

# The ethics of attention: An argument and a framework<sup>1</sup>

Sebastian Watzl

## 1. Introduction

Suppose that you had no other information about a person than facts about her attention. The description you have uses only ‘attention’ and close relatives as a psychological predicate. That is: you have information about what she focuses her attention on; how much attention she pays to certain things; what is salient to her or has a tendency to grab her attention, and what – by contrast – tends to never get her attention. You also know about her capacity to control her attention: how much she has of it and when she exercises it. On the other hand, you are not told what that person believes or desires, what she perceives, what she intends; or what emotions she may have. Only attention; nothing else. I claim that you could still paint an informative picture of her character, her values, and what she cares about; indeed, you could provide decent estimates of many of her other mental states or characteristics. You may infer much about what the person believes or what she sees from what she pays attention to. Information about someone’s attention alone, it seems, allows us to give a rich account of the structure of a person’s mind, her experiences and what kind of life she likely lives. A psychological theory that uses attention as its only primitive psychological term, in other words, would be comparatively rich in its expressive capacities.

The above exercise illustrates how central attention is in an account of mind and agency: it delineates important contours of a person’s life. Several contemporary philosophers (cf. Wu 2011, 2014; Watzl, 2017; Ganeri, 2017; Jennings, 2020) have, relying in part on the large empirical literature on attention in psychology and neuroscience, defended versions of what I will here call the *descriptive centrality of attention*: attention plays an important role in (a) describing an individual’s mind and agency, and (b) explaining many central facts about that individual.

This paper argues for a normative implication of the descriptive centrality of attention: given the descriptive centrality of attention, attention should also be *normatively* central. We should develop an ethics of attention that delineates the normative pressures on that central structure of the mind. In this paper, I argue for this claim and sketch a framework for the study of the ethics of attention.

## 2. The argument for the normative significance of attention: From the ethics of belief to the ethics of attention

I use the ‘ethics of attention’ in this essay in analogy to the ‘ethics of belief’ that is the title of a seminal paper by William Clifford (1877). Since the resulting debate between Clifford and the psychologist and philosopher William James (1896), ‘the ethics of belief’ has come to designate an area of philosophy that discusses which normative pressures, if any, govern our practices of forming, maintaining and relinquishing beliefs and how those normative issues relate to the nature of belief (cf. Chignell, 2018). The ethics of belief has evolved to be a rich, multi-faceted and tightly structured field between epistemology, philosophy of mind, action theory, decision theory, and ethics. Some people in this field argue that belief is essentially normative and constitutively governed by – for example – distinctive norms regarding truth (e.g. Shah, 2003); others argue that the nature of belief exactly precludes the existence of belief norms (e.g. Glüer and Wikforss, 2013); and some think that any rational norms of belief flow from general principles of rationality (Rinard, 2017). To speak of the ethics of belief, in my sense, is not to take a specific stance in this debate; it merely indicates a distinctive study of the normative pressures on belief and how they interact with what beliefs are. The ethics of belief is a field that allows a focused discussion of disagreement.

My argument in this paper is that it is at least equally important to have an ethics of attention. Given the descriptive centrality of attention, we need a field of study of which normative pressures, if any, govern our practices of attention: what norms govern what we should attend to; how we should engage our capacity for attention; or when we should begin and when we should stop paying attention to something. Like the ethics of belief, the ethics of attention connects those normative questions to issues regarding the nature of attention, i.e. to what may or may not be subject to such normative pressures. How, for example, are salience and attention related – and what does that tell us about the respective normative pressures on salience and attention? To what degree can we control our attention? How are different forms of attention connected to each other? An ethics of attention will show how those questions impact and are in turn affected by claims about which norms are appropriate for attention. In my sense, the ethics of attention is not only about *moral* pressures on attention. The relative significance of different sources of normative pressures on attention and their interaction, indeed, is one major topic in the ethics of attention (see Section 8 below). The ethics of attention investigates which patterns of attention are rational, epistemically demanded, morally, aesthetically or epistemically blame-, or praiseworthy, ethically right, just or unjust, or prudentially appropriate. Do all such forms of evaluation exist? What is their significance? These are central issues in the ethics of attention. Speaking of the ethics of attention also does not imply that there are *constitutive* norms of attention. This question, again, is one to be debated within the ethics of attention, a field of study that allows a focused discussion of disagreement.

Why believe that there can be and should be an ethics of attention in this sense?

I suggest that we can find an answer by starting with a more general question: which aspects of the mind, generally, can and should be the focus of a sustained normative assessment? Let us look at some examples. We might, for example, think of virtue theorists as developing an ethics of character traits. They tell us what normative pressures there are for maintaining and relinquishing certain character traits and how those normative issues relate to the nature of character traits and the role they play in our individual and collective lives. And we might think of some deontologists as developing an ethics of the maxims that guide our decision-making: they develop a theory of the normative pressures on those maxims and connect this theory with one of the roles such maxims play in action and decision making. In recent years, many further aspects of the mind have become the subject of a field that resembles the ethics of belief – fields that connect the discussion of the nature of those aspects of the mind with discussions of the norms, if any, that may apply to them. We have, for example, discussions of the nature of desire and the question of whether some desires are irrational or immoral, or whether desire can be a source of epistemic insight (e.g. Tenenbaum, 2010). Similar debates have concerned emotion (e.g. De Sousa, 1990, Brady, 2013). In the case of perception, we have seen an intense debate linking questions about the nature of perception and its epistemic role (cf. Lyons, 2017 for a review), and also whether perception itself might be assessable for rationality (e.g. Siegel, 2016). Finally, consider how questions about the nature of concepts have been linked with questions about whether there might be ethical, political, or epistemic pressures to possess or not to possess certain concepts or to change them – a field that often is explicitly described as ‘conceptual ethics’ (Cappelen et al., 2019).

What, then, is the general principle by which an aspect of the mind becomes a central target of normative analysis? I suggest that the significance of the normative evaluation of an aspect of the mind (or behavior, or social interaction), other things being equal, roughly should be commensurate to its descriptive or explanatory significance. If, for example, one thought that stable, situation-independent, character traits played little role in an account of the mind and in explaining a person’s behavior (which is instead largely explained by social structures or situational features), then the evaluation of such character traits should also not be normatively central (cf. Doris, 1998). If, by contrast, character traits, indeed, deeply shaped a person’s experience, mind, and action, then it would be appropriate to make them a central object of normative evaluation. Similarly, a focus on the normative role of the emotions is justified to the degree to which the emotions, descriptively, deeply shape a person’s life and mind. Finally, the recent focus on conceptual ethics is partly driven by the insight that the concepts we use greatly influences the way we think and feel, how we interact with others, and which lines of inquiry we pursue. Which concepts we use explains many things. Therefore, a normative analysis of concepts is called for. The general principle I suggest is that – other

things being equal – normative significance should follow descriptive significance (I will discuss below what the other relevant factors might be).

Note that the idea here is not that (descriptive) psychology determines ethics, or that an *is* in some other way implies an *ought*. Rather, I'm making a normative claim: when we normatively assess a mind, a person, behavior or experience, we should target those aspects of the mind, the person, the behavior or the experience that play important roles in it rather than those that are only of minor explanatory significance.

Now, we are ready to state *the argument for the normative significance of attention*. I will put it as an argument for the conclusion that *if we should have an ethics of belief, then we should also have an ethics of attention*. So, here is the argument:

1. Attention and salience are explanatorily central in an account of the mind and in the explanation of experience and behavior to a degree that is, at least, roughly commensurate to the explanatory role of belief. (call this the *descriptive centrality of attention*)
2. If the descriptive centrality of attention holds, then attention and salience should also be normatively central in the following sense: answering normative questions about attention and salience is at least as important as answering normative questions about belief (call this *the descriptive-normative link*)
3. So, Answering normative questions about attention and salience is at least as important as answering normative questions about belief. (this is *the normative centrality of attention*).

Given that many philosophers think that answering normative questions about belief is quite important indeed (as the field of epistemology indicates), it follows from the conclusion of this argument, i.e. the normative centrality of attention, that answering normative questions about attention will be quite important as well. We do, in other words, need an ethics of attention.

In the next sections, I will defend the argument just laid out.

### **3. A further reason for why we need an ethics of attention**

Some people might say that, surely, when we, as philosophers, ask about which aspects of the mind or behavior we should make the focus of our normative inquiry, the issue of how much the general public *cares* about getting answers to, or getting guidance regarding normative questions about that aspect of the mind is at least as important as the descriptive centrality of that aspect of the mind. Many people, outside of philosophy, wonder what to believe in an age of fake news; they wonder whether anger is appropriate in a certain situation (cf. Srinivasan, 2018), or what, if anything is wrong with the use of certain concepts (such as racial ones). It is the general public's interest in those questions, one might think, that should drive what philosophers at their normative workbench investigate (just like it's the public's interests that should figure heavily in what kind of machines engineers should develop).

It is not obvious just how heavily the fact that the public is already interested in certain normative questions should figure in which normative questions we as philosophers investigate. Normative innovation might result exactly from seeing the significance of a normative analysis of an aspect of the mind that the general public had tended to overlook. It was, arguably, philosophical insight into the explanatory significance of personal decision-making that first led to philosophical analyses of the normative significance of autonomy and, hence to some of the normative innovations of the enlightenment period.

Yet, a defender of the normative significance of attention can also accept a constraint from public interest in the relevant normative issues. As a matter of fact, the public seems to care a lot more about normative questions regarding attention than normative philosophy does.

In our contemporary context – dominated as it is by mass information and social media—many in the public are faced with pressing challenges about how to organize information, how to individually distinguish what is relevant from what isn't, and how to together agree on what is currently most deserving of collective attention. Accordingly, in the public debate there is no shortage of normative opinions about attention: about what is right, healthy, important, and correct with regard to attention. Social media activity, for example, by many is especially deemed as receiving too much of it. Others argue vigorously that in political debates the particular voices of some persons or groups of persons are getting too much or too little attention. A growing number of self-help and meditation practices, such as mindfulness training, promise help for the proper regulation of attention in an age of distraction (cf. Kabat-Zinn, 2009; Bregman, 2011; Williams and Penman 2011; Newport, 2016). Activists warn of the 'industrialization of [...] attention capture' (Wu, 2017) and some call the liberation of attention the defining moral and political struggle of our 'time' (Williams, 2018). Think also about discussions of spin doctors and troll factories (or certain political actors): they misdirect and distract. Those who criticize either those actors or those who fall for them make claims in the ethics of attention: they argue that it is normatively criticizable to misdirect someone's attention, or that it is wrong to pay too much attention to what is false. Finally, think of the debates about statues, paintings, or history books: the presence of the great men of history in these places gives them a lot of everyone's attention. Our social and physical environments are, literally, engineered and created for certain patterns of attention. Attention in thought, in sympathy, but also in perceptual, in visual or auditory attention: the statues at the central squares are designed to be looked at. Many, in the contemporary public debate, argue that this focus is inappropriate. The painter Titus Kaphar, for example, tries to show with his art that the dreams and hopes of the slave boy, who is depicted in the background of a painting or as leading the horse of a general, they too deserve our attention.<sup>2</sup> Kaphar argues for a claim in the ethics of attention: justice demands changes in our attentional patterns.

The ethics of attention, thus, already is an area of public debate. It is normative philosophy that is lagging behind. While the normative evaluation of

our own and other's attention is central to ordinary normative thought and discourse, the majority of contemporary normative philosophy, from moral theory, to epistemology and decision theory, contains no systematic discussion of the ethics of attention.

So, if the public's interest in answers to normative questions about attention and need for guidance in the public debate is indeed an (additional) constraint on whether we, as philosophers, should develop a focused study of the ethics of attention, then we have, indeed, more reason to focus our energies on (developing) such a field of study. Attention is *both* descriptively central and the public cares about normative guidance with regard to it.

#### **4. Is attention, in fact, descriptively central?**

Let me now look deeper at premise (1) of the argument for the normative significance of attention: the descriptive centrality of attention. Is it really correct that attention is an explanatorily central aspect of the mind and of behavior? If the public's interest in the ethics of attention were driven by a mistaken view of its descriptive significance, then maybe philosophers are right to neglect a normative analysis of attention.

On the face of it, an attack on the descriptive centrality of attention is an uphill battle. In its defense, we can refer to the rich empirical psychology of attention, and to the philosophical accounts mentioned in the introduction. Attention and salience are, arguably, centrally involved in at least the following areas: (a) *Decision making*. A recent review (Orquin and Loose, 2013) emphasizes the 'active role [of attention] in constructing decisions' (p.203), which leads to a host of 'downstream effects' (ibid.). (b) *Intentional agency and agentive control*. Here, Wayne Wu (2011, 2016) has argued for a central role of attention in the explanation of how intentional agency and control over action are possible. (c) *Self-control*. Berkman et al., (2017, p.423) argue that '[a]ttention plays a crucial role in... self-control by gating which options enter the choice set at any one moment and foregrounding their salient attributes' (see also Watzl (forthcoming) where I argue that by appeal to attention, self-control can be fully explained without willpower or a divided mind). (d) *Perception*. Attention is known to affect all aspects of perception, from perceptual precision, response times, to perceptual appearances (see Carrasco, 2011 for a review), and, finally, to whether we perceive certain items at all - as in, for example, Simons and Chabris (1999) famous invisible gorilla experiment. (e). *The Self*. In quite different ways, Dicey-Jennings (2020) and Ganeri (2018) argue that central aspects of the self can be explained by appeal to attention. (f) *Consciousness*. Philosophers have debated which accounts of consciousness are compatible with the effects of attention on consciousness (cf. Block, 2010). And an ongoing debate is about whether attention might be necessary for consciousness (see Watzl (2017) where I argue for a constitutive role of attention with regard to all forms of consciousness). (h) *Memory*. Attention is known to interact with memory in complex ways. What we remember both affects attention, and attention strongly affects what is encoded

in memory (e.g. Chun and Turk-Browne (2007) (i) *Introspection*. Several philosophers have argued that introspection either is or is made possible by a special form of attention (e.g. Chalmers, 2003; Goldman, 2006) or relies on attention in a different way (Watzl, 2017) (j) *Emotion*. Philosophers from De Sousa (1990) to Brady (2013) have argued that certain patterns of salience and attention may be ‘constitutive of emotional experience’ (Brady, 2013, p.181).

To sum up: almost anywhere we look in the mind and in forms of behavior we find plausible arguments to the effect that attention is centrally involved in and complexly interacts with other aspects of the mind. Of course, each particular case to some degree is controversial. But the evidence at least points strongly in the direction of premise (1), i.e. a central role of attention in an account of the mind and of behavior.

Yet, a more subtle attack on the descriptive centrality of attention might be launched. Attention, one might argue, is too variegated to be, itself, descriptively central.

Start by reflecting on the many forms of attention: in one form attention is perceptual: a subject focuses her attention on some perceptually presented items over others. This perceptual attention can be highly focused, or it can be dispersed through a whole scene. But attention also comes in non-perceptual forms: your attention might be occupied by a pain or a feeling; or by a train of thought like planning for your friend’s birthday. Finally, your attention might be occupied by an embodied action: consider playing soccer with all your attention, as opposed to kicking some balls while your mind is occupied with the birthday planning. How could there be a normative theory, one might ask, that treats perceptual focus, getting distracted by a feeling of nausea, and the attentive engagement in sports in anything like a unified manner? To make matters worse, some instances of attention are paradigmatic voluntary actions, like when I, just like that, shift my attention, with mental effort from the melody of a guitar solo to the rhythm played by the drums. Some other cases of attention, by contrast, seem passive and involuntary, like when my attention is drawn, against my will, to a beeping phone that disturbs my musical experience. Attention seems to be not one phenomenon at all; but many phenomena.

This commonsense attack on the unity of attention might be supported further by appeal to the science of attention. It has often been treated as a finding of the cognitive science of attention that it ‘is not a unitary construct’ (Carrasco, 2011, p.1517), best ‘considered as a property of multiple, different perceptual and cognitive operations’ (Chun et al., 2011, p.76), or ‘a label’ for several loosely related cognitive processes (Styles, 1997, p.1) (also e.g. Itti et al., 2005, Anderson, 2011, or Hommel et al., 2019). As Jonardon Ganeri (2017, p.25) puts it, one might argue that ‘attention ... is not a single psychological kind’.

If attention were indeed disunified in this way, then, it may seem, it cannot be, after all, explanatorily or descriptively central. Consider that if attention were not a single psychological kind, then nothing should be lost by replacing the term or concept ‘attention’ with terms or concepts referring to the many more specific

psychological kinds. But if a term or concept is dispensable in this way, then – one might argue – what it picks out cannot be explanatorily central.

Yet, it does not follow from the fact that there are many forms of attention that those forms have nothing interesting in common. One might, for example, think of them as determinates of a single determinable. Or as multiple realizations of attention. Or, relatedly, one might think of those forms as ‘unified by analogy’ (cf. McDaniel, 2010): just like a healthy organism, a healthy piece of broccoli, or a healthy circulatory system have something interesting in common, while also – of course – being interestingly different, so different forms of attention might still have the type of unity required for descriptive centrality. Ganeri (2017, p.12), for example, thinks of all of them as ‘ongoing [forms of] structuring of experience and action’. Attention, unlike belief or memory, is an ongoing – process like – aspect of mentality, and unlike specific desires or specific perceptual experiences, attention is a structural aspect of mind and action.<sup>3</sup> So, even if attention were not a single psychological kind, several psychological kinds may still share enough similarities so that those kinds form a distinctive aspect of the mind that allows for a targeted normative analysis.

In my view, though, a more direct answer to the disunity worry is available (see also Wu, 2011, 2014, 2016) for a related, though distinct, response. According to the priority structure view I develop in Watzl (2017), all forms of attention consist in the agent’s activity of regulating priority structures, which order the parts of the subject’s on-going (occurrent) mental life by their relative priority to the subject. This priority structure view unifies the varieties of attention. Perceptual attention consists in prioritizing some parts of a perceptual state over other parts of that state. If the subject’s attention is, for example, visually focused on an object, then the state of seeing that object is prioritized. Attention to external objects thus gets explained by the relative priority of aspects of the subject’s mind. When attention is focused on a feature (like the color of the object or event) then the state of seeing those features is prioritized. The priority structure view allows for degrees of attention as we move up or down the subject’s priority ordering. When attention is non-perceptual what is prioritized is a non-perceptual aspect of the subject’s on-going mental life. What is prioritized may be a bodily sensation or feeling, mental images, occurring ideas, and occurrent desires that are constitutive of an embodied action or train of thought. The priority structure view thus unifies all forms of attention by taking as the primary notion that of a mental state’s relative priority for the subject. The forms differ only in which aspect of the subject’s mind has the highest relative priority.

The priority structure view also unifies passive and active forms of attention, as well as attention and salience. Attention is a mental activity, a personal level on-going process. A person, on this view, is always engaged in that activity. We are constantly regulating our priority structures. Yet, while this is something we are doing, how we are doing what we are doing is subject to both an active and a passive force. The passive force is salience. The salience of a mental state consists in the degree to which that mental state commands prioritization (a type of



imperative content). Saliency is a feature of subject-level states, but it is not under the subject's control (just like the content of a perceptual state is subject level but not controlled by the subject). The active force interacts with the passive force and shapes the subject's prioritizing to the degree to which she herself guides her own activity. Such guidance can take all forms that it can also take for bodily action. Attention is subject to deliberation, normative judgment, intention, desire and choice, and fine-grained online control.

The central notion in this framework, to repeat, is the notion of a priority structure, where the parts of the agent's mental life are compared by their relative priority or priority weight. The framework can be expanded in many ways depending on how one thinks of the relevant notion of priority. In my own view, relative priority does not have a reductive definition. We have a grip on the notion through its reflection in conscious experience: here relative priority shapes the field of consciousness into center and periphery. One way to take this is to think of our conscious acquaintance with how peripheral or central a part of our mind is as fixing the reference to an underlying type of relation between those parts of our mind.<sup>4</sup>

A detailed defense of the priority structure view of attention, of course, cannot be provided in this paper (see Watzl (2017) for details). One way to understand the argument of this paper is as conditional on a unified theory of attention, such as the priority structure view. *If* such a view is right, *then* the normative centrality of attention follows. Indeed, I believe that the priority structure view is especially well suited to show how attention is descriptively and explanatorily unified – and, as I will show below, it lends itself well to a structured development of the ethics of attention.

On the other hand, as I suggested above, we might not need a unified view like this in order to defend the descriptive centrality from the threat of disunity (unity by analogy might suffice). In this sense, the argument of this paper is not entirely hostage to objections to a unified theory of attention.

## 5. The wrong kind of object objection

Let me then move to premise (2) of the argument for the normative significance of attention: the descriptive-normative link. I've already mentioned the central motivation behind this premise: an explanatory central aspect of the mind, is – because of that – an important lever in the mind. We want to know how to adjust it properly. If the normative importance of an aspect of the mind would not at least roughly follow its descriptive or explanatory importance, then our normative theories threaten to be parochial: we focus our energies on evaluating an aspect of the mind or of behavior that makes little difference. Suppose you found that an aspect of your mind that is of only little descriptive significance is in normatively bad shape. You would likely respond 'who cares!? It doesn't explain much about me or about what I do'. By contrast, suppose that you became unsure about whether a very significant part of your mind is in normatively good shape: this, clearly would be much more disturbing.

Yet, these general considerations cannot be enough to establish the descriptive-normative link. Many aspects of our physiology, for example, arguably play important explanatory roles regarding our behavior, and, indeed, also with respect to how we feel. But that would not make your physiology an interesting target of normative assessment. If attention is to be an important target of normative analysis, then it must be the kind of thing that permits such an analysis. One might, thus, argue that despite its centrality in the mind attention is not the kind of ‘thing’ that is amenable to the relevant type of normative assessment. I will call this ‘the wrong kind of object objection’. In this section, I will answer two version of this objection.

First, one might argue that while attention is descriptively central, it is *not a phenomenon situated at the right level of description* to allow for normative assessment. Theoretical engagement with attention since the mid Twentieth century has been dominated by a complex, and multi-faceted discussion in cognitive psychology and the neurosciences. This leads naturally to thinking of attention in terms of a set of neuronal or psychological mechanisms. But once we do that, questions about good or bad forms of attention, questions of how attention figures in moral considerations, in the assessment of the character of a person, or its proper role in epistemic inquiry, may appear to be misguided: neuronal or psychological mechanisms may seem not to be the right kind of thing for which such normative questions can be properly posed. Indeed, one might, like Jennifer Hornsby, argue that normative questions about, for example, rationality, responsibility, or moral standing arise only with respect to personal level phenomena (e.g. Hornsby, 2000). Beliefs and transitions between beliefs, according to this view, are assessable for rationality precisely and only because they are personal level. Insofar as other aspects of the mind of a person, like her emotions (De Sousa, 1990) or her perceptual states (Siegel, 2016) also impact the person’s normative standing, they, too, must be personal level states or activities.

One could argue against Hornsby that sometimes sub-personal states or activities can also be normatively assessed and impact a person’s normative standing.<sup>5</sup> Yet, there is a more direct response: the fact that there exists a rich sub-personal science of attention does not imply that attention is a sub-personal phenomenon. Just consider that much is also known about sub-personal perceptual processes; and yet it is the person who perceives her surroundings (Burge, 2010); the sub-personal processes underpin and enable the subject’s perceptual state. Similarly, a person’s conscious experience, arguably the paradigm of a personal level phenomenon, is underpinned by sub-personal processes. We should treat attention, like perception and conscious experience, as a personal level phenomenon: attending is something the person does; attentional processes are what underpins the attending. According to the priority structure view, attention, indeed, is a *paradigmatic* personal level psychological phenomenon: priority structures connect all ongoing aspects of the subject’s mind and provide them with personal level unity. In addition, attention is constitutively integrated with other personal-level phenomena such as perception, belief, desire, or conscious experience, that form the parts of priority

structures. The priority structure framework sketched above thus provides a thoroughly personal level view of attention.

A second version of the wrong kind of object objection might appeal to the claim that at least some forms of attention are not under voluntary control. What is salient to you in a specific situation, and hence what captures your attention, evidently is not something you control. In other cases, a subject may have no awareness of, for example, a subtle shift of attention that occurred to her. Insofar as norms are supposed to provide guidance to the subject, talk of norms for a completely involuntary aspect of the mind, or ones the subject is unaware of may appear to be misguided.

This version of the objection at best targets some types of attention norms. Norms of fittingness for emotions, for example, would apply to the emotions whether or not we control emotions or whether or not we are aware of having those emotions. Similarly, norms of fittingness would apply even to involuntary salience or to shifts of attention we are unaware of. In the literature on the ethics of belief, indeed, some people try to explain our *inability* to control belief by appeal to a constitutive truth norm for belief (Shah, 2003). Even when it comes to *responsibility*, it is highly controversial that we are responsible only for what is under our voluntary control (e.g. Smith, 2005). On the priority structure view, even involuntary forms shifts of attention are always changes to an agent's activity: they are ripples in a sea of agency. This arguably makes attention even more amenable to a normative analysis. Be that as it may, though: how the agential and involuntary aspects of attention interact with the normative assessment of attention is a debate that should, like in the ethics of belief, occur within the ethics of attention. It is not an obstacle to the normative significance of attention.

## 6. Content-based and manner-based attention norms

I hope to have convinced that attention in principle can be assessed normatively. Yet one might still believe that no *systematic* normative investigation can be given. One might argue that we need, at least, a proof of concept, to show that an ethics of attention that resembles the field of the ethics of belief is even possible.

The rest of this paper will – in three sections – respond to this worry, by showing how such a systematic investigation may be provided. I will do so by sketching a framework for such an investigation. We can start with the priority structure view. It gives us a clear target for the normative assessment of attention. Norms of attention are norms for the regulation of priority structures. The proposed framework then classifies potential norms of attention along three dimensions: whether they are manner or object-based, instrumental or non-instrumental, and whether its source is moral, prudential or epistemic. I will use this framework to locate proposals for specific normative pressures on attention in the literature, and in order to situate the normative discussion of attention within the wider field of normative philosophy.

The first dimension, then, is between what I will call *content-based* and *manner-based* attention norms.

We can think of the priority structures I have introduced in two tiers. On tier 1, the elements of the structure are mental states that occupy the subject's attention to various degrees, and the relation between them is their relative priority to the subject (we can allow for ties in relative priority). If I'm paying more attention to the drums than to the instrumental solo, then my hearing the drums has strict priority in my mind compared to hearing the solo. Or if I'm so excited about a possibility that I keep thinking about it and notice little of my surroundings, then my excitement and thoughts about the possibility are prioritized over my experience of the world around me. On tier 2, the elements of the structure, by contrast, are external items and the structure orders them by how much the person attends to them. The ordering in tier 1 determines an ordering in tier 2. Tier 2 characterizes the world as it is according to the attention occupiers. We can think of the structure elements at tier 2 as the intentional objects of the structure elements at tier 1. If there are non-intentional, non-directed aspects of the mind (for example, a feeling of nausea or a headache that occupies your attention), then there is more to tier 1 than can be captured in tier 2. Also, if a part of the world is, as it were, attended 'twice over' like when you both look at rustling leaves and listen to them, then we can capture details in the subject's attention on tier 1 that will be lost if we only consider tier 2.

Note that tier 2 structures in the first place provide an ordering of attended items, but probably do not define an interval scale (we can't define:  $x$  receives twice as much attention as  $y$ ). But, if the structures are 'big enough' we can define rough degrees by the number of actual and potential intermediate elements (cf Watzl, 2017, p.82). Sometimes a normative evaluation may want to say that a person ought to attend much more to one thing rather than another, and not merely pay a little more attention to that thing.

Note also that each priority structure already contains information about what is and what is not included in the subject's current mental life. For the moment neglect the structuring relation, and just think of a tier 1 structure as a set of mental states. We can think of that set as a partitioning of the space of *potential* mental states, into those that the subject actually has and those that she does not have (just like the set of all of a subject's beliefs, i.e. every proposition she believes, partitions the space of possible beliefs into those the subject has and those she doesn't have, i.e. propositions believed and those not believed). Two priority structures can differ simply in that one contains an element that the other does not contain. So, by being given a certain priority structure, we already know what is excluded from it. If you 'overthink' a situation, tier 2 attention norms can express the idea that that thought should not have entered your mind. You should not have attended to that situation at all.

With these two tiers in place, we can now define content-based attention norms as those norms that are exclusively formulated at tier 2. Manner-based attention norms, by contrast, are those that pertain directly to how the person's mind is organized at tier 1.

Content-based norms, then, are what is described when we say that a person should pay more attention to X than to Y, that she should not completely ignore the Fs, that it would be appropriate if people paid more attention to X, or that a virtuous person is one who pays enough attention to certain features of the world. Content-based norms will ignore in what *way* a person pays attention to something, and instead are concerned with a coarser level of detail: a person should – in some way or other – pay attention to certain features.

As an example where such coarse-graining appears appropriate, consider a hiring process (see Banaji and Greenwald, 2013; see also Siegel, 2016 for discussion). Suppose that person evaluating various résumés harbors an unconscious prejudice against members of a certain group (the out-group). As a result and as a reflection of that prejudice, the evaluator pays little attention to the strong parts of the résumé of outgroup members, while she pays much more attention to weaker parts of their CVs. By contrast, when she considers the résumé of an in-group member, she pays a lot of attention to the strong parts and comparatively little attention to the weak parts. We may want to say that this distribution of attention is normatively criticizable (below I will discuss ways of thinking of the normative source of that critique). Arguably, though, it is normatively irrelevant *how* she re-distributes her attention. She may make sure to perceptually focus her attention equally on all relevant parts of the CVs, she may make sure to give enough thought to all parts after she has looked at them, or she may counteract her bias by directing her attention in her imagination to the qualities that a person with that CV is likely to have.<sup>6</sup>

As a second example, consider gratitude (see Bommarito (2013), and Chappell and Yetter-Chappell (2016)). Someone who possesses the virtue of gratitude (who is a grateful person) must pay a sufficient amount of attention to the value of what others have done for her. Someone who merely knows that others have done something valuable, or is peripherally aware of that value, but pays no attention to that value is not grateful. '[A]ttention is essential to the virtue. We cannot be grateful by simply being aware that someone has helped us', as Bommarito (2013, p.100) puts it. It is thus plausible that good character is characterized by paying enough attention to the value of what others do for us. But arguably it, again, doesn't matter how we pay attention to it: whether we perceptually take it in, when it is happening, by thinking attentively about it, or by often imagining how valuable those deeds were.

Yet, the example of gratitude also illustrates the limits of content-based approaches to attention norms. Suppose that, as Bommarito observes, that a 'teenager under the influence of Nietzsche might attend to the support of his or her parents but take it to be evidence of how disgustingly weak Mom and Dad are'. (ibid.). Such a teenager is attending to the support of his parents, but does not display the virtue of gratitude. So, not all ways of attending to the help of others are a sign of gratitude. There are two ways of taking this observation. On the one hand, we may leave the relevant norm of attention at the content level and provide a fuller characterization of gratitude by supplementing having the right attention structure with, for example, having the right emotions directed at

the supporting person. On the other hand, we may take Bommarito's observation as showing that the content-based assessment of attention, in this case, is insufficient. If the subject attends to the value of another's help by prioritizing negative emotions directed at this help, then that distribution of attention makes no contribution to her gratitude. Whether we go one way or the other here plausibly depends on whether we think that a positive contribution of attention can be separated out even if the subject's emotions are misguided: is the person at least partially grateful if she attends to the value of others, even if that attention is accompanied by spite? If yes, then we have a content-based attention norm as a partial component of gratitude. If no, then we are here really dealing with a manner-based attention norm.

Let us then get to manner-based attention norms. Here, as I said, we are concerned not with what the subject pays attention to, but with tier 1 that concerns the organization of the subject's mental life. We thus look in more detail at the *way* a subject pays attention to something.

We have already seen one example of how to distinguish such ways: which mental states are prioritized over others. Suppose, for example, we would like to say, as Iris Murdoch (1970) does, that morality requires that a person should pay attention to the concrete other subjects she is engaged with (a 'just and loving gaze directed on an individual reality', p.33). Staring, perceptually, at others plausibly is not a way of fulfilling that normative requirement: rather what should occupy our attention is reflective thought about others as loci of intrinsic value and worth (though, again, the manner based treatment could be debated: if we can see the intrinsic value and worth, then maybe the prioritization of such seeing is a way of encountering the other appropriately; see e.g. Werner (2014) for discussion of whether moral properties are perceptible). As another example, consider the positive evaluation of attentive touch, especially the 'close form of caring touch that forms an essential element of social bonding and human development' (Fulkerson, 2016). There clearly is a good that pertains to the tactile attentive engagement with close others (family, friends, children, and loved ones). Mere inattentive touch does not seem to have the same value; and equally importantly it is attentive *touch* that is valuable, mere attentive thought or seeing does not have the same value (though it may have other values).

Another way of distinguishing ways of attention is by considering *how* a subject's priority structures are regulated. Here we can distinguish, as I mentioned, a passive force on the evolution of those structures in terms of what is salient to the subject and an active force, when a subject guides her attention, for example, with direct and effortful control.

One can argue both that a special value attaches to forms of passively evolving attention structures and to the active control over attention. For the first, consider another example from Bommarito and Chappell and Yetter-Chappell (2016): the role of attention in modesty. Chappell and Yetter-Chappell (2016) argue, correctly I believe, that a modest agent is 'one who is not disposed to find their own positive attributes especially salient' (p.457). She hardly ever finds herself pulled to give attention to those positive attributes. She may have an

accurate view of her positive traits but doesn't find them particularly salient. Maybe, indeed, in line with Yetter Chappell and Yetter-Chappell's observations, a modest person actually does spend much of her time attending to her positive attributes: maybe she is paid to look into them and intentionally and effortfully gives them her attention. But such attention requires effort for her. It's painful. For the modest person, attending to her own positive attributes is not automatic (while the immodest person feels the pull to attend to those attributes). Other plausible examples of salience constituted virtues discussed by Chappell and Yetter-Chappell are: being a good friend might be constituted by finding the friend's positive attitudes more salient than her negative attitudes, someone who is able to 'forgive and forget' finds past wrongs not particularly salient (she need not actually have forgotten about them; she is just not drawn to attend to them). These examples illustrate the positive value that may attach to salience, manners of attention that without control evolve in the right direction.

On the opposite side, consider that we may also positively evaluate being in control of our attention. Immanuel Kant (1775-76, p.62) expresses an extreme version of such an evaluation when he says that '[t]he greatest perfection of the powers of the mind is based on our subordinating them to our power of choice ... For this sake, attention ... [is] only the useful for us, if [it is] under the free power of choice, so that involuntary attentiveness ... produce[s] much harm'. While we might not want to follow Kant that attention is 'only useful' if it is 'under the free power of choice', we may still believe that having enough control over our own attention and being able to guide it according to our own values is an important and valuable aspect of self-control and personal integrity.

It is not my goal here to evaluate such proposals, but only to use them as illustrations of norms of attention that pertain to the manner of attention, and not what it is that we attend to.

## 7. Instrumental and non-instrumental attention norms

I now get to a second dimension of the framework: we can distinguish different types of attention norms. We can, on the one hand, evaluate an agent's priority structures and how she regulates them *instrumentally*, i.e. with regard to an end that is achieved through that form of attention, or we can evaluate attention *non-instrumentally*, i.e. evaluate a form of attention on its own without regard to its upshots. A specific form of a non-instrumental attention norm would be constitutive norms of attention. Some people, as I mentioned above, argue that there are constitutive norms of belief or of specific emotions.<sup>7</sup> One important debate in this area then will concern whether there are any non-instrumental or even constitutive attention norms. If so, what are they? If not, why not?

With Immanuel Kant's negative evaluation of 'involuntary attentiveness' we have already seen a clear case of an instrumental attention norm: Kant seems to think that such involuntariness is problematic because it 'produce[s] much harm'. What is normatively problematic thus concerns the consequences of such a manner of attending. For a similarly instrumental, though less one-sided,

evaluation of the relative costs and benefits of controlled and uncontrolled forms of attention, consider Chandra Sripada's (2018) discussion of the costs and benefits of mind-wandering (i.e. when our cognitive, non-perceptual attention hops uncontrolled from one thought to the other). Such mindwandering sometimes may have negative consequences: it prevents you from achieving routine daily goals and may make you feel bad about yourself (Killingsworth and Gilbert, 2010). On the other hand, arguably there are also good consequences: it may distract from boredom, and might lead to escape routines, and lead to insight and deep learning (what Sripada calls the advantages of 'exploration').

Consider also one way of looking at the hiring case I have discussed. One thing that is problematic about the unequal distribution of attention to the good and not so good aspects of the CVs of in-group and out-group members is that as a consequence, a so prejudiced evaluator will tend to not hire the best person for the job. She will miss out on candidates she would have wanted to hire (Banaji and Greenwald, 2013). Her distribution of attention thus is problematic because it is an ill-calibrated instrument for the purpose to which it is put.

Not all instrumental attention norms will have to do with prudence alone: with what best satisfies the desires of the agent. Consider mindfulness meditation. Such meditation is constitutively characterized by certain attention structures: the Oxford English Dictionary, for example, describes meditation generally as ways of 'focus[ing] one's mind for a period of time'<sup>8</sup> in specifically prescribed ways. Mindfulness specifically has been characterized by the idea that the agent intentionally makes higher-order thoughts about her first order experiences a center of her experience (cf. Latham, 2016).

Mindfulness meditation, by its practitioners and advocates, is often described as a good thing. They advocate an ethics of attention. What, though, is supposed to be good about it? It might serve as an epistemic tool, maybe delivering better insight into the true nature of the soul (though see Ratnayake and Merry, 2018 for a critique), and - of course - it is widely promoted as a practical tool, helping to reduce anything from stress, anxiety, to substance abuse and chronic pain, and promoting increased quality of life and increased prosocial attitudes (whether it succeeds in those regards is another matter. See Goyal et al., 2014 for a recent meta-study suggesting skepticism). Many practitioners and many traditions in which mindfulness meditation is important, though, reject the instrumental approach (cf Purser, 2019 for a critique of such uses of mindfulness). They might advocate mindfulness and meditation practices as a good way of life. Mindfulness training ought, they would say, not just help an agent realize goals she antecedently possesses. The instrumental approach, according to this attitude, does not do justice to the transformative effects of mindfulness and meditations practices. They, so the proposal, help the agent to become a better, more virtuous person: they help the agent, for example, to be more compassionate. Through the relevant training at focusing her mind on her own feelings, experiences, and emotions generally, the agent will be in a position to notice anger toward others before it 'takes a hold of her'. It thus helps her realize moral concern and compassion. Such an evaluation of attention would still fall under the



instrumental approach: priority structures get evaluated only instrumentally, by how well they serve to realize an independently given good or virtuous trait.

Finally, one might also regard attention as an instrument of inquiry. Suppose that you are trying to find an answer to some question, say how many windows there are on a large skyscraper (cf. Friedman, 2020). If you are trying to find the answer to that question, you must keep your attention focused in the right way. If instead, you started focusing attention, on say, the structure of the pavement, or on the clothes of the people on the sidewalk, your inquiry will likely fail. Attention norms thus might be derived from norms of inquiry, where attention serves as an important zetetic instrument (cf. Friedman, 2020 for a detailed discussion of the shape and importance of norms of inquiry).<sup>9</sup>

Are all norms of attention instrumental norms, i.e. norms that evaluate attention according to its consequences? Arguably this is not so.

An instrumental approach treats the relationship between the attention structures and the relevant good (e.g. compassion) as akin to a means-end relation. But arguably the relationship between the attention structures and certain virtues is closer than that: the virtue is *constituted* in part by certain attention structures. What it is for an agent to have that type of virtue is for her to have certain attention structures – both synchronically as well as diachronically. If an attention structure is constitutive of a virtuous character trait, then it can not only be evaluated instrumentally. It is part of what it is to be an agent with that virtue to have specific attention structures. We have already seen some plausible examples of this type: gratitude and modesty may both be constituted by certain forms of attention (on either tier 1 or tier 2).

Attention thus can be evaluated non-instrumentally if it is a constitutive aspect of a feature or state that can itself be evaluated non-instrumentally. What, aside from virtues like gratitude, may be examples of such features or states? One interesting area of investigation concerns the question of whether some attention structures are constitutive of *epistemically* evaluable states. Among these would be positively evaluated states such as knowledge, justified belief, or understanding, but also negatively evaluated states like prejudice.

For a potential example, consider that (at least) sometimes whether an agent knows something may constitutively depend on whether she has attended in the right way. Consider a doctor aiming to discover whether her patient has cancer. She looks through the patient's files and has only a quick glance at an X-ray (but doesn't pay much attention to it). From the evidence she has considered, it is reasonable to conclude that the patient is healthy. But had she paid more attention to the X-ray, she would have noticed a black spot. While the patient is in fact healthy, let us suppose, had the doctor considered the spot this would have undermined her justification to believe that the patient is healthy. One might argue (see Pollock, 1986; Goldberg, 2017) that the doctor in such as case does not know that the patient is healthy *because she had the wrong tier 2 attention structures*. One might take such cases to show that knowledge of certain facts requires that the agent inquires into the matter properly and responsibly. But in order to inquire into the matter properly and responsibly she must pay attention

to the right things (cf. Fairweather and Montemayor, 2017). But now we seem to have a non-instrumentally evaluable attention structure: whether or not the agent cares about knowing anything in this domain, there is a type of epistemic evaluation of the attention structure available.

There is an interesting question of whether some relation less demanding than constitution might also lead to non-instrumentally evaluable attention structures. Return to the hiring example for illustration. One form of evaluation, as we have seen, is instrumental: the evaluator should pay equal attention to the good parts of all CVs because only then will she be able to hire the most qualified person. But arguably there is also another form of evaluation: the prejudiced outlook is irrational (see Siegel, 2016 for further discussion). With no evidence (or indeed against available evidence) the evaluator takes outgroup members to be generally less qualified. Siegel (2016) argues that a tendency to have certain priority structures is itself irrational when it is in a specific sense she defines *based on* having such an irrational outlook: if we would like to say that the irrational outlook is non-instrumentally problematic, then the priority structure that is based on it will itself be non-instrumentally evaluable.

As I mentioned, whether there are any non-instrumental attention norms will be a matter of substantial debate. For now, I hope to have sketched the distinction between instrumental and non-instrumental attention norms enough for such a debate to get started.

## 8. The normative source of attention norms

I now get to a final distinction among different attention norms. This distinction concerns the normative source of the relevant norms.

A first source is prudential considerations. Having certain priority structures, and certain ways of regulating those structures over time can be evaluated according to how good those structures are for the agent who has them. Many instrumental attention norms have such a prudential source. Consider, for example, as we have discussed the costs and benefits of mind-wandering. Mindwandering may get in the way of an agent's pursuit of long-term goals, and in that sense, it is bad for the agent. But it may also be beneficial if it enables deep learning and serves as a source for welcome distraction.

Prudential considerations generally pertain to the wellbeing of the agent: to what is good for her (see Crisp, 2017). Well-being here may be thought of in terms of desire satisfaction: how well does a priority structure help the agent to satisfy her desires? But well-being can also be thought of in terms of the overall balance of pleasure and pain, or in terms of what is objectively good for the agent (one might, for example, hold that friendship or knowledge and understanding is good for the agent independently of whether it is associated with pleasure or pain or whether the agent has the relevant desires). An interesting question in this area concerns whether some forms of attention may themselves be in this way objectively good for the agent: maybe it is, for example, good for the agent to have enough control over her attention. Or consider when Marcus Aurelius

says that ‘those who do not observe the movements of their own minds must of necessity be unhappy’ (Aurelius, *Meditations*, 2nd book). Arguably, Aurelius here thinks that reflective attention to one’s own mind is objectively good for a person, and not just a means for satisfying her desires or because it will improve the overall balance of pleasure and pain.

A second source of attention norms is moral considerations. We have seen such a source at play when we considered the idea that the attention structures characteristic of meditation and mindfulness may serve as an instrument toward instilling the moral virtue of compassion in an agent. And we have also seen examples of moral virtues, like gratitude or modesty, that may be partially constituted by certain forms of attention. In these cases, the fundamental object of moral evaluation is a certain character trait, and attention gets morally evaluated either instrumentally or because it partly constitutes a morally virtuous character trait.

There is another virtue-theoretically inspired view that does not, though, take this indirect route toward the moral evaluation of attention structures. On this view, the morally virtuous person is constituted by certain attention structures not because those are constitutive of specific character traits. Rather, some patterns of attention and salience just are themselves morally virtuous. Good character requires having the right motivational susceptibilities. We find that aspect of virtue and of morality emphasized in thinkers from Aristotle, Iris Murdoch, John McDowell, Lawrence Blum, up to recent work by Seanna Shiffrin.

Which structures then ought the morally virtuous person to have? Arguably, there are several aspects. Arguably, it ought to be salient to her that a situation that has a moral or ethical dimension does have that dimension. For a virtuous person, for example, it will be salient that moral concerns about how we ought to live together are a relevant aspect of her working environment.

Further, the morally virtuous person ought to notice moral wrongs as moral wrongs (they ought to be high enough on her salience ordering). She pays attention to the situation of each individual without needing to consult a moral rule book. Consider a case mentioned by Lawrence Blum (1991). He talks about a situation where two people, John and Joan, are on the subway together. There is another woman who is carrying a heavy shopping bag. John and Joan are both aware of that woman. But the fact that she is uncomfortable is salient to Joan while it isn’t particularly salient to John. Blum thinks that the difference between what is salient to John and what is salient to Joan is of moral significance. While Joan’s attention is captured by the fact that a morally relevant value is at stake, John is not ‘attuned to’ that morally significant fact. Blum thinks that the moral difference between John and Joan is not fully captured in how the difference in what is salient is going to make them act. If John is generally like that then this shows a ‘defect in his character ... [h]e misses something of the moral reality confronting him’.

Finally, Seanna Shiffrin (2016) argues that there are moral requirements of non-negligence. Those requirements include patterns of attention: ‘to be nonnegligent is to be both attentive and responsive in thought and agency to

how the pursuit of one's (permissible) aims and the state of one's agency affect one's ability to satisfy one's other duties and responsibilities'. On Shiffrin's view, it is an agent's moral responsibility to pay enough attention to the potential side effects of her doing.

A third source of attention norms is epistemic considerations. We have encountered this source when we have considered whether certain priority structures are epistemically evaluable because they are constitutive either of certain forms of irrationality or of positive epistemic states such as knowledge or understanding.

Just like for the case of morality, there is also an interesting question of whether a form of attention may be a constitutive feature of an epistemically well-calibrated agent. Maybe, it is part of being epistemically well-constituted that one pays attention to the right things?

Consider the idea that science is not just concerned with knowledge, but with knowledge on matters of interest (Anderson, 1995; Kitcher, 2003). An ideal epistemic agent doesn't just know random facts, but she knows and pays enough attention to *important* facts, facts that objectively merit our attention. This idea is deeply ingrained in how we often think of cultural authority: it purports to identify things that are worth paying attention to – ideas worth considering, books worth reading, creations worth beholding, problems worth solving, histories worth trying to reconstruct – even independent of antecedent social purpose. These things produce pro tanto rational pressure to pay attention to them, regardless of whether or not you have antecedent interest in them, knowledge about them, or desire to pay attention to what someone else says you should pay attention to. The seriousness with which we often debate what exactly belongs in the canon, what it is that 'we all' epistemically speaking ought to pay some attention to here shows that it would be mistaken to simply and quickly attempt to push the relevant epistemic pressures on attention into prudential or moral considerations: it is, at last, a matter of serious debate whether the epistemic value of paying attention to important facts about the world and our place in it can sometimes outweigh moral or prudential considerations regarding what may lead us to improve our lot. Another interesting question here concerns which types of norms are the most fundamental norms. One might, for example, argue that epistemic norms can (or should) be derived from norms of inquiry (Friedman, 2020). Arguably, many attention norms can also be derived from norms of inquiry. If that were so, then epistemic norms and attention norms may be equally non-fundamental. I will conclude this section, by noting that the three sources of attention norms (prudential, moral, and epistemic sources) need not be exhaustive. They help to structure the debate at the current state of the discussion. Yet if, for example, there were primitive and constitutive attention norms about the conditions under which attention is fitting or apt, then those primitive attention norms would not be derivable from any other source. They would instead be constitutive of the nature of attention itself. Whether there are any such constitutive fittingness norms for attention will be one of the central topics a future ethics of attention will need to investigate.<sup>10</sup>

## 9. Conclusion

In this paper, I have argued for the normative significance of attention and attempted to improve our grip on the normative landscape with regard to attention. I have argued that given its descriptive centrality for mind and behavior, attention should be central also in normative philosophy. The public wants answers and guidance with regard to normative questions about attention. Philosophy should develop an ethics of attention that helps to provide guidance and that is commensurate in its richness to the descriptive significance of attention. In the second half of the paper, I have then sketched a framework that may help us to get started at developing the ethics of attention.

## Notes

- 1 This paper has developed out of a presentation initially started with Susanna Siegel. Together we have developed our thinking about attention norms. This material was jointly presented at a workshop on norms of inquiry at NYU. I have also presented it at Umeå University, Stockholm University, and the Rutgers-Columbia Philosophy of Mind Workshop. A presentation resembling more closely the present paper was then presented at a Workshop on Cross-cultural Perspectives on ‘Virtues of Attention’ at NYU Abu Dabi, at the Colloquium of the Institute of Philosophy in Tübingen, and as a keynote address at a graduate conference on Philosophy of Mind, Language, and Cognition at the University of Western Ontario. The paper received its last changes after a presentation to the Oslo Mind Group. Some material was added from applications for an ERC grant and a grant from the Norwegian Research Council. I would like to thank everyone at those occasions for their feedback, and also everyone who provided feedback on the written material. I would like to single out Susanna Siegel for many inspiring and stimulating discussions and excellent feedback. This result is part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement no. 101003208).
- 2 See: [https://www.ted.com/talks/titus\\_kaphar\\_can\\_art\\_amend\\_history?language=en](https://www.ted.com/talks/titus_kaphar_can_art_amend_history?language=en)
- 3 See Watzl, 2017 for some of the points of this paragraph.
- 4 The view allows that once reference is fixed the relevant relation may also hold between unconscious mental states.
- 5 Thanks to Anna Drożdżowicz for helping me to see this clearer.
- 6 Note that I don’t take it as obvious that the critique of the biased assessor of CVs must be situated fully at tier 2. There are probably some ways of attending to the good portions of the ingroup CV that are not an expression of a bias in favor of that group: they might strike a subject as repellent and disgusting and be salient and attract her attention in that way. A person whose attention is distributed in such a way would arguably *not* be normatively criticizable for paying too much attention to the good portions of those CVs (thanks to Max Kippersund for this point). This example illustrates that there are, in a sense, gradations between tier 1 and tier 2 norms: some norms cover a broad variety of ways of attending, while others cover only a very narrow band.

7 E.g. that a belief is correct only if it is true; cf. Shah, 2003; or that ‘anger that p is apt only if p constitutes a genuine moral violation’ (Srinivasan, 2018). Thanks to Anders Nes for drawing my attention to those examples.

8 OED.

9 Thanks to Andrew Lee and Jane Friedman for the connection to norms of inquiry.

10 Thanks for comments by Andrew Lee and Anders Nes that helped me in writing this paragraph.

## References

- Anderson, B. 2011. There Is No Such Thing as Attention. *Frontiers in Psychology* 2(246), 1–8.
- Anderson, E. 1995. Knowledge, Human Interests, and Objectivity in Feminist Epistemology. *Philosophical Topics* 23(2): 27–58.
- Banaji, M. R. and A. G. Greenwald. 2013. *Blind Spot: Hidden Biases of Good People*. Delacorte Press, New York.
- Berkman, E. T., C. A. Hutcherson, J. L. Livingston, L. E. Kahn and M. Inzlicht. 2017. Self-Control as Value-Based Choice. *Current Directions in Psychological Science* 26: 422–428.
- Block, N. 2010. Attention and Mental Paint. *Philosophical Issues* 20(1): 23–63.
- Blum, L. 1991. Moral Perception and Particularity. *Ethics* 101(4): 701–725.
- Bommarito, N. 2013. Modesty as a Virtue of Attention. *Philosophical Review* 122(1): 93–117.
- Brady, M. S. 2013. *Emotional Insight: The Epistemic Role of Emotional Experience*. Oxford University Press, Oxford.
- Bregman, P. 2011. *18 Minutes: Find Your Focus, Master Distraction and Get the Right Things Done*. Hachette UK, London.
- Burge, T. 2010. *Origins of Objectivity*. Oxford University Press, Oxford.
- Cappelen, H. D. Plunkett and A. Burgess (Eds.). 2019. *Conceptual Engineering and Conceptual Ethics*. Oxford University Press, Oxford.
- Carrasco, M. 2011. Visual Attention: The Past 25 Years. *Vision Research* 51(13): 1484–1525.
- Chalmers, D. J. 2003. The Content and Epistemology of Phenomenal Belief. In Q. Smith and A. Jokic (Eds.), *Consciousness: New Philosophical Perspectives* (pp. 220–272). Oxford University Press, Oxford.
- Chappell, R. Y. and H. Yetter-Chappell. 2016. Virtue and Salience. *Australasian Journal of Philosophy* 94(3): 449–463.
- Chignell, A. 2018. The Ethics of Belief. The Stanford Encyclopedia of Philosophy (Spring 2018 Edition). In Edward N. Zalta (Ed.) <https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/>.
- Chun, M. M., and N. B. Turk-Browne. 2007. Interactions Between Attention and Memory. *Current Opinion in Neurobiology* 17(2): 177–184. Chun, M. M., J. D. Golomb and N. B. Turk-Browne. 2011. A Taxonomy of External and Internal Attention. *Annual Review of Psychology* 62: 73–101.
- Clifford, W. K. 1877 [1999]. The Ethics of Belief. In T. Madigan (Ed.), *The Ethics of Belief and Other Essays* (pp. 70–96). Prometheus, Amherst, MA.
- Crisp, R. 2017. Well-Being. *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition). In Edward N. Zalta (Ed.). <<https://plato.stanford.edu/archives/fall2017/entries/well-being/>>.

- De Sousa, R. 1990. *The Rationality of Emotion*. MIT Press, Cambridge.
- Doris, J. M. 1998. Persons, Situations, and Virtue Ethics. *Nous* 32(4): 504–530.
- Fairweather, A. and C. Montemayor. 2017. *Knowledge, Dexterity, and Attention: A Theory of Epistemic Agency*. Cambridge University Press.
- Friedman, J. 2020. The Epistemic and the Zetetic. *Philosophical Review* 129(4): 501–536.
- Fulkerson, M. 2016. Touch, The Stanford Encyclopedia of Philosophy (Spring 2016 Edition). In Edward N. Zalta (Ed.). <<https://plato.stanford.edu/archives/spr2016/entries/touch/>>.
- Ganeri, J. 2017. *Attention, Not Self*. Oxford University Press, Oxford.
- Glüer, K. and Å. Wikforss. 2013. Against Belief Normativity. In. T. Chan (Ed). *The Aim of Belief* (pp. 80–99). Oxford, University Press, Oxford.
- Goldberg, S. C. 2017. Should Have Known. *Synthese* 194(8): 2863–2894.
- Goldman, A. I. 2006. *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford University Press, Oxford.
- Goyal, M. et al. 2014. Meditation Programs for Psychological Stress and Well-Being: A Systematic Review and Meta-analysis. *JAMA Internal Medicine* 174(3): 357–368.
- Hornsby, J. 2000. Personal and Sub-personal: A Defence of Dennett’s Early Distinction. *Philosophical Explorations* 3.1: 6–24.
- Hommel, B., C. S. Chapman, P. Cisek, H. F. Neyedli, J. H. Song and T. N. Welsh. 2019. No One Knows What Attention Is. *Attention, Perception, & Psychophysics* 81(7): 2288–2303.
- Immanuel Kant 1775–76 [2013], *Anthropology* Friedländer (1775–1776). Translated by G. Felicitas Munzel, In: *Lectures on anthropology* (pp. 37–256). Cambridge University Press, Cambridge.
- Itti, L., G. Rees and J. K. Tsotsos. 2005. *Neurobiology of Attention*. Academic Press, Burlington, MA.
- James, William 1896 [1956], *The Will to Believe and Other Essays in Popular Philosophy* (pp. 1–31). Dover Publications, New York.
- Jennings, C. D. 2020. *The Attending Mind*. Cambridge University Press, Cambridge.
- Kabat Zinn, J. 2009. *Wherever You Go, There You Are: Mindfulness Meditation in Everyday Life*. Hachette Books, London.
- Killingsworth, M. A. and D. T. Gilbert. 2010. A Wandering Mind Is an Unhappy Mind. *Science* 330(6006): 932–932.
- Kitcher, P. 2003. *Science, Truth, and Democracy*. Oxford University Press.
- Latham, N. 2016. Meditation and Self-control. *Philosophical Studies* 173(7): 1779–1798.
- Lyons, J. 2017. Epistemological Problems of Perception. The Stanford Encyclopedia of Philosophy (Spring 2017 Edition). In Edward N. Zalta (Ed.). <https://plato.stanford.edu/archives/spr2017/entries/perception-episprob/>
- McDaniel, K. 2010. A Return to the Analogy of Being. *Philosophy and Phenomenological Research* 81(3): 688–717.
- Murdoch, I. 1970 [2013]. *The Sovereignty of Good*. Routledge, London.
- Newport, C. 2016. *Deep Work: Rules for Focused Success in a Distracted World*. Hachette UK, London.
- Orquin, J. L. and S. M. Loose. 2013. Attention and Choice: A Review on Eye Movements in Decision Making. *Acta Psychologica* 144(1): 190–206.
- Pollock, J. 1986. *Contemporary Theories of Knowledge*. Rowman and Littlefield, Maryland.
- Purser, R. 2019. *McMindfulness: How Mindfulness Became the New Capitalist Spirituality*. Repeater Books, London.
- Ratnayake, S. and D. Merry. 2018. *Forgetting Ourselves: Epistemic Costs and Ethical*

- Concerns in Mindfulness Exercises. *Journal of Medical Ethics* 44(8): 567–574.
- Rinard, S. 2017. No Exception for Belief. *Philosophy and Phenomenological Research* 94(1): 121–143.
- Shah, N. 2003. How Truth Governs Belief. *The Philosophical Review* 112(4): 447–482.
- Shiffrin, S. 2016. The Moral Neglect of Negligence (April 18, 2016). In David Sobel, Peter Vallentyne and Steve Wall (Eds.), *Oxford Studies in Political Philosophy*, Oxford University Press, Oxford.
- Siegel, S. 2016. *The Rationality of Perception*. Oxford University Press, Oxford.
- Simons, D. J. and C. F. Chabris. 1999. Gorillas in Our Midst: Sustained Inattention Blindness for Dynamic Events. *Perception* 28(9): 1059–1074.
- Smith, A. M. 2005. Responsibility for Attitudes: Activity and Passivity in Mental Life. *Ethics* 115(2): 236–271.
- Srinivasan, A. 2018. The Aptness of Anger. *Journal of Political Philosophy* 26(2): 123–144.
- Sripada, C. 2018. An Exploration/Exploitation Tradeoff Between Mind Wandering and Goal-Directed Thinking. In Kalina Christoff and Keiran Fox (Eds.) To appear in *Oxford Handbook on Spontaneous Thought: Mind-Wandering, Creativity, and Dreaming* (pp. 23–34). Oxford University Press, Oxford.
- Styles, E. A. 1997. *The Psychology of Attention*. Psychology Press, Oxford.
- Tenenbaum, S. (Ed.). 2010. *Desire, practical reason, and the good*. Oxford University Press, New York.
- Watzl, S. 2017. *Structuring Mind. The Nature of Attention and How It Shapes Consciousness*. Oxford University Press, Oxford.
- Watzl, S. (forthcoming). Self-control, Attention, and How to Live without Special Motivational Powers. In M. Brent and L. Miracchi (Ed). *Mental Action and the Conscious Mind*. Routledge, New York.
- Werner, P. 2014. Moral Perception and the Contents of Experience. *Journal of Moral Philosophy* 13(3): 1–24.
- Williams, J. 2018. *Stand Out of Our Light: Freedom and Resistance in the Attention Economy*. Cambridge University Press, Cambridge.
- Williams, M., and Penman, D. 2011. *Mindfulness: A Practical Guide to Finding Peace in a Frantic World*. Hachette UK. London.
- Wu, W. 2011. Confronting Many-Many Problems: Attention and Agentive Control. *Noûs* 45(1): 50–76.
- Wu, W. 2014. *Attention: New Problems of Philosophy Series*. Routledge, London.
- Wu, W. 2016. Experts and Deviants: The Story of Agentive Control. *Philosophy and Phenomenological Research* 93(1): 101–126.
- Wu, T. 2017. *The Attention Merchants: The Epic Scramble to Get Inside Our Heads*. Vintage, New York.