seem to be very similar*: they both rely on neural mechanisms and information processing to account for the link between a given stimulus and an instruction to act in a certain way. There would still be a relevant difference between both accounts, though: Imperativism claims that it is imperative *content* that motivates you (in virtue of the kind of signal it specifies), and the Affordance Model proposes that *non-contentful* states such as AMIs are what motivate you to act. Our account thus explains perceptual motivation without appealing to a special kind of content, so it should be preferred on grounds of parsimony.⁷ Further assessment of the respective advantages and problems related to this difference could be a task for further work. For now, our aim has been to propose a new candidate on the table that can explain perceptual motivation better than most accounts have done so far.

4 Final Remarks

Automatic action initiations (AMIs) are already one of the components in our model of the mind. We have suggested that we can accommodate the motivational force of itches, perceptual salience, pains and other similar states using this familiar component. There's no need to give these mental states any special content that hooks them up with the actions they motivate, as Imperativists do. And there is no need to appeal to further additional beliefs or desires to accommodate motivational force as Descriptivist, Conative and Dual content/states views do. Instead, these states trigger the relevant action in a quasi-reflexive manner and our experience of these affordances motivates us to perform the actions afforded. On a general note, see that, in contrast with content-based accounts, the Affordance Model is able to explain the close connection of perception to action through the functioning of affordances and AMIs and, improving on reflexive accounts, it allows for action to be voluntary and under the control of the agent. Moreover, the account could accommodate any

⁷ This alone is neutral with respect to the question on whether affordance perception is ultimately representational. Mainstream ecological approaches would deny that it is, but other views that incorporate affordance perception would characterise it as representational (Siegel 2014).

other cases of perceptual motivation that might be found. Even cases where we are perceptually motivated to perform a mental action could be accounted insofar as the relevant AMIs are in place.

Overall we think the Affordance Model is better placed than the other views discussed—with the relevant considerations made for some Imperativists views—to account for perceptual motivation of some perceptual states like itches, pains and perceptual salience. Even if one is not convinced by its merits, we hope at least to have shown that it is a serious contender among possible views.

Funding We would like to thank the audience of the Workshop "Consciousness, Agency and First-Person Representation", organised by Pepa Toribio and Miguel Ángel Sebastián at the University of Barcelona, for helpful comments on a previous version of this paper. We are also thankful to the two anonymous reviewers of this journal. Their comments helped to improve the paper significantly. Financial support for this work was provided by MINECO, Spanish Government, research projects PID2020-115052GA-I00 and PGC2018-093464-B-I00, funded by the Agencia Estatal de Investigación (Spain) and a 2019 Leonardo Grant for Researchers and Cultural Creators, BBVA Foundation (the Foundation takes no responsibility for the opinions, statements and contents of this project, which are entirely the responsibility of its authors).

Declaration

Competing Interests The authors have no competing interests to declare.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anelli, F., A.M. Borghi, and R. Nicoletti. 2012. Grasping the pain: Motor resonance with dangerous affordances. *Consciousness and Cognition* 21 (4): 1627–1639.
- Awh, E., A.V. Belopolsky, and J. Theeuwes. 2012. Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences* 16 (8): 437–443.
- Bain, D. 2011. The imperative view of pain. *Journal of Consciousness Studies* 18 (9-10): 164–185.
- Bisley, J.W., and K. Mirpour. 2019. The neural instantiation of a priority map. *Current Opinion in Psychology* 29: 108–112.
- Davidson, D. 1980. *Essays on actions and events*. Oxford: Oxford University Press.
- Gibson, J.J. 1979. *The ecological approach to visual perception*. New York: Psychology Press.
- Hall, R. 2008. If it itches, scratch! *Australasian Journal of Philosophy* 86 (4): 525–535.
- Hume, D. (1739/2000). *A treatise of human nature* (Oxford philosophical texts), David fate Norton and Mary J. Norton (eds.), Oxford: Clarendon Press.
- Klein, C. 2015. *What the body commands*. Cambridge: MIT Press.
- Levy, Y. 2016. Action unified. *The Philosophical Quarterly* 66 (262): 65–83.
- Martínez, M. 2015. Pains as reasons. *Philosophical Studies* 172 (9): 2261–2274.

- Martínez, M., and C. Klein. 2016. Pain signals are predominantly imperative. *Biology and Philosophy* 31: 283–298.
- Michaels, C.F. 2003. Affordances: Four points of debate. *Ecological Psychology* 15 (2): 135–148.
- Millikan, R.G. 1995. Pushmi-pullyu representations. *Philosophical Perspectives* 9: 185–200.
- Ridderinkhof, K.R., B.U. Forstmann, S.A. Wylie, B. Burle, and W.P.M. van den Wildenburg. 2011. Neurocognitive mechanisms of action control: Resisting the call of the sirens. *Wiley Interdisciplinary Reviews: Cognitive Science* 2 (2): 174–192.
- Siegel, S. 2014. Affordances and the contents of perception. In *Does Perception Have Content?* ed. Berit Brogaard, 39–76. Oxford: OUP.
- Tucker, M., and R. Ellis. 1998. On the relations between seen objects and components of potential action. *Journal of Experimental Psychology* 24 (3): 830–846.
- Watzl, S. 2017. *Structuring mind: The nature of attention and how it shapes consciousness*. Oxford: OUP.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.