Review of Attention is Cognitive Unison: An Essay in Philosophical Psychology, by Christopher Mole

Aaron Henry & Tim Bayne

of intelligibility? Perhaps a more charitable reading of Gutting's comments about obscurity is that they reflect a desire for more informed intercultural commentary of the kind that can enlarge our philosophical horizons. *Thinking the Impossible* is a very worthwhile contribution to this project even if it shows that, in the case of French philosophy since 1960, there is still a way to go.

Paul Patton
University of New South Wales
© 2012 Paul Patton

Mole, Christopher, *Attention is Cognitive Unison: An Essay in Philosophical Psychology*, New York: Oxford University Press, 2011, pp. 224, US $49.95 (hardback).

As Mole's book reminds us, attention was a central topic of interest for both philosophers and psychologists at a time when the boundaries between the two disciplines were not as clearly demarcated as they now are. Alexander Bain, William James, F. H. Bradley and G. F. Stout, amongst others, all had things to say about the nature of attention. But philosophy and psychology have had very different attitudes to attention since they parted ways: while psychologists have made attention one of their central targets, philosophers of mind have typically neglected the topic.

   This period of neglect has now come to an end, and attention is now well and truly back on the philosophical agenda. One of the most elegant of the recent clutch of philosophical treatments of attention is Mole's *Attention is Cognitive Unison*. At the heart of the book is the claim that much scientific work on attention suffers from the fact that theorists fail to appreciate the metaphysical category to which attention belongs. Most theorists adopt a process-based conception of attention, according to which attention is assumed to be a distinctive kind of mental process. Against this view, Mole argues that attention should instead be conceptualized in adverbial terms. To take his main example, saying that someone's actions were done *hastily* is not to advert to a particular process that might be implicated in those actions, but to describe the *manner* in which their actions were executed. Similarly, according to Mole, describing a person's behaviour as attention-involving is not to refer to a particular kind of perceptual or cognitive process that might be implicated in that behaviour, but to the way in which their behaviour was carried out. What exactly is it to perform a task in an attentive manner? Mole's answer is given in the title of the book: the attentive performance of a task is one that involves 'cognitive unison'.

   One can think of the book as having three central aims: to show that the adverbial approach to attention is more plausible than the process-based approach, to develop and defend the idea that attention can be identified with cognitive unison, and to explore the implications of this unison-based account of attention for questions relating to the role that attention might play in an account of consciousness and the grounding of mental content. We shall restrict our comments to the first two of these aims.

   Let us begin with the contrast between process-based and adverbial accounts of attention. First, what exactly is the difference between a process and a manner of occurrence? Processes are defined by Mole as 'the gaining and losing of properties by *objects*' and adverbial phenomena as 'the gaining and losing of properties by *events*' [29]. Mole then distinguishes process-based accounts from adverbial accounts of

attention by means of the following supervenience principle: 'for any two sets of cognitive events that are the same with regard to the processes that they instantiate, if one is attention constituting, then the other is attention constituting too' [33–4]. Whereas process-based accounts are committed to this principle, adverbial treatments of attention reject it.

Mole then employs this supervenience principle to argue against process-based accounts of attention. He argues that it's possible for there to be two subjects that employ the same psychological processes but where only one of them is attending. To take his main example, Mole claims that there is strong evidence that under specific experimental conditions an agent's attention is *wholly constituted* by processes of feature-binding [37–8]. Feature-binding is consequently a good candidate for an attention-constituting process. However, Mole also notes that patients suffering from unilateral neglect appear to be able to bind together visual features presented in the neglected side of their visual field. In other words, neglect patients seem to be able to carry out feature-binding in the absence of attention. He takes this to indicate that attention fails to supervene on feature-binding, and thus to put pressure on the process-based account.

There are a number of points at which this argument might be questioned. First, it is far from obvious that the kind of feature-binding that can take place outside of attention in the context of unilateral neglect is the same kind of feature-binding that requires attention in cognitively unimpaired subjects. The evidence indicates that the kind of binding that requires attention involves integrating features that belong to different dimensions (say, colour and shape), whereas it's possible that the kind of binding that can occur outside of attention is restricted to features that belong to a single dimension (say, form). Second, it is not even clear that the process-based theorist would want to say that attention is constituted by the binding of presented features, as opposed to being a process that is independent of, but causally implicated in, feature-binding; in fact, those who associate attention with feature-binding typically think of feature-binding as one of the *functions* of attention.

Finally, it isn't entirely clear just what *kind* of argument Mole intends to mount against the process-based theorist, in particular whether it is intended as an empirical argument or as a conceptual argument. Because Mole offers experimental evidence against the proposal that attention supervenes on the process of feature-binding, one might conclude that it is intended as an empirical argument. Yet he also adduces hypothetical cases, the mere conceivability of which is said to pose a challenge to the idea that facts about attention supervene on facts about cognitive processes [36, 38–9]. The status of these hypothetical cases is somewhat unclear. They may simply be intended as empirical speculations that Mole offers in lieu of the relevant evidence. Alternatively, he might regard the mere coherence of such cases as sufficient to establish that attention is not a type of process.

Setting the above worries to one side, Mole provides advocates of process-based accounts of attention with much food for thought. Reflection on actual and hypothetical cases leads him to conclude that attention cannot be understood simply in terms of which psychological processes the subject is carrying out, but must also take into consideration the nature of the subject's task and how psychological processing bears on this task. This opens the way for an adverbial analysis of attention according to which the most basic facts about attention are facts about the manner in which cognitive processes occur.

Let us turn now to the second of Mole's aims: that of developing and defending the proposal that attention can be identified with a kind of cognitive unison. His conception of cognitive unison turns out to be negative in character; it is more a matter of what the agent is *not* doing than what they *are* doing. More specifically,

cognitive unison occurs when there is an absence of task-irrelevant processing among the cognitive resources comprising an agent's 'background set', i.e., the set of cognitive processes that an agent can, according to their understanding of their current task, potentially use to perform that task. Significantly, the number of cognitive processes within a task's background set will typically be greater than the number of processes that the agent actually uses in performing that task [67]. Unison is instantiated just in case the processes making up the background set for the agent's current task—including those that only *potentially* serve the task—are not doing anything that does not contribute to its performance. They need not be doing anything at all, but they cannot be occupied elsewhere.

There is much that is attractive about Mole's development of the cognitive unison model. For example, he shows that the approach provides an illuminating treatment of divided attention. According to his analysis, divided attention is not to be explained by appeal to limitations in information-processing capacity *per se*, but is instead a function of the degree to which the background sets for the tasks in question intersect [74–82]. If the background sets for two tasks do not significantly overlap, then the subject may effectively be able to divide attention between them. If, on the other hand, the background sets significantly overlap, then dividing attention between the two tasks will not be possible. Because it is rare for two tasks to have entirely non-intersecting background sets, it is also rare for a subject to be able to fully divide attention between two activities.

Mole also provides an extended and highly illuminating defence of the idea that the cognitive unison account of attention is compatible with our everyday ascriptions of causal efficacy to attention. The core of his defence lies in an appeal to a counterfactual analysis of causally relevant properties—a variant on the counterfactual analysis of causation—which, he argues, is well suited to explain the way attention behaves in commonsense explanations of behaviour. In particular, the counterfactual analysis is able to capture the way that attention typically occupies the causal 'background' in commonsense explanation, functioning as an enabling condition rather than a triggering or mechanistic cause of behaviour.

Although there is much in Mole's approach that we find attractive, we have one worry. To bring this out, we will begin by voicing a different concern one might have with his account, but which on reflection turns out to be unfounded. In particular, one might object that the negative requirement of cognitive unison—that the resources in an agent's background set cannot be doing anything that does not serve the agent's task—is too weak. Suppose, for example, that just before an agent attempts to perform a difficult or 'cognitively fragile' task, i.e., one whose successful execution requires the subject to be paying attention [90, 114], we lesion the regions of the subject's brain underlying the resources in the background set that will not actually be called upon in performing the task. This would seem to be one way of bringing the subject's cognitive processing into unison and thus of ensuring the successful performance of attention-demanding tasks. We take it that this would be a counterintuitive conclusion.

In fact, however, Mole isn't committed to the above conclusion. This is because his definition of what it is for a cognitive process to 'potentially serve' a task, and therefore to be included in the agent's background set for that task, requires that the resource in question be available to be brought online [60–1]. And our imagined lesioning scenario doesn't satisfy this requirement. But we may now ask: why should the mere availability of a resource—one that is never *actually* deployed in performing one's task—be what distinguishes an attentive performance from an inattentive performance and, further, be what's causally relevant to successful performance of attention-demanding tasks? Though Mole addresses worries similar to this one [103

ff.], we would have welcomed a more thorough explanation of why one should expect cognitive unison to improve subjects' performance in difficult or attention-demanding situations.

Although we have raised doubts about some of the details of Mole's analysis of attention, we have no doubt that this is an important book. Indeed, we hope that it will foster the approach to the study of attention that characterized the late nineteenth and early twentieth centuries, in which the boundaries between the disciplines were not as imposing as they have become.

Aaron Henry and Tim Bayne
University of Oxford

Sider, Theodore, *Writing the Book of the World*, Oxford: Clarendon Press, 2011, pp. 336, £30 (hardback).

Ted Sider has given us a tool: *structure*. Like Platonic properties or Lewisian concrete worlds, Siderian structure provides us with a tool for analysing a wide range of metaphysical phenomena and provides a lens through which to view metaphysics itself. 'Metaphysics', says Sider, 'is about the fundamental structure of reality. Not about what's necessarily true. Not about what properties are essential. Not about conceptual analysis. Not about what there is. Structure' [1].

Structure is ideological, not ontological. It concerns the concepts we employ, not the entities which exist. Specifically, it concerns whether the concepts we employ carve reality at its structural joints. Reality comes ready-made with structure. Some of our concepts latch on to this structure better than others. Better than *grue* and *bleen* are *blue* and *green*, but even they aren't perfectly joint-carving: 'To carve perfectly, one must use the most fundamental concepts, expressing the facets of reality that underlie the colours' [5]. The perfectly joint-carving concepts are those from physics, mathematics, and logic, e.g. *up-spin*, *there exists*, *and*. Sider argues for an objective, non-psychological, understanding of structure. Ideology is as much a part of a theory's worldly content as its ontology. Ideology done well (i.e. in joint-carving terms) makes for good theories. We aim at joint-carving for the same reason we aim at truth: so that our beliefs conform to the world.

Sider will be criticized for not providing us with a firm enough cognitive grasp on what structure is and for not adequately defending the claim of its existence. Such criticisms are misplaced. Anyone who understands Lewisian naturalness should be able to understand Siderian structure. It differs from Lewisian naturalness in two ways: (i) It is ideological rather than ontological. Rather than asking, 'Is object *o* perfectly natural?', one asks, 'Does concept *c* carve reality at its joints?' Whereas Lewisian naturalism is entity-based, structure need not be hostage to entities: '[T]he idea that the world has a distinguished structure—that electrons go together and not together with cows, that it is better to think in terms of electrons than in terms of electrons-or-cows, and so on—does not require this [entity] assumption' [85]. (ii) It raises questions of naturalness (i.e. joint-carvingness) for grammatical categories other than the predicate. In addition to asking, 'Does predicate *p* express a perfectly natural property (i.e. is predicate *p* perfectly joint-carving)?', it asks, 'Is function *f* perfectly joint-carving? Is quantifier *q* perfectly joint-carving?' Structure is just Lewisian naturalness extended to ideology. This shift allows one to ask about the