

DOES MOLE'S ARGUMENT THAT COGNITIVE PROCESSES FAIL TO SUFFICE FOR ATTENTION FAIL?

KRANTI SARAN
Ashoka University

Is attention a cognitive process? I reconstruct and critically assess an argument first proposed by Christopher Mole that it cannot be so. Mole's argument is influential because it creates theoretical space for a unifying analysis of attention at the subject level (though it does not entail it). Despite their differing theoretical commitments, prominent philosophers working on attention such as Wayne Wu and Philipp Koralus explicitly endorse it, while Sebastian Watzl endorses a related version. I show that Mole's argument is invalid, amend it to secure its validity, but argue that it still fails. I consider the extent to which the failure of Mole's argument spreads to the versions offered by Wu, Koralus and Watzl. Mole's argument fails because it equivocates between the set of conditions that suffice for constituting attention and the subset of those conditions that are salient, but insufficient, for constituting it. Reflection on this distinction has consequences for the individuation not just of attentional processes but all cognitive processes.

KEYWORDS: attention, unity, cognitive process, sufficiency

1. Introduction

Sebastian Watzl asks, "What, if anything, *unifies* all that is associated with attention?" (2017: 22).

On the one hand when psychologists look at the disparate mechanisms associated with attention they conclude that, "attention cannot be considered a unitary construct in humans" (Robbins 1998: 190). For example, Luck and Veraera (2002: 261) identify six effects of attention which are plausibly subserved by discrete mechanisms: attention directed to a stimuli or task boosts cognitive resources to it; stimuli that are attended to in noisy decision processes are accorded more weight; in bottleneck processes, attended stimuli have preferential

Contact: Kranti Saran <saran@ashoka.edu.in>

access through it; attention binds features whereas unattended features are free-floating; the processing gain on attended stimuli exceeds that on ignored stimuli; and finally, attended stimuli and tasks access working memory preferentially. Psychologists find no unity in the mechanisms responsible for these effects. In his historically informed and philosophically astute review of the literature on attention, Neumann concludes, “there is every reason to believe that the term ‘attention’ does not refer to a unitary entity or mechanism” (1996: 439).

On the other hand, philosophers are optimistic that attention has a unitary nature.¹ Their optimism rests in large part on an argument for the conclusion that attention is not any or all of the disparate mechanisms proffered by psychologists nor is it multiply realised by such mechanisms, because they do not suffice for it. That argument was first offered by Christopher Mole (2011a: 46; 2011b: 74). Mole’s argument is explicitly endorsed by Wayne Wu (2014: 70) and Philipp Koralus (2014: 29), and implicitly endorsed by Sebastian Watzl (2017: 27–32).

Mole’s argument appeals to philosophers because it removes powerful obstacles to a unifying analysis of attention at the subject level. The force of Mole’s argument is that it disposes of two options regarding the nature of attention: first, that attention is a single cognitive process and hence unified at the level of cognitive processing; second, that attention is multiply realised by different cognitive processes and hence disunified at the level of cognitive processing. To be sure, were Mole’s argument successful it would not by itself *entail* that attention is only unified at the subject level, but it would remove two powerful obstacles in the way of that conclusion.

Even if Mole’s argument were successful, it would not dispose of all the obstacles to a unifying analysis of attention at the subject level because it does not dispose of the option that attention is disunified at the subject level. For example, focusing, absorption, vigilance, or the intellectual and sensorial forms of attention that Williams James distinguished, might lack a common nature.² Philosophers seeking to unify attention at the subject level typically have other arguments either to eliminate, as instances of attention, subject level phenomena that are thought to be attentional phenomena (Mole 2011a: 57) or to reduce putatively distinct forms of attention to one (Wu 2014: 94–95). Thus an accurate statement of the dialectic is that were Mole’s argument to succeed we could conclude that were attention unified, it would be unified at the subject level.

As to the substance of the unitary nature of attention at the subject level, there are several proposals on the table: Mole (2011a) argues that it is cognitive unison, Wu (2014) that it is selection for action, Watzl (2017) that it is the activity of regulating priority structures, and Koralus (2014) that its functional role is akin

1. William James is a prominent exception; he recognised the varieties of attention in his chapter on attention in a section titled “The Varieties of Attention” (1890: 416–425).

2. See James (1890: 416–425).

to the relationship between questions and answers.³ I do not aim at adjudicating between these disparate proposals here. I aim, instead, at critically assessing the common argument these philosophers rely on despite their differing theoretical commitments.

I argue that Mole's argument fails. The failure of Mole's argument resurrects the obstacles to a unifying analyses of attention at the subject level—at least until another route is found to the conclusion that attention is not a cognitive process. I build up to Mole's argument in Section 2, reconstruct and amend it in Section 3, lay out the key distinction for analysing Mole's argument in Section 4, criticise Mole's argument in Section 5, and critically assess versions of Mole's arguments offered by other philosophers in Section 6.

2. Mole's Argument

What kind of thing is attention, according to Mole? Mole is an adverbialist about attention: its correct metaphysical category is not a *type* of process but rather the *manner* in which some process takes place. To bring Mole's picture into view, ask yourself, to which metaphysical category do paperweights belong? Mole argues that what makes a thing a paperweight is not what it is made of—its substance—but rather the manner of its use. The metaphysics of paperweights matters, Mole thinks, because correct taxonomisation requires correct metaphysical categorisation. Metaphysical miscategorisation begets an untenable taxonomy, he argues: "What stops us from being able to delineate the paperweights using the distinctions provided by a substance taxonomy is that being made of the same substance as a paperweight does not suffice for being a paperweight. That makes it impossible to rule in all of the paperweights without also ruling in some of the nonpaperweights" (2011a: 27). Mole's point is that if a substance-categorisation drives taxonomisation, it will give the wrong results: on the supposition that paperweights are made of substance X, the resulting taxonomy will include instances of X that are not used as paperweights; this point also cuts against any attempt to argue that paperweights are multiply realisable by a disjunction of types of substances, since there are instances of each type of substance which are not used as paperweights.⁴ According to Mole, what goes for paperweights goes

3. Strictly speaking Koralus (2014: 36) is explicit that his account of the nature of attention is neutral on whether it applies at the subpersonal level. However, all the other philosophers mentioned take Mole's argument as one reason to give a unifying analysis of attention at the *subject* level. Mole and Wu would likely agree with Watzl that, "Attention ... has unity only on the subject level" (Watzl 2017: 37).

4. See Mole (2011a: Section 2.4) on this point. In the next chapter, Mole (2011a: 40) is careful to remark—here I translate talk of attention into paperweights—that one *could* individuate paperweights so finely such that a gerrymandered disjunction gives the correct taxonomy, but at the price that it would no longer be the *substance* of paperweights that drove their taxonomy

for attention: “The claim that attention does not have the metaphysics of a process. . . rules out the possibility that attention can be identified with a disjunction or family of several processes, any one of which is sufficient for the instantiation of attention” (2011a: 31).

It is not the purpose of this paper to assess Mole’s adverbial metaphysics. Rather, the focus of this paper is a particular argument for the conclusion that attention is not a cognitive process that he gives in the context of his larger argument for adverbialism. (I showed in the previous paragraph that Mole argues that adverbialism about attention rules out that attention is a cognitive process. But Mole does not think you need to endorse his adverbialism to think that attention is not a cognitive process. Rather, the structure of Mole’s argument for adverbialism relies, in part, on the claim that attention is not a cognitive process; that claim also happens to be entailed by his view.) To repeat, you do not have to share Mole’s adverbial metaphysics to feel the force of the argument that is the focus of this paper. In fact, his argument is part of what is supposed to motivate the reader to give serious consideration to Mole’s adverbialist alternative.⁵

What does Mole mean by a “cognitive process”? The kinds of cognitive processes mentioned in Mole’s argument all seem sub-personal, for example, Triesman’s feature-binding, and that is how his argument has been read in the literature. But a careful reading of Mole’s official definition of a cognitive process shows that it is not restricted to sub-personal processes:

A cognitive process . . . is a process that operates on representations that encode their contents *for the agent of the task*: In order for a process to count as cognitive, there must be an agent-level contentful state whose content is directly determined, at least in part, by the content of the representations on which that process operates. More precisely, a process is a cognitive process if and only if it is an immediate consequence of the fact that this process is operating on a representation with the content . . . ζ . . . that the *agent* is in a position to believe that . . . ζ . . . , or . . . or to have some other attitude toward the content of the representation in question. (2011a: 57–58)

Notice that the subject-level process of adding 5 and 7 to give 12 operates on representations with the content 5 and 7 such that the agent is in a position to believe that they add up to 12. Given his capacious conception of cognitive processes, it is not surprising to find Mole claiming that “When we deny that attention has the metaphysics of a process we are thereby committed to the claim that the question of whether or not a subject is paying attention is not settled by

but rather the manner of their use, hence conceding the case to adverbialism.

5. For a statement of the dialectical role of the argument that attention is not a cognitive process in Mole’s larger argument for his adverbial account of attention, see (2011a: 35–36).

the facts about which *processes* are going on at *any* level of description" (2011a: 32). Though Mole does not say so explicitly it should be clear that cognitive processes include processes that are defined *functionally*. The key point according to Mole is that the instantiation of any brain process or functionally characterised process does not suffice for attention, because what matters is the *manner* of their occurrence. (I will table more discussion of Mole on cognitive processes for Section 5, when I look at the reception of Mole's argument.)

Mole's strategy to establish that attention is not a cognitive process starts with attributing a supervenience claim to his opponent, one that Mole thinks is implausible:

The process-first view is committed to the claim that the instantiation of the same processes as are instantiated in a case of attention is sufficient for being an instance of attention, and so it is committed to claiming that no change from attention to inattention is possible without a change in the underlying cognitive processes. . . . The process-first view is committed to saying that nothing at all, natural or otherwise, could make a difference to whether or not a set of processes constitute attention, other than by changing the facts about which processes those are. That claim is so strong as to be implausible. (2011a: 36)

Mole thinks that attributing this supervenience claim to this opponent is plausible because his opponent thinks that a cognitive process *suffices* for attention, so surely the very same cognitive process must constitute attention in every case or it would not suffice for attention. Why does Mole think his opponent's supervenience thesis is implausible?

The supervenience thesis rules out cognitive processes constituting attention in one context but not in another. Mole thinks that what supervenience rules out, reality rules in. Mole (2011a: 38) argues that Triesman's view that feature-binding constitutes attention is just such a process. In the normal case it constitutes attention. However, there is evidence from unilateral neglect patients that they bind features without those processes constituting attention. Since they neglect items on the left side of space, they are not attending to them. But the evidence suggests that feature-binding occurs in the neglected area because such subjects show priming effects that depend on the semantic content of figures that must be bound together to have that content. Another source of evidence is these patients suffer from visual illusions when illusory figures are presented on their neglected side, but producing illusory effects requires binding the figures.⁶

The contrast between the normal case of attention and unilateral neglect suggests to Mole a 'recipe' for discovering pairs of cases in which a cognitive process constitutes attention in one case but not another:

6. Mole (2011a: 38) provides references for these claims.

Once one of these simple contexts has provided us with an example of a simple attention-constituting process we can then construct a new example where we keep that simple piece of processing fixed but situate it in a more complex context, where the subject is engaged in something more sustained and cognitively elaborate. Where there is sufficient other processing to swamp the process that we started with, or where, as in the unilateral neglect case, there is some deficit that prevents that process from being properly integrated with the rest of the subject's cognition, we have an example where it is no longer feasible that that process constitutes the subject's attention. This is a recipe for finding a pair of cases in which the same piece of processing is in one case attention constituting and in the other case not. (2011a: 41)

By finding pairs of cases instantiating the same process only one of which constitutes attention, Mole undermines what he takes to be the implausibly strong supervenience claim that attention supervenes on the processes that instantiate it:

In each case, it is the nonsufficiency of instantiating the same process as an instance of attention that enables us to infer that the matter of whether or not some event instantiates attention is not settled by the facts about the cognitive processes that the event instantiates, and it is on this basis that we conclude that attention is not a cognitive process. (2011a: 46)

Thus at the heart of the pattern of cases that Mole considers is *sufficiency to constitute attention*: despite instantiating the same process, one of the cases does not instantiate attention, hence attention is not a cognitive process.

But what exactly is Mole's argument? Here is his argument reproduced verbatim from (2011a: 46):

- (1) If attention is a cognitive process, then, for all events x and y , if x and y instantiate the same cognitive processes, then if either one of them is an instance of attention, the other is too.
- (2) There are some events that are instantiations of the feature-binding process and that are instances of attention.
- (3) There are some events that are instantiations of those same feature-binding process and that are *not* instances of attention.

Therefore:

- (4) Attention is not a cognitive process.

Mole takes Treisman's work on feature-binding to support (2). (3) rests on the evidence that some unilateral neglect patients must be binding unattended features.⁷

Mole (2011a: 40) is careful to say that his argument is not a "conclusive refutation" of the view that attention is a cognitive process, for two reasons: first, because a sufficiently fine-grained individuation of cognitive processes could give the correct taxonomy for attention via a gerrymandered disjunction (though likely at the cost of having the processes drive the taxonomy—see Footnote 4); second, the claim that feature-binding processes constitute attention rests on an inference to the best explanation, which is hostage to empirical fortune—a better explanation may yet emerge. My own objections to Mole's argument rest on neither of these grounds and I do not contest Mole's empirical premises in either of these two ways. Rather, my focus will be on (1).

What exactly does (1) mean? In the next section, I sharpen Mole's argument, show that it is invalid, and suggest a fix. This sets the stage for my core philosophical objection to Mole's argument: even the patched-up version of Mole's argument suffers from an equivocation between sufficient conditions for attention and salient but insufficient conditions for attention.

3. Reconstructing Mole's Argument

Recall that Mole's first premise reads as follows:

- (1) If attention is a cognitive process, then, for all events x and y , if x and y instantiate the same cognitive processes, then if either one of them is an instance of attention, the other is too.

Let's start with the antecedent clause of the outermost conditional. The antecedent of (1) says something weaker than the property of *being an instance of attention* is identical to that of *being a cognitive process*, because not all cognitive processes are cognitive processes of attention. It says something stronger than anything that is an instance of attention is an instance of a cognitive process. A charitable reading of the antecedent is that attention is a *determinant* or *type* of cognitive process: it says that being an instance of attention is a type of being a cognitive process.

There are two further details of (1) to be ironed out. First, what does Mole mean by the "same cognitive processes" instantiated in two events? Mole might

7. See Mole (2011a: 46). Though cast in terms of Treisman's Feature Integration Theory, Mole is clear that the form of the argument lends itself to any mechanism; for example, he considers the analogous argument in which 'feature binding' is substituted by 'task switching' with suitable changes to the empirical premises (2) and (3).

mean that every cognitive process whatsoever instantiated in one event is duplicated in the other. However this cannot be Mole's intended meaning. Recall Mole's recipe, quoted earlier, for generating pairs of cases instantiating the same attentional processes only one of which constitutes attention.⁸ Mole's recipe turns on taking the same cognitive processes *of attention* and embedding them in a different 'surround' so that they no longer constitute attention. If Mole's recipe involved cases that duplicated every cognitive process whatsoever then every instance of attention in one would be an instance of attention in the other. Hence the more charitable interpretation of Mole is that by the "same cognitive processes" in the two cases, he does not mean that every cognitive process whatsoever instantiated in one event is duplicated in the other, but rather only the same cognitive processes *of attention*. Second, given that by "same cognitive processes" Mole means the same cognitive processes of attention, we can sharpen (1) further by specifying that Mole intends that the two cases instantiate *all* the same cognitive processes of attention. (The alternative reading, that the two events only share *some* of the same cognitive processes of attention, is too weak to secure the consequent of (1).)

If my interpretation of (1) is correct, we can rewrite it as:

- (1') If being an instance of attention is a type of being a cognitive process of attention, then, for all events x and y , if x and y instantiate all the same cognitive processes of attention, then if either one of them is an instance of attention, the other is too.

and rewrite the conclusion, which is the negation of the antecedent, as:

- (4') It is not the case that being an instance of attention is a type of being a cognitive process of attention.

An informal counter-example shows that Mole's argument is invalid. Suppose that the domain of events consists of E_1 , E_2 , and E_3 , all of which involve instantiations of cognitive processing. Suppose there are not one but two distinct attentional processes P_1 and P_2 . In E_1 , P_1 instantiates attention, but in E_2 , P_1 does not. In E_3 , P_2 instantiates attention. Premise (2) is true because in E_1 , P_1 instantiates attention. Premise (3) is true because in E_2 , P_1 does not instantiate attention. Now consider the antecedent of (1') 'being an instance of attention is a type of being a cognitive process'; it is true because both E_1 and E_3 , the events instantiating attention, are cognitive processes. Since (4') is just the negation of the true antecedent of (1'), it is false. Now turn to the consequent of (1'), which is itself a conditional. It says that for any two events such that they instantiate the same cognitive processes of attention, one is an instance of attention if and

8. See Mole (2011a: 41).

only if the other event is an instance of attention. The antecedent of this latter conditional is false when we consider E_1 and E_3 which instantiate two *distinct* processes of attention; hence this latter conditional is vacuously true. Hence both the antecedent and consequent of (1') are true. Hence we have an interpretation on which all the premises are true but the conclusion is false.

While Mole's argument is invalid there is a valid argument in the vicinity that accords with the spirit of Mole's view and preserves the first premise and conclusion:

- (1') If being an instance of attention is a type of being a cognitive process of attention, then, for all events x and y , if x and y instantiate all the same cognitive processes of attention, then if either one of them is an instance of attention, the other is too.
- (2') If any two events x and y instantiating all the same cognitive processes of attention and such that one instantiates attention if and only if the other does, then it cannot be case that there exist two events instantiating the same cognitive processes of attention (i.e., feature-binding processes) only one of which is an instance of attention.
- (3') There exist two events instantiating the same cognitive processes of attention (i.e., feature-binding processes) only one of which is an instance of attention.
- (4') It is not the case that being an instance of attention is a type of being a cognitive process of attention.

I think both (2') and (3') accord with the substance and spirit of Mole's original argument: (2') captures the thought that the supervenience relation (that Mole rejects) entails that there cannot be a certain pattern of cases, while (3') captures the thought that precisely such a pattern of cases exists. This version of the argument is valid, but it still fails.

4. The Key Distinction: Sufficient versus Salient but Insufficient Conditions

Why does Mole's argument—even the patched up version—fail? To set the stage for understanding why Mole's argument fails, we need the distinction between the set of conditions that suffice for constituting attention and the subset of those conditions that are salient but insufficient for constituting attention.

How do the sufficient conditions for constituting attention differ from the salient but insufficient conditions? Before fleshing out that distinction, I distinguish it from another more specific and influential distinction with which it is

apt to be conflated. I argue that because that more specific and influential distinction is entangled with disputes tangential to my purpose, and because Mole does not appeal to it, clarity and charity demand the more general distinction for my analysis of Mole's argument.

The more specific and influential way of distinguishing sufficient conditions from salient but insufficient conditions is in terms of the *core* realiser and the *total* realiser for the instantiation of a property introduced by Shoemaker (1981). The core realiser makes a salient but insufficient contribution to realising a property. The total realiser consists of the core realiser plus supplemental conditions and suffices for realising the property. Shoemaker illustrates the distinction as follows:

C-fiber stimulation in a Petri dish will not realize pain, or any other mental state. At best C-fiber stimulation will be what I call a *core realizer* of pain, rather than a *total realizer* of it. A total realizer of a property will be a property whose instantiation is sufficient for the instantiation of that property. A core realizer will be a property whose instantiation is a salient part of a total instantiation of it. In the case where having C-fiber firing occurring in one is of the core-realizers of being in pain, the total realizer will be something we might call "C-fiber-stimulation-plus"—having C-fiber firing occurring in one and having a brain that is so wired that C-fiber stimulation in it has the standard causes and effects of pain. (2007: 21–22)

What principle determines what counts as a core realiser? There is no single principle in the sense that what is salient shifts depending on the context of inquiry. For example, the core realiser may differ depending on the level of inquiry (e.g., neurons, functional profiles, etc.) and whether our interest is in prediction, control, intervention, or some other aim.

A consequence of the context-sensitivity of the core realiser is that we can apply it not just to pain characterised as a neural event of C-fibres firing, but also to pain characterised *functionally* as having the standard causes and effects of pain. Just as the neural core realiser (C-fibres firing) is distinct from the total realiser of pain, the core *functional* profile of pain is distinct from the 'surround' without which it would not realise pain. The claim that the functional core realiser of pain differs from the 'surround' that is part of its total realiser rests on the plausible assumption that the former is characterised to exclude other neuro-functional states distinct from pain but necessary for its occurrence, such as the states responsible for having any consciousness at all. Thus, whether the feature integration process is characterised at the neural or functional level, Shoemaker's distinction still gets a grip.

What principle determines what counts as a total realiser? What counts as

a total realiser is *everything* that needs to be instantiated to *suffice* for the mental state. How to characterise everything that suffices for a mental state—its total realiser—engenders at least two disputes the opening moves of which are sketched below.

The first dispute centres on 'everything' or the actual elements that go into sufficing for a mental state. For example, take the debate about the minimal metaphysically sufficient conditions for perceptual experience. Block (2005) holds the 'orthodox' view that nothing outside the brain is part of it, while O'Regan and Noë (2001) and Noë (2004) argue that it includes both the brain and bodily activity. Block and his opponents disagree about what elements go into sufficing for a mental state, but they share a common understanding of what a 'sufficient condition' amounts to.

The second dispute centres on 'sufficiency'. Is it context-sensitive like the core realiser? For the affirmative, consider the work of Wilson (2001; 2004) who relativises the concepts of core and total realisers to the state of a subject, thereby excluding from the total realiser what he calls "background conditions" in the environment like oxygen without which the subject could not exist; thus his total realisers are "metaphysically context-sensitive" and he concedes that they are not by themselves sufficient for the instantiation of a mental state (2001: 9).⁹

In contrast, those who deny that the total realiser is context-sensitive can motivate their view by an analogy with Lewis on causation. Lewis (1973: 558–559) is not interested in the single cause we select as "the" cause of an event because we think it noteworthy or under our control, he rejects such "principles of invidious discrimination" that filter all the causes down to one. Lewis writes,

We sometimes single out one among all the causes of some event and call it "the" cause, as if there were no others. Or we single out a few as the "causes," calling the rest mere "causal factors" or "causal conditions." Or we speak of the "decisive" or "real" or "principal" cause. We may select the abnormal or extraordinary causes, or those under human control, those we deem good or bad, or just those we want to talk about. I have nothing to say about these principles of invidious discrimination. I am concerned with the prior question of what it is to be one of the causes (unselectively speaking). (1973: 558–559)

Just as Lewis rejects invidious discrimination about causes because it reflects our selective interests rather than the unselective metaphysical facts, the opponent of context-sensitivity rejects invidious discrimination about metaphysically

9. Though not cast in terms of core and total realisation, Burge (2010: 57–59) also brackets 'background conditions' on the grounds that in the context of constitutive questions, we are only interested in the subset of sufficient conditions that *explain* the nature of the state. Thus were Burge to speak of the total realisers of mental states he would consider them *epistemically* context-sensitive.

sufficient conditions because it reflects our selective interests rather than the unselective metaphysical facts.

While settling these disputes is beyond the scope of this paper, two crucial points stand out: first, everybody agrees that the core realiser does not suffice for what it realises; second, nobody disputes that a line can be drawn between the core and total realiser, even though how and where to draw the outer limit of the total realiser is deeply controversial.

Given the distinction between core and total realisers, why avoid using it in favour of the more general distinction between sufficient conditions and salient but insufficient conditions? After all, the latter distinction inherits analogous controversies about what suffices for a mental state and whether sufficiency is context sensitive, so there is no particular gain to switching on that score. Nevertheless, dropping Shoemaker's nomenclature extracts the core of the distinction and puts daylight between it and irrelevant controversies specific to Shoemaker's views. It also allows me to couch my objection in terms that Mole himself uses.

Shoemaker's distinction between core and total realisers is cast in the terms of his theory of property realisation: *properties realise other properties*. Both components of his view are controversial. Whereas most philosophers think that causation is a contingent relation, Shoemaker has long held that it is not: properties have their causal profiles essentially.¹⁰ According to Shoemaker (2010: 127) properties are individuated by their forward-looking causal profiles.¹¹ On his 'subset' model of realisation, property *A* realises property *B* when *B*'s forward-looking causal features are a subset of *A*'s forward-looking causal features. Philosophers have worried that his account does not solve the causal exclusion problem (Audi 2012; Pineda & Vicente 2017), that it results in the overdetermination of the mental by the physical (Morris 2011), and that it does not adequately explain why physical events entail mental events (McLaughlin 2007; Tiehen 2014).

Rather than reinterpreting the core and total realiser out of the context of Shoemaker's controversial account of properties and realisation, it is simpler to break from Shoemaker and use the more general distinction that fits better with Mole's own language.

Sufficiency to constitute attention is at the heart of Mole's argument. The distinction between sufficient conditions and salient but insufficient conditions helps us to keep in mind that a single isolated process, whether characterised neurally or functionally, does not suffice for a mental state. As Shoemaker's

10. See Shoemaker (1980). Shoemaker also thinks that if properties have their causal profiles essentially, then, "causal laws are metaphysically necessary... and... causal necessity is metaphysical necessity" (1998: 66).

11. Shoemaker (2010) recants his earlier claim in (2007: 12) that the backwards-looking causal profile of a property partly individuates it as he now thinks that if two properties share all their forward-looking causal features then they must also share all their backward-looking causal features. Thus he reverts to his view in (2001: 77–78).

example of C-fibre firing illustrates, it is a myth to think that some small subset of brain processing *suffices* for pain. I argued that it is equally a myth to think that the functional profile of pain suffices for pain. What goes for C-fibre firing and pain goes for feature-binding and attention. If sufficiency is what we are after, we have to expand the scope of our concern from a single isolated mechanism to a mechanism in the entire context of brain activity. A mechanism seen in the entire context of brain activity is a mechanism whose efficacy is susceptible to changes in what Shoemaker would call the 'surround'. Let me return to the details of Mole's patched-up argument.

5. A Critical Assessment of Mole's Patched-Up Argument

Does 'all the same cognitive processes of attention' in (1') refer to the sufficient conditions or just the salient but insufficient conditions for constituting attention? If it refers to the latter then (1') is vacuously true, because by definition the salient but insufficient conditions will not suffice to constitute attention—it is only the most salient part of what suffices for attention. Thus the intended reading of (1') must refer to the sufficient conditions for constituting attention.

Does the 'cognitive process of attention' in the consequent of (2') refer to the sufficient conditions or just the salient but insufficient conditions for constituting attention? I think Mole is forced to adopt the latter reading, especially when the consequent of (2') is read in the context of (3')—it must be the same particular and restricted attentional process at issue in both premises. Mole surely does not mean all the processes of any type, for then both cases would be whole-brain cognitive process duplicates, which is clearly not the scenario Mole envisages. On the other hand, the more restricted 'attentional' processes are at best the salient but insufficient conditions for constituting attention, because to think otherwise would be to fall into the trap of thinking that feature-binding processes can constitute attention without the contribution of the 'surround'. So these premises must mean by 'feature-binding processes', processes that are salient but, as I have argued, insufficient to constitute attention.

If my analysis of the first two premises is correct, Mole's argument equivocates between the sufficient conditions for constituting attention in (1')—and the antecedent of (2')—and the salient but insufficient conditions for constituting attention in the consequent of (2'). If we reason about (3') as we did about the consequent of (2'), *mutatis mutandis*, then we recapitulate an equivocation between the salient but insufficient conditions at issue in (3') and the sufficient conditions at issue in (1').

Finally, the sense of 'cognitive process' in the conclusion (4') cannot be the sufficient condition sense in (1'), because if it were, then the process *would* suffice for attention. Hence it must be the salient but insufficient condition for consti-

tuting attention. Again, we have an equivocation, this time between (1') and (4').

Mole's ground clearing argument fails and its failure teaches us to keep a sharp eye on how individuating a mental state by the conditions that are salient but insufficient for constituting it differs from individuating it by the conditions that suffice to constitute it.

6. Mole's Argument in Wu, Koralus, and Watzl

Mole's argument for the conclusion that attention is not a kind of cognitive process is significant because versions of it are recapitulated by other philosophers working on attention. In his final section, I discuss the extent to which my critique of Mole's argument infects the versions of it offered elsewhere in the literature.

6.1. Wu

Wu reproduces Mole's argument verbatim, except that he drops "cognitive" from Mole's "cognitive process". Wu adds, "One can replace 'process' with 'neural mechanism' and in place of feature-binding substitute any appropriate mechanism" (2014: 70). He goes on to argue on the same page, using Mole's argument, that divisive normalisation does not suffice for attention.

It should be clear that as presented in Wu, the argument is invalid. Further, even if we were to render it valid in the style adopted for Mole's original argument, it would still suffer from equivocating between sufficient conditions for attention and the salient but insufficient conditions for attention, except this time the process in question would not be feature-binding but divisive normalisation.

The failure of the argument has repercussions for the motivation behind Wu's project. Wu (2014: 71) argues that genuine progress in studying attention requires unifying principles. One candidate unifying principle is that attention is a type of neural process. But Wu thinks the argument undercuts such a unifying principle. It is the failure of the process view that motivates, at least in part, Wu's identification of the unifying principle as *selection for task*. To recap, in Wu's hands, the conclusion that attention is not divisive normalisation motivates him to look for the additional ingredient X which when added to divisive normalisation suffices for attention; Wu thinks that X is *erving selection for task* (2014: 71). But Wu has not established that attention is not divisive normalisation, so the dialectical motivation for the search for X is shaky.

6.2. Koralus

Koralus (2014: 29-30) endorses the *task-switching argument* found in Mole (2011a:

39), which is just a particular instance of the general pattern of cases in which an acknowledged case of attention instantiates a process while in the second case the same process is instantiated in what is putatively not a case of attention, hence the process cannot suffice for attention:

suppose you observe someone who first has an active psychological process A corresponding to one strategy of generating numbers and then switches to a psychological process B corresponding to another strategy of generating numbers. If that person is following the task of generating random numbers, then the observed switch could well be indicative of the fact that she is following her task attentively... By contrast, if that person is following the task of calculating a certain numerical series, then the switch between strategies could well be indicative of the fact that her attention is slipping. Mole argues that it is very plausible that we have pairs of tasks of this sort, where a switch from a psychological process A to a process B corresponds to a lack of attention *if we are pursuing one task*, but corresponds to full attention, if we are pursuing another task. If we have such pairs of tasks, then attention is not a matter of what type of psychological processes are active. (Koralus 2014: 29)

Though the task-switching argument introduces the notion of task performance, at heart it turns on the argument we analysed and rejected. For all the reasons advanced against Mole's argument, the task-switching argument should also be rejected.¹²

As with Wu, the failure of the argument has repercussions for the motivation behind Koralus's project. Koralus banks on Mole having established that attention is not a type of cognitive process, "that attention is not a matter of whether a particular kind of cognitive process is active, but of how the activity of psychological processes relates to a task" (2014: 30). Koralus's project is then to give a novel account of what it is for psychological processes to relate to task in such a manner that they constitute performing a task attentively. But since Mole has not established that attention is not a cognitive process, Koralus is banking on what he ought not to bank on and the dialectical motivation to locate attention in the manner of the relation between psychological processes and task is shaky.

12. There are independent reasons that make trouble for the task-switching argument: it seems to implausibly collapse the distinction between attending to a task but botching it and performing a task inattentively. For example, suppose the task is carrying a full glass of water from the kitchen to the dining table without spilling it: spilling it while moving carefully seems to differ from spilling it while moving carelessly. The task is unsuccessfully performed in both cases, but only in the latter case is it clear that the task was performed inattentively. For all Mole has said about the task-switching case, both involve attention but the person just botches the task in one case.

6.3. Watzl

In contrast to Wu and Koralus, Mole's argument is less central to the presentation of the positive view argued for in Watzl (2017). This is not surprising. Whereas Wu and Koralus's projects seem consistent with Mole's adverbialism¹³, Watzl's project is *inconsistent* with Mole's adverbialism. Watzl's view is that attention is a certain high-level process—the process of regulating priority structures, but recall Mole's remark quoted earlier, "When we deny that attention has the metaphysics of a process we are thereby committed to the claim that the question of whether or not a subject is paying attention is not settled by the facts about which *processes* are going on at *any* level of description" (2011a: 32).

Suppose we restrict, as Mole does not, "cognitive processes" to a class of sub-personal processes. Thus restricted, Watzl endorses the conclusion of Mole's argument. Watzl writes, "the anti-reductionist position agrees with the adverbial view that attention is not identical to a type of neuronal or computational process or mechanism (or set of such processes or mechanisms)" (2017: 33), which entails that attention is not a type of sub-personal cognitive process. But what about Mole's argument? How does it show up in Watzl?

Pinning down the exact form of Watzl's overall argument for Mole's conclusion is a demanding hermeneutical task best left for another occasion; I will only focus on the parts of Watzl's argument salient to Mole's argument. On the one hand Watzl offers what he calls, "a strong inductive, and empirically based, argument against identifying reductionism about attention" (2017: 27) while on the other hand he casts his argument as a dilemma for any attempt to argue that attention is a kind of cognitive process. We find an echo of Mole's argument in both of Watzl's argumentative strategies. Consider the first horn of Watzl's dilemma:

The first horn of the dilemma: suppose one attempts to identify attention with a fairly specific neuronal or computational mechanism. In this case we are faced with two problems. *First, we are confronted with a wealth of findings that suggest that many neuronal or computational mechanisms that are tightly associated with attention in some contexts, in other contexts operate without attention [my italics].* Given current evidence it is highly unlikely that any one specific mechanism is operative in all scenarios in which attention is employed. *Second, current evidence suggests that it is highly unlikely that any specific mechanism can explain most of the central fea-*

13. Wu writes, "I propose that cognitive unison [Mole's adverbial view] conjoined with selection for action provides a more complete account of attention" (2014: 77) while Koralus accepts adverbialism but differs from Mole in giving, "a novel account of the relation that has to hold between a task and various psychological processes in order for an agent to count as performing a task attentively" (2014: 30).

tures of attention. Rather, different mechanisms seem to explain different features of attention. (2017: 27)

A version of Mole's argument crops up again in Watzl's inductive argument. After reviewing a number of candidate processes for attention and illustrating pairs of cases instantiating the process where one of the pair lacks attention, Watzl argues,

Wherever we look we seem to find the same assessment. The attempt to identify attention with a specific set of neuronal or computational mechanisms fails, because these mechanisms often operate also in the absence of attention, and even when they are implicated in attention the mechanisms tend to explain only a subset of attentional effects. (2017: 31)

These passages from Watzl suggest that his view is susceptible to the objection against Mole's argument: it is a mistake to think of the pair of cases as involving a particular process sufficing for attention in one case but not in another. A particular process, certainly a process individuated as finely as these authors consider, *never* suffices for attention. Instead, these processes are at best salient but insufficient conditions for attention. Therefore, it is illegitimate to infer that the process does not constitute attention *because* in some cases it occurs without attention; rather, any such process fails to suffice for attention because sufficing for attention requires the contribution of the 'surround'. What can be concluded from such pairs of cases is rather weak: that process P in the context of the 'surround' S suffices for effect X whereas the same process in the context of a different 'surround' S' does not.

To what extent do these objections impact Watzl's project? A definitive judgement must wait on a careful and detailed reconstruction of Watzl's arguments. But we can make a provisional observation. Like Mole, Wu, and Koralus, Watzl takes the establishment of the conclusion that attention is not a cognitive process as an essential ground-clearing move to providing his alternative, a unifying analysis of attention at the subject level. To the extent then that Watzl's argument for that conclusion falters in its reliance on a version of Mole's argument, the dialectical impetus for Watzl's own view weakens.

7. Conclusion

I have argued that Mole's argument for the conclusion that attention is not a cognitive process fails. It fails by equivocating between the sufficient conditions for attention and the salient but insufficient conditions for attention. Its failure is not merely a matter of local interest to Mole aficionados in two ways. First, my analysis shows that we must pay careful attention to how individuating a mental

state by the conditions that are salient but insufficient for constituting it differs from individuating it by the conditions that suffice to constitute it. Second, the failure of Mole's argument ripples outwards, weakening the dialectical impetus for the positive projects of both Wu and Koralus. It also seems to weaken, though perhaps to a lesser extent, the impetus for Watzl's view. With the failure of Mole's argument, proponents of the view that attention has a unitary nature at the subject level need a new argument to show that attention is neither a single cognitive process nor multiply realised by different cognitive processes.

Acknowledgements

Thanks to Maya Saran and Scott Dixon for helping me think through the logic of Mole's argument. Thanks too to the anonymous referees whose comments helped me to greatly improve the paper. Finally, I am grateful to Ashoka University for granting me research leave and to Harvard University for providing me access to scholarly resources. Without their support this paper could not have been started, let alone completed.

References

- Audi, Paul (2012). Properties, Powers, and the Subset Account of Realization. *Philosophy and Phenomenological Research*, 84(3), 654–674. doi:10.1111/j.1933-1592.2010.00476.x
- Block, Ned (2005). Review of *Action in Perception* by Alva Noë. *The Journal of Philosophy*, 102(5), 259–272.
- Burge, Tyler (2010). *Origins of Objectivity*. Clarendon Press.
- James, William (1890). *The Principles of Psychology* (Vol. I). Henry Holt & Co.
- Koralus, Philipp (2014). The Erotetic Theory of Attention: Questions, Focus and Distraction. *Mind & Language*, 29(1), 26–50. doi:10.1111/mila.12040
- Lewis, David (1973). Causation. *The Journal of Philosophy*, 70(17), 556–567.
- Luck, Steven J. and Shaun P. Vecera (2002). Attention. In Hal Pashler and Steven Yantis (Eds.), *Steven's Handbook of Experimental Psychology* (235–286). John Wiley & Sons.
- McLaughlin, Brian P. (2007). Mental Causation and Shoemaker-Realization. *Erkenntnis*, 67(2), 149–172. doi:10.1007/s10670-007-9069-7
- Mole, Christopher (2011a). *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford University Press.
- Mole, Christopher (2011b). The Metaphysics of Attention. In Christopher Mole, Declan Smithies, and Wayne Wu (Eds.), *Attention* (60–77). Oxford University Press.

- Morris, Kevin (2011). Subset Realization, Parthood, and Causal Overdetermination. *Pacific Philosophical Quarterly*, 92(3), 363–379. doi:10.1111/j.1468-0114.2011.01401.x
- Neumann, Odmir (1996). Chapter 10: Theories of Attention. *Handbook of Perception and Action*, 3, 389–446.
- Noë, Alva (2004). *Action in Perception*. The MIT Press.
- O'Regan, J. Kevin and Alva Noë (2001). A Sensorimotor Account of Vision and Visual Consciousness. *Behavioral and Brain Sciences*, 24(05), 939–973.
- Pineda, David and Agustín Vicente (2017). Shoemaker's Analysis of Realization: A Review. *Philosophy and Phenomenological Research*, 94(1), 97–120. doi:10.1111/phpr.12251
- Robbins, Trevor W. (1998). Arousal and Attention: Psychopharmacological and Neuropsychological Studies in Experimental Animals. In Raja Parasuraman (Ed.), *The Attentive Brain* (189–220). The MIT Press.
- Shoemaker, Sydney (1980). Causality and Properties. In Peter van Inwagen (Ed.), *Time and Cause* (109–135). D. Reidel.
- Shoemaker, Sydney (1981). Some Varieties of Functionalism. *Philosophical Topics*, 12(1), 93–119.
- Shoemaker, Sydney (1998). Causal and Metaphysical Necessity. *Pacific Philosophical Quarterly*, 79(1), 59–77. doi:10.1111/1468-0114.00050
- Shoemaker, Sydney (2001). Realization and Mental Causation. In Carl Gillett and Barry Loewer (Eds.), *Physicalism and Its Discontents* (74–98). Cambridge University Press.
- Shoemaker, Sydney (2007). *Physical Realization*. Oxford University Press.
- Shoemaker, Sydney (2010). Reply to My Critics. *Philosophical Studies*, 148(1), 125–132. doi:10.1007/s11098-010-9499-y
- Tiehen, Justin (2014). Subset Realization and the Problem of Property Entailment. *Erkenntnis*, 79(2), 471–480. doi:10.1007/s10670-013-9493-9
- Watzl, Sebastian (2017). *Structuring Mind: The Nature of Attention and How it Shapes Consciousness*. Oxford University Press.
- Wilson, Robert A. (2001). Two Views of Realization. *Philosophical Studies*, 104(1), 1–31. doi:10.1023/A:1010358523824
- Wilson, Robert A. (2004). *Boundaries of the Mind. The Individual in the Fragile Sciences: Cognition*. Cambridge University Press.
- Wu, Wayne (2014). *Attention*. New Problems of Philosophy. Routledge.