

ATTENTION AS SELECTION FOR ACTION DEFENDED

Wayne Wu

Carnegie Mellon University

Attention has become an important focal point of recent work in ethics and epistemology, yet philosophers continue to be noncommittal about what attention is. In this paper, I defend attention as selection for action in a weak form, namely that selection for action is sufficient for attention. I show that selection for action in this form is part of how we, the folk, experience it and how the cognitive scientist studies it. That is, selection for action pulls empirical and folk-psychology together. Accordingly, philosophers who take seriously either source have reason to work with selection for action as their starting conception of attention. This conception provides a way to bridge empirical and philosophical concerns where attention is central. The theoretical advantages of selection for action have been obscured by the common opinion that it is easily refuted. I defend the position against many of the published objections and then deploy it to provide a foundation for the intuitive, but inchoate idea of attention being more or less, focusing on recent work in the ethics of attention.

1. Introduction

Penultimate Version (please quote from forthcoming article).

There has been a resurgence in interest on attention among epistemologists and ethicists (see the essays in (Archer, 2022); also (Fairweather & Montemayor, 2017; Mole, 2024; Munton, 2023; Panizza, 2022; Saint-Croix, forthcoming)). Ella Whiteley (2023) has recently examined harms due to a surplus of attention, and Georgi Gardiner has made the case for a central role for attention in epistemology, emphasizing the acquisition of appropriate attentional habits in individuals and groups (Gardiner, 2022). In these rich discussions, attention illuminates central philosophical issues.

Notably, philosophers maintain neutrality regarding the *metaphysics* of attention. Opting for less controversy, many draw on introspection and folk psychology. These approaches, however, have not fully leveraged the science of attention where attention is arguably one of the best understood psychological phenomena. This leaves reflection on the normative significance of attention at the surface of psychological understanding. The science can provide further depth, but a bridge is needed that makes available empirical understanding to ethical and epistemic investigation. This bridge is provided by attention's central role in guiding action, mental or bodily.

Attention is fundamentally connected to action (Mole, 2011; Wu, 2011b). Consider responding to a reason, say that a person has collapsed before you or deciding whether p is true given evidence E . William James (1890) stated what everyone knows about attention: attention is *the mind's taking possession of a target to deal with*, for example, responding to a person in need to aid them or to evidence in order to draw the correct conclusion. Attention is central to many, arguably all, actions. Consider the standard philosophical disruptor, an evil demon or neuroscientist, who abolishes attention at an inopportune moment, just as one is about to aid or infer. One is now prevented from selecting the person to help or the evidence to reason from.

Penultimate Version (please quote from forthcoming article).

The person in need becomes just one object among many, the evidence just one proposition among others (Campbell, 2002). The agent has lost track of the target: *Which person is in need? What evidence was I considering?* The agent's action requires the agent's attending to targets to inform a response. An adequate account of actions of ethical and epistemic concern requires understanding attention.

This essay defends attention as selection for action in the version I have defended as the starting point for philosophical reflection on attention. However, I stop short of my metaphysical commitments which are not needed to fruitfully inform ethics and epistemology (Wu, 2011a, 2011b, 2014). It seems to me that philosophers have wrongly viewed selection for action as simply refuted. Against this, I argue that a weak version of the view provides the right foundation for a philosophical psychology of attention. Section 2 sets a *bridge* principle from cognitive science to folk and philosophical psychology by demonstrating that William James' functional articulation of what attention is grounds the standard empirical paradigms. Since many philosophers working on attention draw on empirical work, they are committed to this functional account. Section 3 refutes a recent argument against the centrality of the functional account in science. In doing so, it underscores how that account is central to understanding attention and identifying *relevant* neural correlates. Section 4 addresses the inadequacy of published counterexamples to selection for action that appeal to folk-psychology and introspection, responding to many of the published objections to it. Finally section 5 deploys the bridge principle to examine the ethical significance of attention. I focus on Whiteley's recent discussion of harmful *surpluses* of attention. Gradability in respect of attention has not been given sustained, psychologically-grounded philosophical analysis. I use the bridge principle to initiate such an analysis, clarifying the intuitive but inchoate idea of *quantifying* attention.

Penultimate Version (please quote from forthcoming article).

2. The Jamesian Conception in Empirical Methodology

Let's begin with situating William James' ideas within the science of attention.¹ James noted:

Everyone knows what attention is. [*Attention*] is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others (James, 1890, p. 204, my italics).

I will call this the *Jamesian conception* understood not as James's own considered views but as his conceptualization of the folk psychological view of attention. The core idea (in italics) is correct: attention is the subject selecting a target to deal with. Here is a list of ordinary cases: recalling a memory of a past event or a fact to ponder or report (episodic and semantic remembering); looking at something to figure out what it is or to commit it to memory; listening to your partner at a loud party or instead surreptitiously surveilling the conversation behind you;

¹ I go beyond my earlier defense of the principle (Wu, 2014). What is needed is to show, as I do here, concrete cases of the empirical sufficient condition in actual experimental paradigms. The functional conception *unifies* these cases, and that theoretical unification is a strong argument for accepting the conception as a basis for organizing empirical results for philosophical work through it.

Penultimate Version (please quote from forthcoming article).

aiming at the center of the dartboard by focusing on it; scanning a crowd to locate (or avoid) a person or keeping an eye out in one's periphery; considering reasons to answer a theoretical or practical question and so on. We can easily multiply cases drawing on our own experiences in attending. What unifies these mental and bodily actions is the phenomenon that James rightly noted that we know. In each, a subject exemplifies mental selection to deal with something. This form of attention is sometimes spoken of as voluntary, top-down, or goal-directed attention, so:

Subject *S* attends to *X* iff *S* selects *X* for further processing in order to respond to *X*.

Explicit endorsement of this principle is not hard to find in empirical work.²

² Selection for action is routinely endorsed in the sciences: Chica et al (2013) p. 108 (though they speak of it as a mechanism, see Section 8); Sani et al. (2021) p. 1, Moore and Zirnsak (2017) p. 48; Shulman and Corbetta (2012) p. 113; Rungratsameetaweemana and Serences (2019) p. 181 (see further citations by Carolyn Dicey Jennings (2020), p. 8). Empirical instantiations of the conception (empirical sufficient conditions) can be derived from the textbook definition of attention as *selection (of a target) for further processing* (Bahmani et al., 2019, p. 129; Bear et al., 2020, Chapter 21; Holt et al., 2019, p. 225; Koester et al., 2021, Chapter 25; Workman & Taylor, 2021, p. 50). If selection of information for further processing is sufficient for attention, then so is selection for a specific type of processing, namely to generate action. Despite this, Denis Buehler claims that cognitive science “rather than being committed to [the Jamesian conception], rejects it” (Buehler 2023, 1679).

Penultimate Version (please quote from forthcoming article).

We are searching for a less controversial bridge principle, so extract from the biconditional the form of attention manipulated in the cognitive science of attention:

If S selects X to guide response, then S attends to X

This yields specific conditionals for experimental tasks as instances, what I have called *empirical sufficient conditions*:

If a subject selects X to inform doing Task T , then the subject is attending to X

Empirical sufficient conditions are instantiated relative to tasks T and are concrete expressions of the Jamesian conception in science. Consider these tasks and their associated empirical sufficient conditions whose formulations can be easily stated using the above frame conditional:

- *Dichotic listening*: Subjects select one of two auditory verbal streams to verbally shadow (repeat), so auditorily attend to the targeted stream.
- *Symbolic Spatial cueing*: On cue, subjects visually select a spatial location to monitor, then visually select the subsequent target to inform report, so visually attending to space and target.

Penultimate Version (please quote from forthcoming article).

- *Visual Search*: Subjects identify a specific target among a distractors. Subjects select stimuli to compare with a target template and, if there is a match, report it, thereby visually attending to the target.
- *Simple Span Tasks*: Subjects retain a sequence in memory, then later report the remembered items. Subjects mnemonically encode specific targets in working memory for later test, thereby cognitively attending to them.
- *Inattentional Blindness*: subjects count the number of ball passes among one group of players, ignoring a ball passed by a second group. While subjects visually select the ball for counting, a gorilla appears. The selection of the ball shifts attention away from the gorilla, so subjects attend to the ball but not the gorilla.
- *Attentional Blink*: subjects scrutinize a rapid serial presentation of stimuli at fixation, and are tasked with reporting two possible targets by first remembering them (storing in working memory). On some separations (“lags”) between the two targets, $T1$ and $T2$, the subject’s detection of $T1$ leads the subject to fail to detect $T2$. Engagement of attention in selecting $T1$ for encoding and report leads to failure to select $T2$: attention *blinks*.

These are a subset of attention paradigms, but they each exemplify the functional structure in the Jamesian conception.

This is a disparate lot, and unification is provided through a *functional frame*: correct task performance requires that the subject select a target to respond to as instructed. The intentional

Penultimate Version (please quote from forthcoming article).

action that subjects perform involves a specific mental selection in cognition or perception to guide instructed behavior, thus instantiating attention. There is a shared functional role exemplified across tasks. Just as we the folk direct attention by giving people instructions or commands in mundane situations, experimenters manipulate (set) attention through task instructions. Further research uncovers the implementation of such selection in neural processing.

Methodologically, the sufficient condition also provides the basic structure for new paradigms. Any paradigm has to ensure that the subject is paying attention appropriately. Experimenters design a new task with a desired target such that attention is fixed on it by having the subject respond to it in an instructed way. In a well-designed task, a necessary condition on correct task performance is that the subject selects the requisite target to respond to as per instructions. For example, visual spatial attention has been demonstrated in rodents, birds and bees (Eckstein et al., 2013; Sridharan et al., 2014; Wang & Krauzlis, 2018) by adapting the paradigms that are used with primates, namely spatial cueing. Note that bees do not have a neocortex, the target of much research on the neural basis of attention in primates. We attribute attention to bees because of their ability to respond selectively to a spatial location. Bees and primates might have different neural implementations of selection for task.

Reflecting on experimental design identifies a *normative* condition both in the task and in its implementation. In the implementation, a subject that fails to do the required selection, say due to intransigence, boredom, distraction, or a failure to understand instructions subverts the experiment by failing to pay appropriate attention. Similarly, a design that fails to ensure the appropriate setting of attention can be criticized in internal lab discussions or in the journal or grant review process. This leads to redesign of the experimental task.

Penultimate Version (please quote from forthcoming article).

In affirming the standard paradigms, any theorist of attention endorses the methodological principle that regulates the design of those paradigms. Extant scientific research is unified because we know that attention as studied in science instantiates the Jamesian conception expressed concretely in specific empirical sufficient conditions. Because the science is relevant to much philosophical reflection on attention, this part of selection for action should be the philosopher's *starting conception* of attention. After all, it is for cognitive science.

3. The Centrality of the Jamesian Conception in Science Defended

Denis Buehler (2022c) has argued that cognitive science rejects the empirical sufficient condition and so, the Jamesian conception (but see (Wu, 2024)).³ His argument has two steps: First, he provides behavioral evidence that there is selection for action in the absence of attention. He then connects this behavioral observation with the absence of what he calls a “neural criterion” for attention, arguing that the (inferred) absence of a specific neural signature strengthens the claim that attention is absent. Each step is incorrect.

The task he focuses on is the attentional blink. This is puzzling. To make empirical sense of the metaphor of attention *blinking*, the standard interpretation of the paradigm makes use of the Jamesian conception (see previous section). In attentional blink, we have to establish that the subject is visually attending to the first stimulus which in this case is the visual selection of a

³ Buehler has perhaps been my view's most vocal critic. While I will strenuously push back on his claims in this paper, I think his work on executive control and attention in many ways compatible with my approach to attention (see (Buehler, 2022b, 2022a)).

Penultimate Version (please quote from forthcoming article).

target, *T1*, to remember for report. This describes what visual attention in this task comes to. Sometimes, visual attention fails in regards to a second *T2* stimulus. What does attention failing (blinking) come to? Blinking is based on a postulated failure to store the *T2* stimulus in memory. Yet this must be failure of the type of selective attention deployed for *T1*. If so, attention as deployed on *T1* sometimes “blinks” on *T2* (the inference seems to assume that failure of selection entails failure of attention). So, attention is disrupted. These points depend on the Jamesian conception and the corresponding empirical sufficient condition.

Against this, Buehler claims that researchers deploying the attentional blink reject the related empirical sufficient condition. He focuses on a study by Joo et al. (2009) as a counterexample. Joo et al. use the attentional blink. Interestingly, in their experiment 2, they provide evidence that if the *T2* target consists of two arrays of circles (sets of objects) presented left and right of fixation where the subject must report which array has the larger mean size, the attentional blink does not arise. That is, subjects are able to accurately identify the array with larger mean size. Since this involves selection of a statistical property of the *T2* target for report, we have selection for report. Yet the escape from blinking occurs under conditions that normally induce an attentional blink when the selected target is at fixation. Buehler argues that the selection for task does not amount to attention, emphasizing Joo et al.’s conclusion: “evidently the refined stimulus information used for computing mean size remains available even in the absence of focused attention” (Joo et al. 2009, 1) (see Buehler (2022c, pp. 1678–1679)).

The invocation of *focused* attention should give pause. Joo et al.’s emphasis on focused attention follows a conventional dichotomy in the field where focused attention is contrasted with *distributed* attention. Accordingly, it is a *non sequitur* to move from Joo et al.’s claim that

Penultimate Version (please quote from forthcoming article).

focused attention is not involved at *T2* in experiment 2 to the claim that attention *simpliciter* is not involved. After all, distributed attention could be deployed. This is what Joo et al. claim:

In our work, people come into the laboratory with refined ability to perform mean set size judgments, with practice being unnecessary and error feedback providing little further improvement (Chong & Treisman, 2003). Moreover, people are best at this task when it is performed under a *distributed* attention mode *rather than with focused* attention (my emphasis).

I suggest that Buehler misreads the critical finding. Stated roughly: some forms of attention do not blink. In conditions where attention to localized targets blink (focused attention; see the control experiment, (Joo et al., 2009, p. 8)), attention to statistical properties (distributed attention) does not. Put more carefully, the study provides evidence that distributed attention as operationalized by selection of ensemble statistics to perform the instructed task (here, mean array size), is resistant to the attentional blink under the conditions instituted in their study that typically cause focused attention to blink. The contrast between distributed and focused attention is necessary for a correct interpretation of the work's significance.⁴

⁴ This renders problematic Buehler's appeal to other studies, for they are working with this contrast. Alvarez and Oliva (2008) speak of selecting "statistical summary features" outside of the "*focus* of attention". They build on the work of Ariely (2001) and Chong and Treisman (2003, 2005) on similar summary statistics specified in task instructions as targets for report. Buehler cites a review by Cavanaugh and Alvarez in support of his objection, yet in the first

Penultimate Version (please quote from forthcoming article).

Distributed attention is itself an interesting phenomenon deserving further philosophical reflection (see Adrienne Prettyman (2023) for a philosophical discussion). The contrast between focal and distributed might suggest that attention is like a spotlight (a limited resource) that can be localized or spread. That idea, driven by a common metaphor, requires defense and analysis. Note, however, that in the literature at issue, the focal/distributed distinction is operationalized in a specific way. Focal and distributed attention are fixed by different kinds of *targets* selected for task performance. Focal attention is directed at standard objects: localized shapes, letters, representations of middle-sized dry goods suitably sized (e.g., faces, chairs and so on). Distributed attention is directed at properties of sets of such objects such as the mean size of arrays in the Joo et al. paper. In other experiments, diffuse attention targets perceptual *gist*, say whether an image is of an urban or a natural scene (Alvarez & Oliva, 2008).⁵

sentence of their abstract, Cavanaugh and Alvarez emphasize that attention allows a subject to “select information from [objects and locations] for report and storage” (Cavanagh & Alvarez, 2005, p. 349) which is the selection for task deployed in the attentional blink. Cavanaugh and Alvarez begin with the Jamesian conception. We should not take them as contradicting themselves.

⁵ Noting work of Brosowsky et al. (2021), a referee asks whether implicit learning might be a potential counterexample: selection in implicit learning is not attention,. This large topic deserves sustained exploration that I cannot undertake here. I will note that Brosowsky et al. write: “there has been considerable debate as to whether implicit learning relies on attentional resources. Nevertheless, numerous results converge on the view that implicit learning does not require attentional resources. Explicit learning, however, requires executive processing and

Penultimate Version (please quote from forthcoming article).

It is important not to lose the forest for the trees. This section brings forward the centrality of the Jamesian conception instantiated in specific empirical sufficient conditions in understanding the significance of experimental data. The researchers cited do not misconstrue

attentional resources to guide attention, maintain task goals, and reduce interference” (3). Notice that the authors distinguish attention from resources guiding it. It is compatible with this perspective that selection for implicit learning might be *automatic* attention distinguished by being less affected or unaffected by the resource depletion that undercuts top-down attention in explicit learning. It would require detailed argument to show that this selection in implicit learning is not attention. This cannot be simply asserted on pain of begging the present question. Nor is the denial established just from the fact that top-down attention is not deployed in implicit versus explicit learning. There are other forms of attention.

The dialectic is similar to claims that habitual actions do not require attention. Carolyn Dicey Jennings (2020) claims that scientists studying habits “universally” deny attention in habitual actions. I submit that the data only clearly demonstrates that top-down or focused attention is absent. Selection for action in habits might be distinctively automatic (for defense, see (Bickel, 2024); for attentional habits, see (Jiang & Sisk, 2019)). In both implicit learning and habits, selection for action might be automatic attention. I do not claim that these remarks settle the issues, only that matters remain to be settled. Note that the textbook definition of attention as generically selection of information/targets for further processing would entail that we have attention in implicit selection for learning and for guiding habitual response.

Penultimate Version (please quote from forthcoming article).

that significance, for their tasks rely on distinguishing between focal attention, selection at typical objects like letters, and distributed attention, selection of statistics, each for memory and report. This underscores selection for task (action) as at the center of methodology and empirical inference.

What of Buehler's second step? While I have blocked his argument at its first step, I also dispute the second step in order to highlight the central role of empirical sufficient conditions in identifying relevant *neural correlates* of attention. Buehler considers the use of correlated electroencephalography (EEG) signatures during the attentional blink. He asserts that "Behavioral paradigms like the attentional blink pursue, roughly, the following strategy for studying...response-coupling [without attention]":

Subjects perform task *T* by response-coupling perceptual information *X*. Extensive neuroscientific research suggests that behavioral criteria for attentional selection are *merely* heuristic, and should be superseded by neuroscientific criterion NC. Performing task *T*, subjects do not meet NC; they hence exhibit inattentive perceptual response-coupling (Buehler, 2022c, p. 1687).

Buehler invokes work by (Dell'Acqua et al., 2006) as illustrative of this strategy. Where one has identified a neural criterion, in this case what is called the *N2pc* (*pc*: posterior contralateral) EEG signature, task structures are "merely heuristic". I find this principle odd, but the main problem is that it is based on a misreading of the significance of the N2pc. It cannot be a neural criterion in Buehler's sense, something that can be disconnected from the particularities of the attentional blink task as an experimental construct. This is because the N2pc is an intended artifact of a

Penultimate Version (please quote from forthcoming article).

particular way of implementing the attentional blink for the sake of collecting neural data in EEG.

Buehler's appeal to the N2pc is to *event related potentials* (ERPs) in EEG studies during the attentional blink. EEG provides a window on neural activity during attention. ERP measures focus on EEG signatures that are time locked to a specific event, say the presentation of a task stimulus. Buehler aims to show selection of a target for response that does not entail attention. Given his strategy, he aims to identify a neural criterion. He takes ERP studies to support his principle applied to the N2pc:

Research in cognitive neuroscience, such as the study by Dell'Acqua et al. [(Dell'Acqua et al., 2006)], illustrates a neuroscientific criterion for attentional selection. In our case, the criterion consists in the presence of the N2pc-component in the ERP-wave. This criterion is in principle independent of any behavioral criterion for deployments of attention (Buehler, 2022c, p. 1687).

The reason Buehler requires this is that he wants to strengthen his claim that in the Joo et al. experiment, attention is absent in selecting statistics to inform report. Given successful task performance, he concludes that there is inattentive response-coupling.⁶

⁶ Buehler makes a significant assumption in the absence of actual data, namely that the Dell'Aqua et al. results will carry over to the Joo et al. paradigm. Both studies use the attentional blink but with different stimuli. Buehler provides no empirical evidence for this claim.

Penultimate Version (please quote from forthcoming article).

Unfortunately, inferring from the absence of an N2pc in the attentional blink to the absence of attention is fallacious. The N2pc is a component of the N2 peak in ERP recordings, a negative deflection in the EEG signal. It is tied to posterior (back of the head) electrodes on the scalp. Buehler is correct that it is standardly interpreted as a correlate of attentional selection. A number of points, however, must be kept in mind. In ERP experiments, the electrical signatures are interpreted relative to an event *that is defined by the task structure*, typically the presentation of a task-relevant target that the subject must selectively respond to. There is no ERP without a relevant event. Second, a very specific task structure is crucial for the effect since the N2pc is a *relative* signal extracted by taking the difference between the response of distinct posterior electrodes namely those corresponding to the left and to the right visual field (see (Zivony & Lamy, 2021, Figure 1)). To extract an N2pc, EEG researchers probe the attentional blink by presenting two target streams, one in each visual hemifield. A contrast between the two visual hemifields is *necessary* to see an N2pc.

This means that Buehler's claim that in "performing task T [$T = \textit{Attentional blink}$], subjects do not meet NC [i.e. N2pc]; [subjects] hence exhibit inattentive perceptual response-coupling" is invalid (Buehler, 2022c, p. 1687). After all, there is a straightforward case where attention is deployed without an N2pc, namely performance of the attentional blink task with all stimuli presented at fixation. Here, given no lateralized difference in the visual stimuli, there can be no N2pc. This case, however, just is the standard implementation of the attentional blink (classically, (Chun & Potter, 1995; Raymond et al., 1992)). It follows, by Buehler's treating the N2pc as a neural "criterion", that there is no attention in any of these cases, just inattentive response-coupling. No attention blink researcher would accept that consequence since it undercuts the standard interpretation of attention blinking. Something has gone wrong.

Penultimate Version (please quote from forthcoming article).

If the N2pc as neural criterion renders the behavioral criteria merely heuristic, then these criteria, such as task structure, cannot be adduced to fix the problem. Buehler cannot say that the N2pc is a neural criterion for specific task implementations of the attentional blink. This is to concede that the N2pc *is* tied to a specific task structure. Accordingly, it would not then be a neural “criterion” of attention in Buehler’s sense. We should reject Buehler’s principle. The N2pc is relevant to understanding attention because of the specific task structure that is imposed to generate it. It is observable only under specific empirical sufficient conditions, namely where targets of selection are lateralized.⁷ This is hardly to take the task structure as “merely heuristic”.

Again, let us not lose the forest for the trees of technical details. There is a general point in the background. Given underdetermination, correlated neural data are interpreted within a task structure through the empirical sufficient condition. This structure is what allows us to identify certain neural correlates as putatively relevant to attention, given that they are correlated with implementation of the empirical sufficient condition in task performance. In general, for every task where neural correlates of attention are measured, those correlates are taken to be

⁷ A further reason that the N2pc will not serve Buehler’s purposes in talk of a neural criterion is that the N2pc is an electromagnetic signature based on a difference between the signal of scalp electrodes. It is a higher-order causal consequence of attention, not a causal factor on it. Buehler intends something stronger than neural correlate when he speaks of a criterion, but the N2pc is likely only a correlated effect not a constituent.

Penultimate Version (please quote from forthcoming article).

“attentional” because they are observed during task performance.⁸ This bridging between the neural and psychological is a powerful upshot of the relevant empirical sufficient conditions.

4. Problematic assumptions in debates about attention

Philosophers often have assumed that the selection for action account is easily refuted, yet as I argue in this section, their responses teeter between begging the question or making faulty, certainly questionable, assumptions about the epistemic standing of folk psychological judgements regarding the phenomena they adduce, specifically about *automatic* attention.

⁸ Peter Fazekas and Bence Nanay (2020) argue that the core realizer of attention is “the not-stimulus-triggered amplification of the presynaptic signals providing stimulus-specific inputs of normalization mechanisms is uniformly applicable to both lower- and higher-level perceptual processing” (p. 307) and that *amplification*, not selection, is central to attention. The neural data they invoke is collated by imposing the Jamesian conception in appropriate tasks, and the relevance of their proposal should be tailored to explain selection for task. They recognize that for amplification to be relevant to attention, it must be *selective* (see their section 5.3), yet they conclude, “it is hard to see why such selective allocation would render attention itself selective” (Fazekas & Nanay, 2020, p. 315). Yet every case of amplification that is a core realizer of a subject’s selection for task must also be a case of *selective* amplification. Otherwise, the amplification will not be a core realizer of *selective* task performance. I do not find it hard to see why *selective* amplification is also a core realizer, so why selection must be central to their account.

Penultimate Version (please quote from forthcoming article).

Automatic attention can be captured by the Jamesian conception, namely as when the mind is *taken possession of* by the target which drives the subject to deal with it. For example, a loud bang grabs auditory attention driving the subject to orient (turn) towards it.

The standard resistance to the sufficiency claim presents concrete cases of selection for action without attention. The absence of attention, however, cannot be merely asserted on pain of begging the question. It is not enough in the present case to appeal to the resources of folk psychology, say adducing how people speak about attention, their intuitions about it, and appeals to plausibility or naïve observations regarding it and thereby take the matter to be settled. When these appeals clash with claims in the sciences, we should not merely assume that the former have adequate epistemic weight. Given the empirical and theoretical work that validates automatic attention, we should validate the epistemic credentials of folk psychologically motivated counterclaims against such attention.

There are good reasons to question folk access to automatic attention versus intention-based attention. Consider the automaticity of visual attention as seen in saccadic eye movements that lead the agent to foveate targets within the visual field. Theorists of visual attention speak of this as *overt* attention. Subjects are estimated to make one to three saccades per second. Let's say you get seven hours of sleep a night (more or less). Given 17 waking hours, you make about 61 – 184,000 thousand eye movements each day. Many of these movements, moments in an agent's looking, are automatic. The agent is typically unaware of them. Where agents are aware of a specific eye movement, they often intend to make that eye movement, say when one is directed to hold one's gaze in an experiment. These intentional or top-down forms of attention are cases where agents can reliably track attention, but it is also clear to agents that they are attending. In contrast to intentional cases, we are largely unaware of automatic covert visual attention. How

Penultimate Version (please quote from forthcoming article).

many times have your eyes moved in the last two seconds? On what did your eye land?

Arguably, introspection is not reliable in detecting much of automatic perceptual attention and, given the lack of access to the bulk of such attention, we are not in a position to speak authoritatively about it.

Having raised questions about the availability of the required access to automatic attention, consider proposed counterexamples. Sebastian Watzl notes:

consider coming home and switching on the light in a familiar room. Again, you might be peripherally conscious of switching on the light (as well as being peripherally conscious of the switch). Furthermore, you are selecting the light switch as the target of your bodily action. But again, in a case like this, your act of switching on the light might be so automatic that your attention need not be focused on the light switch (Watzl, 2011, p. 154).

Watzl claims that in switching on the light, one has selected the light switch but attention is not directed to it. This would be a *non sequitur* if Watzl assumed that focused attention is the only form of attention, so I shall not read the response in that way. Unless questions are begged, focused attention is presumably top-down attention, and on this Watzl is right. But attention can be automatic, so the automaticity of the switching can indicate both the absence of focused attention and the presence of automatic attention. In any event, the case is underdescribed. Watzl would agree that the imagined reach is not random but directed at the appropriate spatial location. The bodily movements must be guided by an sufficiently accurate spatial representation of the switch's location, either a visual or mnemonic representation. If no such spatial

Penultimate Version (please quote from forthcoming article).

representation is involved, we don't have the requisite selection needed for the case to be a counterexample to our sufficient condition. But if a subject-level visual or mnemonic spatial representation is involved, why isn't this guidance due to the subject's visual or mnemonic attention, specifically an automatic form? This cannot just be denied without begging questions.⁹

What would show that attention is absent? Are there other empirical conceptions of attention that might provide neutral arbitration? Yet on the textbook characterization of attention in cognitive science, attention as selection of targets for further processing (see footnote 2), we *do* have attention, the subject's selecting a spatial location in vision or in memory for further processing. So, we have a clash between the philosopher's invocation of folk psychology and a

⁹ Henry Taylor (2015) suggests that a host of background beliefs are selectively involved in any type of behavior. Thus, this would mean that we have too much attention. The case imagined does not satisfy the antecedent of the relevant condition since attention entails access not just accessibility. The inputs to behavior must actually guide response and this means that relevant content in those inputs is responded to. Taylor speaks of the beliefs as "quiescent" and says, "these beliefs are surely not attended to when we perform these actions, rather they are entirely ignored" (45). Yet if they are ignored by the subject, then they are not selected for action. They are merely accessible. Only when the action must be explicitly "made sense of", say when the agent is queried for her reasons for acting, then the reasons are accessed in subsequent rationalization. Only then are they attended to.

Penultimate Version (please quote from forthcoming article).

fundamental assumption in science that allows for automatic attention. Since most of the challenges to selection for action presented in the literature involve automatic attention and given plausible doubts about reliably accurate subject access to automatic attention, the burden is on the opponents of selection for action to show that in their specific cases, there are good grounds to deny that automatic attention is present on the basis of the epistemic sources they rely on, be this intuition, folk psychological assumptions, or introspection. The accuracy of their denials of (automatic) attention's engagement cannot just be assumed.¹⁰

Cat Saint-Croix has suggested that:

when ascending a staircase, you perceive the distance between one stair and the next, and you select this perception for the action of placing your foot on the next step. This selection for action is evident in the fact that slight variations in these distances frequently cause missteps. But, we are almost never conscious of doing this, and certainly not with the degree of precision we in fact employ. So, it is reasonable to conclude that selection for action is not sufficient for attention (Saint-Croix, 2020, p. 506).

I agree that in this case, we do not notice our access to spatial information to which we are visually tuned to guide behavior. This selection, however, is automatic. Saint-Croix's argument

¹⁰ Buehler in his presentation of counterexamples consistently avers that the presence of attention is "intuitive", "plausible" and "obvious". In my (2018), I respond to Buehler's earlier work in a similar way as I do here (Buehler, 2018) as well as to criticisms by Carolyn Dicey Jennings and Bence Nanay (Jennings & Nanay, 2016).

Penultimate Version (please quote from forthcoming article).

seems to be that something counts as attention only if we are conscious of it to some degree. I reject this principle as inconsistent with the empirical theory of attention. Much automatic attention goes unnoticed as the eye movement case demonstrates.

Philosophers assert in their counterexamples that a subject selects for action yet does not attend. They must show that the epistemic sources that they take to support their claims about attention provide reliably accurate access to attention or its absence in the specific cases invoked. Specifically, they must show that they are not simply failing to recognize automatic attention and, to the extent they do note it, that their sources are accurate in ruling it out. At the same time, the science that they often rely on in other contexts endorses the principle they reject. That is the crucial point, that in the attempts to provide counterexamples, we lose site of the forest, namely the unification afforded by the Jamesian conception in linking empirical work across all levels of analysis and in anchoring experimental methodology.

If a principle such as the Jamesian conception and its entailments in empirical sufficient conditions play a central theoretical and methodological role, sets a regulative principle for interpreting behavioral and neural data, showing the relevance of neural correlates like the N2pc in ERPs, provides a norm for the design of experimental paradigms, constitutes a functional principle that ties together multiple levels of analysis in behavior and neurobiology, and has a reliably accurate folk-psychological source in awareness in intentional action (James' famous description), then these structural features should command our theoretical attention as a *starting point*. Ethicists and epistemologists need not be neutral about attention's nature. They can use the Jamesian conception.

5. The Ethics of Attention

Penultimate Version (please quote from forthcoming article).

Can there be harmful surpluses of attention? The question presupposes a clear notion of an attentional surplus or deficiency. Such claims can sometimes be overly metaphorical, so I will anchor the idea to a common approach in psychology that identifies attention with some limited resource *R*. The idea of attention as a limited, gradable resource, while common in empirical and folk psychology, has not been given sustained philosophical analysis that is psychologically grounded. Since selection for action provides the starting place for the science of attention, in this section I deploy it to explicate gradability in respect of attention as revealed in harms tied to the amount of attention. As selection for action is central to empirical methodology, it provides a psychological basis for the analysis. Specifically, it opens up the possibility that resulting philosophical posits can be measured and manipulated.

Ella Whiteley has compellingly argued that there are harms due to a “relative attentional surplus on the wrong property” (Whiteley, 2023, p. 506). Rebecca Buxton and Lisa Whiting emphasized that “a result of [the association between maleness and philosophy] is that women are often remembered as women first: they are seen more as women than they’re seen as philosophers....we forget that they are principally philosophers” (Buxton & Whiting, 2021). Whiteley notes that in certain contexts, focusing on women qua women, we pay attention to the wrong property leading to a *relative* surplus of attention to gender as opposed to occupation that can lead to specific harms. We pay relatively *more* attention to women as women and relatively *less* attention to them as philosophers. Yet how should we understand an attentional surplus? In what sense does attention allow for such quantification?

Penultimate Version (please quote from forthcoming article).

Whiteley also appeals to a *pattern* or *structure* of attention. Patterns and surpluses are bridged through *salience*. Patterns of attention make targets of attention (more or less) salient where this can involve a problematic surplus of attention. Patterns of attention are explained through the structuring of contents, “how properties are foregrounded and backgrounded in our attention by giving some relative salience over others” (Whiteley, 2022, p. 194). An attentional pattern can also involve a specific structural organization of one’s psychology and consciousness. So, a subject’s visual experience of a scene could exhibit different attentional patterns to the extent that patterns *A* and *B* differ in what visible item is at the center of one’s consciousness, what is “foregrounded” rather than “backgrounded”. An object *a* can be more central or more foregrounded than *b* or vice versa (cf. Watzl’s (2017) priority structures). Perhaps Whiteley is relying on introspection, but we should not assume that the introspection of a phenomenal foreground versus background is accurate or uncontroversial (cf. the conflict between **my** and Watzl’s introspection **of** such phenomenology (Wu, 2014, Chapter 4.5)).

Yet why should a pattern or structure *on its own* be ethically significant? Why does putting something in a center of a structure make it problematic? As Whiteley notes in discussing an objection to her appeal to the moral significance of attentional patterns centered on a person’s race: “Attending primarily to a person’s...racial identity cannot be morally problematic, so goes this critique, since there are many contexts in which this pattern of attention is not wrongful or harmful” (Whiteley, 2023, p. 515). For example, the same priority structure centered on race can occur in a racist context or it can occur in a discussion group involved in a constructive discussion on race. Whiteley responds: “While an attentional pattern that makes an individual’s (e.g.) racial identity (or identity as a person who has experienced rape, in the other case) their most salient feature is a red flag, indicating a potentially problematic form of attention, I argued

Penultimate Version (please quote from forthcoming article).

that information about the particular case and context is needed to be able to judge whether that form of attention is indeed morally problematic” (Whiteley, 2023, p. 516). Yet this seems to concede in a basic case that attention as such is not what matters. The same pattern or amount of attention can be good or bad, depending on context.

How might we bridge talk of structures and surpluses? Theorists often rely, even if implicitly, on a familiar metaphor, the spotlight. Now, imagine that we had a visual scene or an array of thoughts in working memory (**say** a representational structure of these **targets**). Take the “light” of attention, a limited and gradable resource, and apply it to the representation of visibilia or thoughts to render them more or less salient. After we apply the light of attention, we end up with a pattern, a structure of differentially illuminated objects. We could, of course, have applied the light in a different way or with different rules of application. The metaphor shows that a link between resource and pattern/structure is coherent, yet we gain no psychological understanding nor do we see how harms arise from this distribution. If the light, or whatever resource one invokes as attention, is applied to my gender or my race, why should that harm me? The picture at best describes what is to be explained.

Whiteley’s discussion is important but leaves an open question. How should we understand the appeal to a surplus or deficiency of attention as a psychological phenomenon? We should not leave it unanalyzed, merely describing it or capturing it with suggestive metaphors. In building on Whiteley’s discussion, I shall deploy selection for action to explain relevant gradability tied to attention. Selection for action provides psychological grounding for the analysis.

Penultimate Version (please quote from forthcoming article).

Some might assume, however, that there is a simple way to capture surpluses, namely by appeal to attention as itself a limited resource. I believe that this should not be the first option. We can agree that the appeal to resources is natural when we confront gradable *behavior*. For example, in psychology

(a) Performance in attention tasks (action) is gradable along various dimensions.

This is the data. It can be explained by what we can call the *identity* view:

(b) Identity: Attention is a limited (gradable) resource, R , that can be distributed and exhausted.

While the identity claim might seem obvious, Odmarr Neumann (1987) admonished us to not move immediately from (a) to (b). That inference is problematic given a second explanation:

(c) Causal dependence: Attention's selective operations are modulated by a limited (gradable) resource, R , that can be distributed and exhausted (e.g. energy).

Note that both (b) and (c) appeal to the gradability of resource R to explain the gradability of performance in (a). In principle, they can each explain equally well the data. Arguably, however, (c) is weaker since it appeals only to causal dependence rather than identity. The selection for

Penultimate Version (please quote from forthcoming article).

action view opts for (c), and so in this section, I use it to explain harmful effects of surplus and deficiency regarding attention.

Given two viable explanations of (a), (b) is not obligatory. Indeed, the default empirical explanation seems to be (c). Consider as illustrative, Donald Broadbent's (1958) seminal filtering theory of attention. Broadbent was keen to explain certain limitations in performance in the dichotic listening paradigm where given the subject's parroting a verbal stream played in one ear (selecting heard words for utterance), subjects were not able to remember the higher level properties of the voice in the other ear such as expressed semantics even as they could recall the pitch and timbre of the voice. Broadbent's explanation of this limitation in performance invoked a concrete limited resource, *informational channel capacity*, which can be precisely quantified in Shannon information theory. Given limited channel capacity, Broadbent explained the decline in performance by positing that the attentional filtering must act *early* in processing. Others argued that filtering acts late in processing, early processing having sufficient capacity to process all stimuli (Deutsch & Deutsch, 1963). Channel capacity, a limited resource, explains the stage of attention's operation. Either way, Broadbent's filter (his attention) is not a gradable resource. Rather, its operations depend on the availability of an informational resource.¹¹

¹¹ In the Load Theory of Attention (better of attentional performance), Nilli Lavie and co-workers (Lavie et al., 2004) take perceptual and cognitive load to *modulate* attention. Attention is *not* load. So, they endorse causality, not identity.

Penultimate Version (please quote from forthcoming article).

Is there a good empirical argument favoring (b) over (c)? Daniel Kahneman influentially affirmed (b) in identifying attention with effort. He motivated the resource position as follows:

there is more to attention than mere selection. In everyday language, the term "attention" also refers to an aspect of amount and intensity. The dictionary tells us that to attend is to apply oneself presumably to some task or activity. Selection is implied, because there are always alternative activities in which one could engage, but any schoolboy knows that applying oneself is a matter of degree. Lulled into a pleasant state of drowsiness by his teacher's voice, the schoolboy does not merely fail to pay attention to what the teacher says; he has less attention to pay. A schoolboy who reads a detective story while his teacher speaks is guilty of improper selection. On the other hand, the drowsy schoolboy merely suffers from, or perhaps enjoys, a generally low level of attention (Kahneman, 1973, p. 3).

Kahneman moves beyond attention as selection to attention as a resource by invoking language use and folk observation. Yet these sources are likely based on the acknowledged data in (a), namely performance limitations, and how to explain (a) is precisely what is at issue. To describe the sleepy student, we can capture their decline in performance by speaking of their paying less attention (a resource), yet that need only indicate that folk psychology has chosen (b) to explain

Penultimate Version (please quote from forthcoming article).

(a). This begs the question against (c), so folk psychological data does not provide a distinct argument that shows the explanatory priority of (b) over (c).

My point is not to deny that attention is a resource. Instead, I push back against a natural assumption, that since performance is gradable, correlated attention must itself be gradable, say a limited resource. Yet there has always been an alternative explanation that appeals to selection and its draw on distinct gradable resources. Since scientists of attention are methodologically committed to selection for action, then they are already assuming it to explain performance. Arguably, for them, (c) should be (and has often been) the starting point for explaining the gradability of performance. It should then be ours as well.¹²

¹² Are there cases where the resource perspective is more natural? Might the result in Joo et al. discussed in Section 3 be construed as involving more or less attention across the arrays? One could interpret the results in this way since (b) can explain gradable performance as (c). That said, Joo et al. do not take attention to be a resource. Rather, they take focal and distributed attention to differentially draw on limited “attentional” resources where in their task, distributed attention is not affected by resource depletion in the same way that focal attention is. Thus, they opt for (c) over (b) (causal dependence is also deployed in distinguishing between internal versus external (memory versus perceptual) attention in (Kiyonaga & Egner, 2013)). A resource theorist can reinterpret Joo et al., but the reinterpretation, in clashing with the perspective I am arguing for here, would require defense.

Penultimate Version (please quote from forthcoming article).

Multitasking might provide a case that naturally entails a resource (identity) approach, for one task might be primary, the other secondary, and so, the former apportioned more attention than the latter. What behavioral differences would result? For sake of argument, here's one (inspired by (Frankfurt, 1978)): if the primary task were perturbed, the agent's corrective response would be faster than if the secondary task were perturbed. That is, performance is graded. The resource theorist explains this by saying that more attention is given to the primary task. Alternatively, a selection theorist proposes that attentional selection is split in that there are two distinct selections for task, each drawing differentially on the same limited resource. For example, some psychological theories take working memory as what controls attention and that working memory capacity, which is a limited resource, is a measure of attentional control (Engle, 2002, 2018). So, difference of performance across individuals is correlated with measured working memory capacity. Talk of the distribution of attention as a resource is a description waiting for precise analysis. In contrast, the case of differential control on attentional selection has a quantitative aspect in working memory capacity. The control of attention grades with working memory capacity. This is to work with causal dependence on a limited resource (for a defense of working memory as a basis for understanding control of attention, see my (Wu, 2023)).

Attention as resource is invoked in the psychology of reading such as in (Snell et al., 2018)).

Visual attention and reading is a very large topic that I cannot take up here, but I make some preliminary remarks. Even if a resource view in reading holds, it is an open question whether that conception will explicate Whiteley's notion of a surplus in cognitive cases. A second issue is empirical: given that (b) and (c) both can explain gradable performance, have psychologists of

reading given a compelling argument that favors identity over the causal dependence or have they, instead, just assumed identity as Kahneman did? Snell et al. (2018) note that on a parallel processing model of reading, “the attentional distribution would follow a Gaussian shape centered on the attentional focus, such that while processing of the fixated word would normally be the strongest, surrounding words and letters may nonetheless exert some influence” (1). They ground the distribution account in a model of saccade generation in reading by Engbert and colleagues (Engbert et al., 2002, 2005). A core feature of that model is *spatially distributed processing* of words at a fixation, implemented in a dynamic field of activations to the words in a spatial window. This distributed activation is analogized to a *saliency map* (Engbert et al., 2005, p. 782). The saliency map explains visual attention, and activations on the map are conceived as graded. Yet in typical models, saliency maps are inputs into a winner-take-all mechanism that *sets attention*: attention moves to the location of highest saliency (Koch & Ullman, 1985). Thus, the output is selection at a location. Accordingly, the gradability of saliency values does not entail gradability of attention. Instead, saliency explains shifts of attentional selection to specific locations. If we take saliency maps to be a gradable factor influencing selection, namely a dynamic spatial field of activations, this comports with (c), not (b). On its face, a resource reading is not obligatory. There is obviously more that needs to be said, but developing the case must examine whether the resource view provides a better account than causal dependence as an explanation of the data. That said, we need not settle this issue to show that selection for action provides an analysis that explicates Whiteley’s insights about harms. I thank a referee for pressing these cases. The referee does note a worry about an “attention explosion”, namely that

Penultimate Version (please quote from forthcoming article).

With that in mind, let us transpose Whiteley's insights about the harms of attention to the weaker key provided by selection for action.¹³ First, drawing on attention's link to action, a pattern of attention is assessed by understanding what it is *for*, where the response attention guides is part of understanding what attention is, functionally speaking. So the mere mental selection of someone's gender need not be harmful but the selection of it for a specific response can be. Focusing on a woman's gender for the purposes of a joint project of writing a paper on

selection for action will license a proliferation of attention. I am less concerned, but this depends on a concrete explication of *action* in selection for action. On this, see (Wu, 2023).

¹³ As Whiteley invokes Watzl's (2017) account, let me make some critical remarks (see also (Wu, 2019)). Arguably, Watzl's account is an instance of selection for action. The agent's activity in attending is to regulate priority structures comprised of mental states. Watzl conceives of a salient fire alarm as commanding the agent to put the alarm at the top of a priority structure. The transition is from the prior mental structure where the alarm is low priority to the new one where it is high priority. Regulation of priority is set within agency, yet why should attention be *limited* to just this form of selection for action? The view seems overly restrictive. Second, Watzl's formal definitions of priority allow for degrees of attention. He criticizes selection for action as unable to explain degrees since selection does not come in degrees (Watzl, 2017, p. 112). Here, he presupposes, but does not argue, that attention is itself gradable (e.g. attention is a resource). His formal definitions of various notions of priority provide a way to more precisely characterize degrees of attention, but they do not show that we should favor the resource account over causal dependence on a resource.

Penultimate Version (please quote from forthcoming article).

gender, say one that uses her experiences as a starting point, need not be harmful. Yet said focus for a different purpose can be harmful, say in informing explicitly or implicitly biased negative responses. The norms of attention can be explained through the norms of action. This is compatible with Whiteley's own response to queries about the normative significance of patterns and structures, but it draws on attention as selection for action, exemplified in the situations Whiteley considers.

Now turn to the idea of a surplus of attention, central to Whiteley's diagnosis of attentional harms, drawing on attention as selection for action. I approach quantifying attention in terms of a diachronic perspective.¹⁴ If attention serves action, then like action, attention has a temporal dimension (time is a gradable, limited resource in action). Attention can be quantified in at least two temporal dimensions. First, one can spend too much time attending to something in order to inform a specific response, say continued deliberation on someone's gender or race in thought over time, i.e. *duration*. Second, one can carry out too many *instances* of the action

¹⁴ Does Whiteley prefer a synchronic approach? In Buxton and Whiting's example, we have a synchronic perspective: We remember that the agent is a woman, and simultaneously, forget that she is a philosopher. But this is just selection of one feature with failure of selection of another in recollection. We don't need further grades of attention to understand the case, and it seems odd to speak of a subject, at a time, remembering *more* that an agent is a woman and remembering less (forgetting more?) that she is a philosopher. I leave development of a synchronic approach open, but the availability of a diachronic approach deepens our understanding of Whiteley's claims.

Penultimate Version (please quote from forthcoming article).

which attention serves, such as one's looking at someone's body multiple times for some purpose, say in a medical context or in something harmful like leering.

So consider *saliency* understood in this way: T is salient to S in time t if T is attended to.

Using the Jamesian conception, we can recast this as

T is salient to S at temporal range t if T is selected by S to mentally guide response R at t .

To act is to generate R guided by selection of T . So, if S stares at someone's chest during a medical exam, then S visually guides clinical assessment by visually selecting T to inform judgment over time. The nature of saliency in terms of harm can, in the first instance, be understood in terms of the action at issue, individuated in part by the response R . We can now quantify attention. First, we can measure the amount of attention by the amount of time attention is deployed, namely measuring the value of t . For example, how long does one think about an agent's gender relative to her being a philosopher? Is one *fixated* on the issue, focusing on it for long periods of time? Second, we can quantify the amount of attention by the number of separate instances within a stretch of time τ . For example, how often does one mention or think about the agent's gender relative to mentioning or thinking about her being a philosopher?

So, consider looking which need not, in itself, be harmful. Say S is trying to change their lecherous visual tendency and by habit, looks at a person's chest. S might avert their eyes immediately away in horror, realizing what they've done. What attention is *for* matters. If we think harms are also tied to quantity (a "surplus"), we can probe this in the ways noted. S 's inappropriate gaze might linger too long (duration) or might keep habitually happening during the course of time of a conversation (instances). We can relativize assessment by comparing this

Penultimate Version (please quote from forthcoming article).

case to past cases of *S*'s visual behavior with respect to durations or instances (for example, assessing if *S*'s attempts at debiasing their gaze is improving). These notions can then be used to actually measure attention. This quantification is made available by the Jamesian conception, perhaps not surprisingly since that conception grounds experimental investigation of attention where measurement of performance is possible.

The temporal dimensions need not be the only ones to consider. Another quantitative dimension can be extracted in Whiteley's noting that the

form of attentional fixation might be the extreme excess of attention many people with visible disabilities receive; one testimony from a social worker with cerebral palsy compares their experience of being disabled to being a celebrity: "Because being famous—I assume without being famous myself means that you receive a high level of attention whether you like it or not and in every possible situation."

First, apply the previous dimensions to this description: (a) that attention will harm in this case depending on the associated response that attention informs, (b) that unvarnished quantitative talk of attention can be more precisely explained in terms of duration and (c) that talk of attention in "every possible situation" can be understood in terms of instances of attention over stretches of time. That said, a fourth dimension is suggested by talk of "every possible situation."

This brings us back to contexts. For an additional quantitative dimension can be found in individuating types of situations that are ethically relevant, the set of normatively assessable contexts for certain forms of attention. We can then think of instances of attention relative to specific instantiations of context types. Contexts will be partly individuated by the relevant

Penultimate Version (please quote from forthcoming article).

response informed by attention. For example, we can speak of a subject's paying more attention in terms of the frequency of deploying attention instances within problematic contexts relative to neutral or positive contexts.

What appeal to the empirical (Jamesian) conception of attention has facilitated is recognition of four dimensions for which there can be complex interactions in the assessment of attentional harms: (1) the relevant response that is central to attention's guiding function, (2) attention's duration, (3) its instances, and (4) relevant contexts within which it is exercised, this being partly individuated by (1). This is an advance in the analysis of gradability talk regarding attention. These are not the only dimensions that can be relevant, but in making clear that attention brings with it multi-dimensional axes of potential normative significance, we provide a concrete basis for evaluating harms. What the analysis also shows is that the *interaction* between these factors will be crucial to understanding the good and ill of attention.

This returns us to the harm that Whiteley captures by speaking of woman first, philosopher second. To describe one natural and familiar case, consider a philosopher, f who is salient to a community of philosophers P in that some members of P attend to f in professional contexts C by selecting f 's being a woman to inform professional response, say when filling a slate of conference speakers. One might argue that for the relevant responses in C , f 's being a woman is not a relevant feature. Rather, it is her being a philosopher of X that should come first (be attended to). If there is a harm in an attentional surplus, we can, for example, quantify this over countable professional context instances c_i , and in doing so make literal talk of a surplus of attention. For example, across contexts of choosing speakers, the agent's being a woman, and not her being a philosopher, is the basis for selecting her as a conference speaker. With this in place, we can concretely investigate the precise harms that attention brings. The Jamesian conception

Penultimate Version (please quote from forthcoming article).

deployed in the analysis will also allow us to engage scientific research on attention to deepen our psychological understanding of the mental basis of the noted harms. Why do people recollect her being a woman and not her being a philosopher? In this way, the Jamesian conception in the weakened versions of selection for action advances philosophical discussion.

Let me conclude with a pointed challenge to those who invoke a notion of more or less attention as part of philosophical psychology: that idea, while intuitive, does not have a clear empirical grounding beyond the data of gradable performance. If the idea of more or less attention matters to one's arguments, then I have put forward a concrete, empirically grounded set of dimensions for gradability. To ignore it would, in respect of serious psychology, be to fall back on intuition and metaphor.

6. Conclusion

The goal of this essay was to argue that the welcome resurgence of interest in attention in ethics and epistemology should take a stand on what attention is to ensure an adequate philosophical psychology. A weakened version of selection for action was put forward given the Jamesian roots of empirical paradigms for studying attention. This principle was defended against a host of objections, and its theoretical utility in considering attentional harms was demonstrated by providing an analysis of talk of a surplus of attention in concrete, quantifiable dimensions.¹⁵

References

¹⁵ Acknowledgments: I am grateful to Alon Zivony for many discussions about the science and metaphysics of attention including at a Vision Sciences Society 2024 symposium we contributed to, to Ella Whiteley for discussing her views with me, and to two anonymous referees for their extremely helpful and constructive comments.

Penultimate Version (please quote from forthcoming article).

Alvarez, G. A., & Oliva, A. (2008). The Representation of Simple Ensemble Visual Features Outside the Focus of Attention. *Psychological Science*, *19*(4), 392–398.

<https://doi.org/10.1111/j.1467-9280.2008.02098.x>

Archer, S. (Ed.). (2022). *Saliience: A Philosophical Inquiry*. Routledge.

<https://doi.org/10.4324/97811351202114>

Ariely, D. (2001). Seeing Sets: Representation by Statistical Properties. *Psychological Science*, *12*(2), 157–162. <https://doi.org/10.1111/1467-9280.00327>

Bahmani, Z., Clark, K., Merrikhi, Y., Mueller, A., Pettine, W., Isabel Vanegas, M., Moore, T., & Noudoost, B. (2019). Prefrontal Contributions to Attention and Working Memory. In T. Hodgson (Ed.), *Processes of Visuospatial Attention and Working Memory* (pp. 129–153). Springer International Publishing. https://doi.org/10.1007/7854_2018_74

Bear, M., Connors, B., & Paradiso, M. A. (2020). *Neuroscience: Exploring the Brain, Enhanced Edition: Exploring the Brain, Enhanced Edition*. Jones & Bartlett Learning.

Bickel, L. (2024). Why the performance of habit requires attention. *Mind & Language*, *39*(2), 260–270. <https://doi.org/10.1111/mila.12467>

Broadbent, D. E. (1958). *Perception and communication*. Pergamon Press.

Brosowsky, N. P., Murray, S., Schooler, J. W., & Seli, P. (2021). Attention need not always apply: Mind wandering impedes explicit but not implicit sequence learning. *Cognition*, *209*, 104530. <https://doi.org/10.1016/j.cognition.2020.104530>

Buehler, D. (2018). A Dilemma for ‘Selection-for-Action.’ *Thought: A Journal of Philosophy*, *7*(2), 139–149. <https://doi.org/10.1002/tht3.378>

Penultimate Version (please quote from forthcoming article).

Buehler, D. (2022a). Agential capacities: A capacity to guide. *Philosophical Studies*, 179(1), 21–47. <https://doi.org/10.1007/s11098-021-01649-6>

Buehler, D. (2022b). Explicating Agency: The Case of Visual Attention. *The Philosophical Quarterly*, pqac034. <https://doi.org/10.1093/pq/pqac034>

Buehler, D. (2022c). Seeing Circles: Inattentive Response-Coupling. *Ergo an Open Access Journal of Philosophy*, 9(62), Article 0. <https://doi.org/10.3998/ergo.3587>

Buxton, R., & Whiting, L. (2021, February 4). *Women or Philosophers? - The Philosophers' Magazine*. <https://www.philosophersmag.com/essays/230-women-or-philosophers>

Campbell, J. (2002). *Reference and Consciousness*. Oxford University Press.

Cavanagh, P., & Alvarez, G. A. (2005). Tracking multiple targets with multifocal attention. *Trends in Cognitive Sciences*, 9(7), 349–354.

Chica, A. B., Bartolomeo, P., & Lupiáñez, J. (2013). Two cognitive and neural systems for endogenous and exogenous spatial attention. *Behavioural Brain Research*, 237, 107–123. <https://doi.org/10.1016/j.bbr.2012.09.027>

Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision Research*, 43(4), 393–404. [https://doi.org/10.1016/S0042-6989\(02\)00596-5](https://doi.org/10.1016/S0042-6989(02)00596-5)

Chong, S. C., & Treisman, A. (2005). Attentional spread in the statistical processing of visual displays. *Perception & Psychophysics*, 67(1), 1–13. <https://doi.org/10.3758/BF03195009>

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, 21(1), 109.

Penultimate Version (please quote from forthcoming article).

Dell'Acqua, R., Sessa, P., Jolicœur, P., & Robitaille, N. (2006). Spatial attention freezes during the attention blink. *Psychophysiology*, *43*(4), 394–400. <https://doi.org/10.1111/j.1469-8986.2006.00411.x>

Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review*, *70*(1), 80–90.

Eckstein, M. P., Mack, S. C., Liston, D. B., Bogush, L., Menzel, R., & Krauzlis, R. J. (2013). Rethinking human visual attention: Spatial cueing effects and optimality of decisions by honeybees, monkeys and humans. *Vision Research*, *85*, 5–19. <https://doi.org/10.1016/j.visres.2012.12.011>

Engbert, R., Longtin, A., & Kliegl, R. (2002). A dynamical model of saccade generation in reading based on spatially distributed lexical processing. *Vision Research*, *42*(5), 621–636. [https://doi.org/10.1016/S0042-6989\(01\)00301-7](https://doi.org/10.1016/S0042-6989(01)00301-7)

Engbert, R., Nuthmann, A., Richter, E. M., & Kliegl, R. (2005). SWIFT: A Dynamical Model of Saccade Generation During Reading: Psychological Review. *Psychological Review*, *112*(4), 777–813. <https://doi.org/10.1037/0033-295X.112.4.777>

Engle, R. W. (2002). Working Memory Capacity as Executive Attention. *Current Directions in Psychological Science*, *11*(1), 19–23. <https://doi.org/10.1111/1467-8721.00160>

Engle, R. W. (2018). Working Memory and Executive Attention: A Revisit. *Perspectives on Psychological Science*, *13*(2), 190–193. <https://doi.org/10.1177/1745691617720478>

Fairweather, A., & Montemayor, C. (2017). *Knowledge, Dexterity, and Attention: A Theory of Epistemic Agency*. Cambridge University Press.

Penultimate Version (please quote from forthcoming article).

Fazekas, P., & Nanay, B. (2020). Attention Is Amplification, Not Selection. *The British Journal for the Philosophy of Science*, 72(1), 299–324. <https://doi.org/10.1093/bjps/axy065>

Frankfurt, H. G. (1978). The Problem of Action. *American Philosophical Quarterly*, 15(2), 157–162.

Gardiner, G. (2022). Attunement: On the Cognitive Virtues of Attention. In M. Alfano, J. de Ridder, & C. Klein (Eds.), *Social Virtue Epistemology* (pp. 48–72). Routledge.

Holt, N., Bremner, A., Sutherland, E., Vliek, M., Passer, M., & Smith, R. (2019). *Psychology: The Science of Mind and Behaviour* (Fourth). McGraw Hill.

James, W. (1890). *The Principles of Psychology, Volume 1*. Henry Holt and Co.

Jennings, C. D. (2020). *The Attending Mind*. Cambridge University Press.

Jennings, C. D., & Nanay, B. (2016). Action without attention. *Analysis*, 76(1), 29–37. <https://doi.org/10.1093/analys/anu096>

Jiang, Y. V., & Sisk, C. A. (2019). Habit-like attention. *Current Opinion in Psychology*, 29, 65–70. <https://doi.org/10.1016/j.copsyc.2018.11.014>

Joo, S. J., Shin, K., Chong, S. C., & Blake, R. (2009). On the nature of the stimulus information necessary for estimating mean size of visual arrays. *Journal of Vision*, 9(9), 7.

<https://doi.org/10.1167/9.9.7>

Kahneman, D. (1973). *Attention and Effort*. Prentice-Hall.

Kiyonaga, A., & Egnér, T. (2013). Working memory as internal attention: Toward an integrative account of internal and external selection processes. *Psychonomic Bulletin & Review*, 20(2), 228–242. <https://doi.org/10.3758/s13423-012-0359-y>

Penultimate Version (please quote from forthcoming article).

Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, 4(4), 219–227.

Koester, J. D., Mack, S. H., & Siegelbaum, S. A. (Eds.). (2021). Visual Processing for Attention and Action. In E. R. Kandel (Ed.), *Principles of Neural Science* (6th ed.). McGraw Hill.
neurology.mhmedical.com/content.aspx?aid=1180642180

Lavie, N. N., Hirst, A., de Fockert, J. W., & Viding, E. (2004). Load Theory of Selective Attention and Cognitive Control. *Journal of Experimental Psychology General*, 133(3), 339–354.

Mole, C. (2011). *Attention is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford University Press.

Mole, C. (2024). Emancipatory Attention. *Philosophers' Imprint*, 24(1), Article 1.
<https://doi.org/10.3998/phimp.3163>

Moore, T., & Zirnsak, M. (2017). Neural Mechanisms of Selective Visual Attention. *Annual Review of Psychology*, 68(1), 47–72. <https://doi.org/10.1146/annurev-psych-122414-033400>

Munton, J. (2023). Prejudice as the misattribution of salience. *Analytic Philosophy*, 64(1), 1–19.
<https://doi.org/10.1111/phib.12250>

Neumann, O. (1987). Beyond Capacity: A Functional View of Attention. In H. Heuer & A. F. Sanders (Eds.), *Perspectives on Perception and Action* (pp. 361–394). Lawrence Erlbaum Associates, Publishers.

Panizza, S. C. (2022). *The Ethics of Attention: Engaging the Real with Iris Murdoch and Simone Weil*. Routledge. <https://doi.org/10.4324/9781003164852>

Penultimate Version (please quote from forthcoming article).

Prettyman, A. (2023). What is diffuse attention? *Mind & Language*, *38*(2), 374–393.

<https://doi.org/10.1111/mila.12365>

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, *18*(3), 849–860. <https://doi.org/10.1037/0096-1523.18.3.849>

Rungratsameetaweemana, N., & Serences, J. T. (2019). Dissociating the impact of attention and expectation on early sensory processing. *Current Opinion in Psychology*, *29*, 181–186. <https://doi.org/10.1016/j.copsyc.2019.03.014>

Saint-Croix, C. (2020). Privilege and Position. *Res Philosophica*, *97*(4), 489–524.

<https://doi.org/10.11612/resphil.1953>

Saint-Croix, C. (forthcoming). The Epistemology of Attention. In K. Sylvan, E. Sosa, J. Dancy, & M. Steup (Eds.), *The Blackwell Companion to Epistemology*, 3rd edition. Wiley Blackwell.

Sani, I., Stemmann, H., Caron, B., Bullock, D., Stemmler, T., Fahle, M., Pestilli, F., & Freiwald, W.

A. (2021). The human endogenous attentional control network includes a ventro-temporal cortical node. *Nature Communications*, *12*(1), Article 1.

<https://doi.org/10.1038/s41467-020-20583-5>

Shulman, G. L., & Corbetta, Maurizio. (2012). Two attentional networks: Identification and function within a larger cognitive architecture. In M. I. Posner & L. Huang (Eds.), *Cognitive Neuroscience of Attention*, Second Edition (pp. 113–128). Guilford Press.

Penultimate Version (please quote from forthcoming article).

- Snell, J., Mathôt, S., Mirault, J., & Grainger, J. (2018). Parallel graded attention in reading: A pupillometric study. *Scientific Reports*, *8*(1), 3743. <https://doi.org/10.1038/s41598-018-22138-7>
- Sridharan, D., Ramamurthy, D. L., Schwarz, J. S., & Knudsen, E. I. (2014). Visuospatial selective attention in chickens. *Proceedings of the National Academy of Sciences*, *111*(19), E2056–E2065. <https://doi.org/10.1073/pnas.1316824111>
- Taylor, J. H. (2015). Against Unifying Accounts of Attention. *Erkenntnis*, *80*(1), 39–56. <https://doi.org/10.1007/s10670-014-9611-3>
- Wang, L., & Krauzlis, R. J. (2018). Visual Selective Attention in Mice. *Current Biology*, *28*(5), 676–685.e4. <https://doi.org/10.1016/j.cub.2018.01.038>
- Watzl, S. (2011). Attention as Structuring of the Stream of Consciousness. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 145–173). Oxford University Press.
- Watzl, S. (2017). *Structuring Mind: The Nature of Attention and how it Shapes Consciousness*. Oxford University Press.
- Whiteley, E. (2022). Harmful Salience Perspectives. In S. Archer (Ed.), *Salience: A Philosophical Inquiry* (pp. 193–212). Routledge.
- Whiteley, E. (2023). “A Woman First and a Philosopher Second”: Relative Attentional Surplus on the Wrong Property. *Ethics*, *133*(4), 497–528. <https://doi.org/10.1086/724538>
- Workman, L., & Taylor, S. (2021). Attention and perception. In *Cognitive Psychology: The Basics* (pp. 50–83). Routledge.

Penultimate Version (please quote from forthcoming article).

Wu, W. (2011a). Attention as Selection for Action. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 97–116). Oxford University Press.

Wu, W. (2011b). Confronting Many-Many Problems: Attention and Agentive Control. *Noûs*, 45(1), 50–76. <https://doi.org/10.1111/j.1468-0068.2010.00804.x>

Wu, W. (2014). *Attention*. Routledge.

Wu, W. (2018). Action always involves attention. *Analysis*, 36, 311–395.

Wu, W. (2019). Structuring Mind: The Nature of Attention and How It Shapes Consciousness, by Sebastian Watzl. *Mind*, 128(511), 945–953.

Wu, W. (2023). *Movements of the Mind*. Oxford University Press.

Wu, W. (2024). We know what attention is! *Trends in Cognitive Sciences*, 28(4), 304–318. <https://doi.org/10.1016/j.tics.2023.11.007>

Zivony, A., & Lamy, D. (2021). What processes are disrupted during the attentional blink? An integrative review of event-related potential research. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-021-01973-2>