

## WHAT IS ‘MENTAL ACTION?’\*

[penultimate draft; final and authoritative version forthcoming in *Philosophical Psychology*]

*Abstract.* There has been a resurgence of interest lately within philosophy of mind and action in the category of mental action. Against this background, the present paper aims to question the very possibility, or at least the theoretical significance, of teasing apart mental and bodily acts. After raising some doubts over the viability of various possible ways of drawing the mental act/bodily act distinction, the paper draws some lessons from debates over embodied cognition, which arguably further undermine the credibility of the distinction. The insignificance of the distinction is demonstrated in part by showing how the focus on ‘inner’ acts hampers fruitful discussion of Galen Strawson’s skepticism of mental agency. Finally, the possibility is discussed that a distinction between covert and overt action should supplant the one between mental and bodily action.

### I

Recent discussions recognize that philosophy of mind and action has been unduly preoccupied with physical or bodily acts in the past decades. The continued omission of mental action is said to have impoverished debates in action theory and beyond, and a corrective is offered. For example, in the Introduction to a recent volume that has been instrumental in bringing the topic of mental action back to the fore, Matthew Soteriou states that the papers collected show some of the “various ways in which a focus on mental action can contribute to our understanding of a range of different issues and themes — in the philosophy of action, the philosophy of mind, and epistemology” (Soteriou 2009a: 1). The present paper rejects Soteriou’s enthusiasm as unwarranted. No doubt drawing attention to certain act-types that have been previously ignored is in itself a welcome development. But the significance of focusing on a demarcated class of mental acts is dubious. It is unclear that a distinction between mental and bodily acts can so much as be *drawn* in a fruitful and stable way. *A fortiori*, such a distinction carries little explanatory value.

---

\* I am grateful for the very helpful comments I received from audiences at the ‘Mental Action and the Ontology of Mind’ conference at London’s Institute of Philosophy, and a workshop on intentional action at Tel Aviv University. Special thanks to my respective commentators there, Maria Alvarez and Elvira Di Bona, for their careful and insightful remarks. Thanks also to John Hyman, Noa Leibowitz, Lucy O’Brien, Matt Soteriou, and an anonymous referee for comments and discussions that vastly improved this paper. Research for this paper was supported by the Israel Science Foundation (grant no. 1120/17).

In fact, a focus on mental action can serve to obscure the deeper underlying issues. Properly coming to terms with the place of mental action, then, should actually lead us to dispense altogether with ‘mental action’ as a meaningful category — at least as this term is generally understood. (Unless otherwise stated, all ensuing occurrences of ‘mental action’ should be read as ‘*intentional* mental action’. There are likely also *nonintentional* mental acts, done accidentally and without any purpose or intention in mind. But these can be safely ignored here since, first, for obvious reasons, it is the intentional variety that attracts far more attention in the literature; and second, unintentional acts may be subsumed under the intentional, given relevant alternative descriptions of the act (Davidson 1980, essay 3).

The specious distinction between mental and bodily action is rooted in the entrenched and arguably wrongheaded Cartesian ‘internalism’ about the mind in general — the view on which mental states, attitudes, processes, experiences etc. supervene on one’s bodily state. ‘Internalism’ in this sense contrasts with so-called ‘vehicle externalism’ (Hurley 1998) – the thesis that some mental attitudes (i.e. vehicles of mental content) are partly constituted by stuff in one’s environment – rather than with ‘content externalism’. But externalists in the present sense include also, among others, Williamson 2000 and Gibbons 2001, who may or may not be aptly classified as vehicle externalists. Williamson and Gibbons argue that *knowing* – which partly supervenes on the environment – is nevertheless a genuinely mental state.<sup>1</sup>

Internalism so understood has been the target of sustained attack in recent decades, which has managed to diminish somewhat its near-universal adherence. But as the renewed interest in the time-honoured distinction between mental and bodily acts reveals, at least within the domain of action and practical agency, internalism is still very much the orthodox position. At least among the mainstream who endorse physicalism (more on this in a moment), mental action is typically seen as equivalent to action performed ‘inside one’s head’, so to speak, illustrated by such examples as visualizing images before the mind’s eye or engaging in silent inner monologue. The present paper challenges this piece of orthodoxy.

---

<sup>1</sup> It should be clear that the use of ‘externalism’ and ‘internalism’ in the text is independent of how these terms are commonly used in epistemology and moral psychology.

Its central lesson is captured by the following disjunctive claim: Either the distinction between mental and bodily action is specious, or else it is of little explanatory value.

The argumentative strategy proceeds by making a case for each disjunct in turn. Sec. *II* aims to collapse the distinction between mental and bodily action. Several stages of the argument here turn on accepting some form of physicalism about the mind (those who deny physicalism of any stripe will anyway tend to reject the supervenience of the mental on the bodily and so, for their own reasons, need not commit to ‘internalism’ as defined above). It is argued that if we accept that mental events comprise bodily events, it is hard to see how the distinction between mental and bodily acts is meant to be drawn. Secs. *III* & *IV* continue the assault on the tenability of the distinction. The pair of sections cumulatively draw out and incorporate some key lessons from the sizeable literature on so-called Embodied Cognition, and specifically the doctrine of the Extended Mind, according to which mental attitudes are sometimes constituted by stuff external to the subject’s head (e.g. her limbs) and more broadly her body (objects in the environment). While the idea of extended states of mind may be problematic, as some have argued, a surrogate thesis proposing a more conservative extension of only acts of mind is shown to be considerably less objectionable.

For the reader not convinced by that stage of the futility of trying to tease apart mental and bodily acts, Sec. *V* waves the previous objections to the feasibility of this manoeuvre but casts doubt over its theoretical significance. The section highlights the parity of explanatory role between instances of the two categories, and demonstrates how one central controversy — surrounding Galen Strawson’s skepticism about mental agency — is in fact ill-served by an arbitrary focus on internal mental action. Finally, and much more briefly, the section explores the suggestion that an intensionally different yet extensionally equivalent distinction between covert and overt action is a viable and fruitful means of carving up the terrain. Some doubts are raised about the alleged equivalence to the mental act/bodily act distinction, but the possibility is left to stand that contrasting covert and overt action can be illuminating and conducive to fruitful debates. Section *VI* then rounds off the discussion.

## II

Start by noticing the following point. From a pre-theoretical point of view, it is actually unclear how the distinction between mental and bodily action is supposed to be drawn.<sup>2</sup> A first naïve stab is that bodily acts do, whereas mental acts do not, involve bodily motions. But this cannot mark the feature that tells the two apart, since motions of the *brain* are of course involved in performing mental acts. Appealing to one's *limbs* as the circumscribed bodily region involved only in bodily acts does not help. For, first, the type of organ involved in performing an act seems a rather thin basis on which to rest a purportedly consequential distinction. Second, and more decisively, it is simply false that only limbs are involved in bodily acts. Those tempted by this suggestion need only think of such acts as pulling a face or sticking one's tongue out.

A different tack focuses instead on the use of musculature. On this suggestion, bodily acts are distinguished from mental acts by involving muscle contractions. However, the proposal struggles to accommodate certain mental acts of perceptual attention. Consider *looking*. Staring idly into the distance, one is directed by a bystander to notice something on the rooftop straight ahead. As one obliges, and so intentionally turns one's visual attention to the area, muscles in one's eye contact to focus the gaze. One performs a mental act of looking at or attending to some object in a way that involves muscle contractions, contrary the present proposal. The proposal thus fails to identify a disanalogy with bodily acts: Both may be performed by contracting the relevant muscles.<sup>3</sup>

Friends of the mental act/bodily act distinction may do better by appealing to something like the following principle:

**[BvsM]** *Act type V is a bodily act-type iff it involves the agent of V moving her body; Otherwise, V is a mental act-type.*

---

<sup>2</sup> Compare Soteriou (2009b: 232-3), who briefly makes a related point.

<sup>3</sup> The discussion so far may raise the possibility of drawing the distinction between mental and bodily acts in the following way. Perhaps bodily acts *necessarily* or *constitutively* involve motions of the body, whereas mental acts do not necessarily do so. This proposal essentially concedes the point that the mere lack of bodily movements is not a distinguishing feature of mental action. It also leads to the suspicion that the latter is not a *deep* feature of mental acts — a point argued for in section V below.

[BvsM] is open to different interpretations depending on the theory of action one subscribes to. On the orthodox ‘causalist’ reading (Davidson 1980; Searle 1983; Velleman 1992; Setiya 2007; D’Oro 2012), the subject’s ‘moving her body’ would be understood in causal-reductive terms to mean that the subject’s mental state — her desire or intention — causes (in the right way) the corresponding motion of her body. Causalists would therefore give the principle something like the following gloss:

**[CAUS BvsM]** *Act type  $V$  is a bodily act-type iff it involves the occurrence of a bodily event caused (in the right way) by the relevant rationalizing mental state(s) of the agent’s; Otherwise,  $V$  is a mental act-type.*

As stated, however, [CAUS BvsM] fails, as its first clause applies not only to bodily but to *mental* action as well, given how the production of the latter involves the causation of brain motions. (It should be clear that this failure has nothing to do with the notorious problem of ‘deviant causal chains’. The causation of mental events ‘in the right way’ also of course involves motions of the brain.)

The suspicion might arise that this problem is an artefact of a gratuitous failure to specify the connection that must obtain between the intention and the event it causes. The thought is natural enough. After all, the way causalists propose to understand ordinary acts of e.g. raising one’s arm is precisely as cases where one’s intention to raise one’s arm causes one’s arm to rise; only an intention with content that ‘matches’ in this sense the event it causes could rationalize the act. Hence, perhaps it is a mistake to formulate [CAUS BvsM] as above in terms of ‘a bodily event caused...’. But how *should* the principle be reformulated to capture the requisite matching? The answer will plainly not come from suggesting that a bodily act-type  $V$  involves the occurrence of a bodily event *of  $V$ -ing*. One cannot purport to explain the act of raising one’s arm as causally involving the event of raising one’s arm. An alternative proposal has it that bodily act-type  $V$  involves the causation of the event that *constitutes  $V$ -ing*. But this proposal could only be made to work at the cost of rejecting physicalism. For, if some form of physicalism is true, mental act-types will also involve the causation of (bodily) events that constitute  $V$ -ing. And it is not clear how

else to secure the match in question in a way that would prevent mental action from conforming to the definition of bodily action. The interim conclusion could thus be stated by saying that at least if some form of physicalism is true, it is unclear how to distinguish mental acts from bodily acts.

Theories that do not, unlike causalism, countenance a reductive understanding of action and agency may be immune to the above problem. Consider the idea of agent-causation, on which nothing short of the agent herself could be the cause of her acts (Bishop 1983; Markosian 1999, 2012; O'Connor 2000; Mayr 2011; Brent 2017). Theories that subscribe to this idea will tend to uphold the original principle [B<sub>vs</sub>M] and reject its reductive surrogate [CAUS B<sub>vs</sub>M], which may allow them to sidestep the problem. For they could then concede that mental act-types involve the causation of bodily events in the brain, yet crucially deny that such events are caused *by the agent*. Rather, they may wish to insist that the brain event is caused by some sub-agential mechanism. However, in the present dialectical context, this result is not so much a point in favour of the doctrine of agent-causation as a strike against the distinction between mental and bodily action. Recall that the point being presently made is that there is no clear *pre-theoretical* way to distinguish the two kinds of act. Our credence in the reality of the distinction should rise if it were applicable independently of theory choice, and correspondingly drop otherwise. And it now appears that the distinction does indeed fail to exhibit theory neutrality. While agent-causationists may be able to make sense of it, the (orthodox, very widely received) causalist position struggles to do so, as we have seen.

The next section sets out further reasons to regard the distinction as illusory. These reasons emerge from reflecting on the 'Embodied Cognition' research programme.

### III

A central strand within the much discussed Embodied Cognition approach is the 'extended mind' doctrine (henceforth, ExM. For defenses of the view, see Clark & Chalmers 1998; Wilson 2004; Menary 2007; Clark 2008). The striking doctrine proposes that cognitive states and processes are sometimes

partially realized<sup>4</sup> by pieces of external physical machinery in the subject's environment, thereby obliterating the traditional boundary of skin and skull as what circumscribes the mental. For example, ExM states that dispositional *beliefs* are sometimes stored in physical artefacts, devices, and gadgets. Other views, such as 'embedded cognition' and 'situated cognition', are closely related to ExM and the broader Embodied Cognition programme (Robbins and Ayedale 2009; Dawson 2014; Shapiro 2011, chs. 3-5). These latter views are somewhat less contentious, primarily insofar as they do not countenance the *constitution* of mental attitudes by objects outside the brain or even the body, as ExM does. Rather, they 'merely' highlight some crucial ways in which mental and cognitive processing seems inextricably entangled with the workings of the body — a position which in itself represents a deep challenge to standard philosophical and cognitive-scientific theorizing, which are built around the computation of inner representations.

Drawing on these somewhat milder versions of embodied cognition may well be sufficient for the purposes of the present discussion. Nevertheless, in what follows, the focus will be on ExM, mainly given the aim of exploiting some of the particular argumentative resources it deploys to shore up its provocative conclusion. As will become clear, the point here is *not* to defend ExM but rather to highlight some insights that can be gleaned from its supporting arguments — without necessarily fully endorsing their intended conclusions — in a way that undermines the serviceability of the mental/bodily action distinction.

\* \* \*

A key argument in Clark & Chalmers' (C&C) classic case for ExM turns on what they dub the Parity Principle (1998: 8):

---

<sup>4</sup> There is a potentially significant issue concerning the formulation of ExM that rarely receives due attention in the literature. This is the question of the precise metaphysical relation that ExM claims holds between the mind and the environment: Is it constitution, supervenience, or what? Clark & Chalmers 1998 is itself not sufficiently clear on this matter. Chalmers 2008 quickly dismisses the idea that ExM requires any substantial metaphysical commitments, while Shapiro 2011 (Ch. 6) seems to opt for the constitution option without much argument. In what follows, I shall bypass this issue and invoke *constitution* or *realization* as the operative relation in the rough-and-ready formulations of ExM adverted to, in keeping with most discussions. But when presenting the views defended here (which, as will become clear, are not equivalent to ExM) the relation invoked will be supervenience, which seems to be the more apposite suggestion.

If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is (for that time) part of the cognitive process.

Comparing several pairs of scenarios illustrates how PP is put to work in the service of ExM. Thus Inga, a normal and healthy resident of New York, hears about an interesting exhibition over at the Museum of Modern Art, and decides to visit. She thinks for a minute, recalls that the museum is located on 53rd Street, and sets out in that direction. Inga is then compared to Otto, an Alzheimer's patient who regularly uses his notebook to store and retrieve valuable information, treating it as a substitute for his failing biological memory. When Otto hears about the exhibition and decides to go, he consults his notebook which says that MoMA is on 53rd St. He heads in that direction and goes into the museum. C&C suggest that if one allows that Inga has a (dispositional) belief that MoMA is on 53rd St. then by PP, one should also allow that Otto has a similar dispositional belief stored in his notebook. Another famous scenario wielded by C&C involves an ordinary Tetris player who uses her imagination to rotate the geometric images sliding down her screen, to see if they fit into vacant slots in the emerging structure. The scenario is then compared with another player whose device includes a button which rotates the shapes for her. Once again, the suggestion is that, by PP, the second player is (also) undergoing a mental process, one partly realized by a physical device.

C&C lay down several criteria that scenarios such as Otto's must meet if they are to yield extended mental states and processes:

- 1) "That the resource be reliably available and typically invoked."
- 2) "That any information thus retrieved be more-or-less automatically endorsed. It should not usually be subject to critical scrutiny. [...] It should be deemed about as trustworthy as something retrieved clearly from biological memory."
- 3) "That information contained in the resource should be easily accessible as and when required." (Clark 2010: 46)<sup>5</sup>

---

<sup>5</sup> C&C (1998: 17) tentatively adduce a fourth criterion, discussed below.



Still, even with these criteria in place, the scenarios are unlikely to impress those who do not see the appeal of PP in the first place. Why accept the principle? Several commentators, friends as well as critics of ExM, see it as underwritten by an allegiance to functionalism, given how it seemingly invokes functional role as the mark of the mental. In fact, however, it is possible to view PP more innocuously as supported by a piece of quasi-functional reasoning, whereby an attitude that displays functional sameness with some paradigmatic mental attitude carries a *defeasible presumption* of mentality, absent some salient difference that could act as a defeater (cf. Chalmers 2008: xv). Put differently, functional role is at least *one* important mark of the mental. Framed this way, PP carries a much thinner commitment, which falls considerably short of full blown functionalism and should hence appear less pernicious to its non-functional detractors. And importantly for present purposes, if PP is allowed to stand, it implies that a whole range of acts fail to align neatly along the traditional mental/bodily divide.

To see this, start with the original scenarios from C&C. The second Tetris player appears to be *imagining* or *visualizing* in part by pressing the ‘rotate’ button; and Otto emerges as *recalling* the location of MoMA — two instances of mental acts partly constituted by physical machinery.<sup>6</sup> For some different examples, consider doing a calculation on a piece of paper (or a calculator) instead of in one’s head (Wilson 2004); developing one’s ideas by writing or talking rather than just by internal thought; deliberating on paper (writing down the pros and cons of *V-ing*); talking to yourself silently vs. out loud (or again, silent speech with and without *mouthbing*); ordinary reading by a sighted person vs. a blind person reading Braille (or again, a sighted person reading with and without the help of eye saccades); watching with a naked eye vs. using a magnifying glass or a microscope; etc. etc. Notice how the above sample of examples go well beyond ExM’s extension of dispositional beliefs and cognitive processes. The possibility

---

<sup>6</sup> One may wish to deny that these are genuine instances of mental acts, insisting that they are better understood as some amalgam of mental and non-mental processes/events/acts. This objection is discussed in Sec. *IV* below.

of extended practical *deliberation* was already noted. Similarly, while one can make a mental note — *decide*<sup>7</sup> — to  $V$ , one can also do so on a piece of paper or an e-calendar.

Consider Grotto, who only settles on  $V$ -ing once he makes a physical note to  $V$ ; a purely internal decision to  $V$  does not display for Grotto the functional attributes traditionally thought of as characteristic of intention, viz. temporal and psychological stability, a reliable connection to action, a basis for intra- and inter-personal planning and deliberation, and so on. It seems then that, by PP, Grotto's scribbling his note to  $V$  amounts to his deciding to  $V$ . Indeed, the idea of extended decisions should seem natural enough to the many adherents of the received planning conception of intention, as developed most comprehensively by Bratman (1987). For they need only reflect on the way in which we regularly settle on the route to our destination by means of such software as Google Maps, Waze etc.

As noted, C&C and other champions of ExM take the argument from PP to demonstrate that, at least if criteria (1) - (3) above are met, a subject's mental *states* — primarily, her dispositional beliefs — may be constituted by extra-bodily objects and artefacts. This is a highly contentious thesis, which faces some serious objections (some of which are discussed in the next section). However, recall that the point of retracing the reasoning behind the argument from PP is importantly *not* to defend ExM, but rather to question the utility of teasing apart mental and bodily acts. Given this different focus, a somewhat more conservative take on what the argument from PP actually shows would serve present purposes equally well.

One such take (another is discussed in sec.  $V$  below) consists in the idea that it is only the dynamic, not the stative, aspects of the mental economy that are unbounded by skin and skull, occasionally encompassing more than what goes on within the brain or the body. Thus the mental acts of deliberating or calculating can consist partly in scribbling on a piece of paper; 'inner' monologues can consist partly in moving one's lips; deciding can consist partly in hitting the 'go' button on a smartphone; and so on. Importantly, countenancing such extensions of mental acts need not incur a commitment to the

---

<sup>7</sup> It is assumed here that decisions are *bona fide* acts. Even as staunch a denier of (most) mental activity as Strawson (2003) allows that this is true of at least some decisions. However, if the reader disagrees, she should feel free to delete this example from the present list of extended mental acts.

corresponding extension of mental states. There need be no pressure to view one's intention to walk to the office via the park, for example, as partly constituted by the navigation software or the device on which the software runs. To be sure, the *information* regarding one's intended route could be stored in the device, which may raise the question of whether one's *belief* that one is walking via the park is constituted by the device. The present suggestion need not take a stand on this issue. Taking another example, when calculating the product of 233 and 57, one's act could extend to incorporate the physical means of pen and paper without also implying that any belief with the content  $\langle 233 \times 57 = 13,281 \rangle$  is similarly extended.

ExM is evidently taken by the bulk of its supporters and critics alike as promoting the extension of *both* mental states and mental acts or processes (though see Carter et al. 2016). In contrast, the present suggestion proposes a more moderate extension of only the latter category. On this suggestion (call it Mod-ExM), *some mental acts supervene<sup>8</sup> on the subject's extra-bodily environment*. As noted briefly above, Mod-ExM encounters less intuitive resistance and is overall theoretically more economical by comparison. The next section demonstrates this claim by looking at how Mod-ExM fares with respect to some of the most challenging objections that have been levelled against ExM.

#### IV

One way to quickly verify that Mod-ExM is indeed the more palatable of the two theses is by considering the reluctance felt by many towards extending the mind in the sense ExM proposes, on the grounds that doing so entails a counterintuitive corresponding extension of the *self*. For example, if Otto genuinely has beliefs which contents are stored in his notebook, then it seems the notebook should be considered a part of him as much as a part of his mind. And that seems outlandish.<sup>9</sup>

---

<sup>8</sup> See n. 3 above.

<sup>9</sup> Rupert 2004. C&C 1998: 18 allow that the self may extend with the mind.

Now notice that, whatever the force of this as an objection to ExM, it does not seem to find a parallel grip on Mod-ExM. This is because, once again, the objection arises as a consequence of the feature of ExM which Mod-ExM crucially does not endorse, viz. the extension of states of mind. If an agent performs a calculation using pen and paper, this is understood by Mod-ExM as an extended mental act of hers. But this understanding creates no intuitive pressure to suppose that the pen and paper are in some sense a part of the agent. Not every object one picks up becomes a part of oneself. A prosthetic arm is plausibly a part of a subject whose biological arm it (permanently) replaces, whereas a prosthetic arm that a different subject tries on for size and then puts down is not. The intuitive asymmetry seems to stem at least in part from the physically and functionally *stable* integration of the arm with the former subject's body, compared with the physically and functionally *transient* integration of the arm with the latter subject's body.<sup>10</sup>

What goes for physical extensions of the body should go for mental states, attitudes, etc. Mental states – in particular, alleged dispositional beliefs like Otto's – are typically rather stable and often long-term properties of subjects. Hence, they are naturally seen as constitutive of the persistent core that defines the person who holds them. Faced with the Otto scenario, this then prompts the (intuitively disagreeable) suspicion that, if the intended lesson of the scenario is conceded, the notebook must be seen as quite literally a part of him. But mental *acts*, unlike mental states, are far less stable and robust: They are transient entities, possibly occurring very rarely and sometimes just the once; even when they recur, they tend to be performed in different ways, on different objects, for different purposes, and so on. They are much more naturally seen as things a person (sometime) does, not part of who she is. And if there *is* some recognizable sense in which one's deliberating is a part of one, this seems to hold equally for bodily acts such as one's building a tree-house. Countenancing the extension of mental acts, then, does not underwrite a similar suspicion that Mod-ExM is committed to outlandish ideas about personal identity.

---

<sup>10</sup> I'm grateful to an anonymous reviewer for getting me to clarify this point.

Consider next the natural concern raised by Gertler (2007), on which the state of the coupled entity comprising Otto and his notebook cannot be a genuine belief since it lacks the essential mark of *first personal access*. After all, Otto uses *perception* to look up the location of MoMA in his notebook, just like any observer would. Now however troubling this objection actually is to ExM, it does not seem to threaten Mod-ExM. For as a matter of fact, it is very widely (though not universally) held that we *do* have something like privileged and peculiar access to our intentional acts, including our ordinary bodily acts. It remains a hotly debated question how to best explain what is sometimes referred to as our ‘practical knowledge’ of what we are doing when we act intentionally; but it is generally agreed that the phenomenon itself is real enough.

Another potentially serious objection to ExM is the worry about so-called ‘cognitive bloat’. Detractors charge ExM with promiscuously allowing the mind to leak very extensively into the world and encompass any number of contents, rendering the mind an unrecognizably huge and sprawling entity. For example, those who share this concern point out, as a *reductio*, that if we ascribe to Otto the belief that MoMA is on 53rd st., we should be equally happy to ascribe to someone with regular access to a phonebook or a phone directory true beliefs about every phone number listed therein (Rupert 2004).

C&C anticipate this concern, and offer in response criteria (1) - (3) cited above by way of restricting the type of externally stored content that could qualify as one’s dispositional belief. It is particularly a fourth criterion, which C&C tentatively endorse, that is meant to block the absurd implication whereby a subject could have “beliefs about all of the phone numbers available to him through directory assistance ... so long as he remembers how to dial up the operator” (Rupert 2004: 403). According to this fourth criterion, which Rupert dubs ‘past-endorsement’, “the information in the [external resource] has been consciously endorsed at some point in the past, and indeed is there as a consequence of this endorsement” (C&C 1998: 17). However, the past-endorsement criterion faces at least two worries. First, as C&C themselves acknowledge (hence their *tentative* endorsement of the criterion), we sometimes seem to acquire ordinary beliefs through non-conscious processes. And second, the criterion may actually run against the spirit if not the letter of ExM, since “if an extended (or any) belief requires conscious

endorsement in order to be a genuinely held belief, and conscious endorsement is ultimately an internal process ... then the traditional subject is privileged in a deep sense, after all” (Rupert 2004: 404).

Be that as it may, however, once again there does not seem to be an equally problematic version of the present objection for Mod-ExM. This is because the source of the worry is ExM’s insistence that *non-conscious* states (viz. primarily dispositional beliefs) are genuine constituents of the mind. But when the focus turns to *intentional* (mental) action, non-consciousness drops out of the picture, for the simple reason that we are generally consciously aware of our intentional acts. For the same reason, Mod-ExM entirely sidesteps the controversy over whether the mental encompasses the non-conscious— a thesis which some regard as the deep mistake in ExM (Gertler 2007).

However, there is at least one perennial reaction to the idea that the mind extends beyond brain and body that seems to arise naturally enough whether the candidates for extension are mental states or mental acts. The thought here is that what ExM sees as a unified entity consisting of an inextricable compound of biological body and bits of the surrounding environment, is in fact no more than two, certainly deeply entangled and inter-dependent — yet crucially, perfectly separable — entities. Thus many would deny that Otto (or the coupled entity comprising Otto and the notebook) genuinely believes that MoMA is located on 53rd St. Rather, he is better understood as having a perfectly ordinary belief that the museum’s address is written down in his notebook, where he proceeds to look it up. This understanding sits better with common sense, is more conservative, and less disruptive of the received philosophical and cognitive-scientific outlook.<sup>11</sup> To infer with C&C that Otto does believe that MoMA is on 53rd St. is said to commit what Adams & Aizawa (2008; 2009) dub ‘the coupling-constitution fallacy’ — i.e., to confuse a causal for a constitutive relation, as they illustrate with the following nice example: “[i]t is the interaction of the spinning bowling ball with the surface of the alley that leads to all the pins

---

<sup>11</sup> Some critics (Rupert 2004, Adams & Aizawa 2009) complain that ExM renders the mind inapt for sustained scientific theorizing, undercutting the hugely successful research agendas in the cognitive sciences. This charge is not discussed here, since it targets particularly ExM with its focus on extended *cognition*, a focus not shared by Mod-ExM.

falling. Still, the process of the ball spinning does not extend into the surface of the alley or the pins” (Adams & Aizawa 2009: 83).

To pin the charge of committing a coupling-constitution fallacy on Mod-ExM, the opponent must first demonstrate that a clean break can be achieved between the inner and the outer elements of such mental acts as cited above when motivating Mod-ExM. (Demonstrating this would not yet settle the matter at hand; there would remain the question of which of the rival pictures, the casual or the constitutive, actually provides the *best* understanding of the relevant phenomenon. But demonstrating feasibility is a precondition for the objection to stick.) Now with some such acts, effecting the requisite separation can seem straightforward enough. Consider Grotto again. What has been conceived above as his extended decision might be taken as his making an ordinary ‘inner’ decision, followed by the (bodily) act of recording the decision in his calendar. But other acts are far less amenable to a similar decomposition.

For example, when developing one’s ideas by speech or in writing, the ongoing process continuously loops backwards and forwards between the ‘outer’ scribbling and the ‘inner’ thought; it is not as though one first thinks out the idea in one’s head, and then proceeds to write it down or say it out loud. The writing/speaking and the ‘inner’ thinking are far more integrated than that — indeed, so integrated as to prompt the thought that one’s writing/speaking *just is* one’s thinking (hence the expression, ‘thinking out loud’). This represents a serious obstacle for an adequate understanding of one’s behavior by piecemeal decomposition of the imagined sort. The coupled entity comprising the agent + pen and paper *itself* constantly creates its own inputs for the activity, which then feed back into the activity, thereby influencing the course it takes, which in turn generates further inputs, and so on. This makes it hard to achieve a clean break between the supposedly different stages of the process along the internal/external divide in a way that would render the process in its entirety recognizable as an integrated activity.

Consider working out which route to take from Bilbao to San Sebastian by tracing the different options on a map. One way of doing this would be to mentally visualize, ‘in one’s head’ as it were, the

different routes, plot each one on the map, compare and then choose. But one alternative (arguably more natural and effective) way this exercise might go involves one setting out some initial segment of route A on the map, which makes one realize that going that way would regrettably bypass quaint little town T, which then leads one to revise the initial segment, which in turn makes accessible several other sites one had not thought of, which then... It seems forced to insist that one's planning in this scenario only occurs away from the map. One's thought process is not linear in a way that may have lent itself more naturally to piecemeal decomposition. (This is in sharp contrast, note, with Adams and Aizawa's bowling example, where there are no similar content-generating feedback loops between ball, pins, and alley that may hamper decomposition [Shapiro 2011: 180]).<sup>12</sup>

Similar points apply to e.g. deliberating practically and reading Braille, among others. One may attempt to explain deliberating whether to  $V$  or not as consisting in such discrete 'inner' elements as intentions to deliberate by writing; reasons for and against  $V$ -ing occurring to one; an ultimate decision to refrain from  $V$ -ing, and so on; alongside such 'outer' elements as writing and erasing motions of the fingers. But once again, explaining the subject's behavior by such piecemeal decomposition risks losing sight of the mental forest for the trees. To deny that the writing subject is genuinely deliberating, and insist instead that her behaviour be understood as an aggregate of successive truncated episodes defined by their supervenience base is to distort the nature of her (non-linear) activity beyond recognition. A much more natural way to make sense of what the agent is doing is by seeing all such events as continuous steps within a broader process — in this case, the process of deliberation.<sup>13</sup> The objection from a coupling-constitution fallacy, at least as it is levelled against the extension of mental acts, thus fails. But the discussion may seem to give rise to a different, though likewise deflationary objection to Mod-ExM. Considering it helps to bring out the central lesson this paper aims to establish.

---

<sup>12</sup> I'm grateful to an anonymous reviewer for urging me to clarify this.

<sup>13</sup> Compare Wilson and Clark, 2008.



One might protest that the foregoing defense of extended mental acts risks degenerating into an idle dispute over choice of terminology. The discussion so far has been at pains to establish that reading with one's fingers is a genuine form of reading; that deliberating using pen and paper is a genuine form of deliberating; that deciding by hitting the 'go' button on a smartphone is a genuine form of deciding; and so on. But an opponent may be happy to grant all that. Indeed, where at least some of these examples are concerned, claiming otherwise may seem rather frivolous: How can one plausibly deny that some of us sometimes read sentences in Braille or develop our ideas in speech? The attribution of certain particular act-types cited above may strike the opponent as less apposite than others, but this too may be a merely verbal dispute not worth worrying about. If the claim is that deliberation can sometimes take place (partly) outside the head, the present rejoinder will be that this non-standard usage of 'deliberation' can be accommodated without giving up on the traditional distinction between mental and bodily acts, by recognizing *two* kinds of deliberation — one purely inner and another partly outer. Indeed, the opponent may have no objection to the latter kind usurping the title 'mental', as long as the intended meaning is made clear.

However, this deflationary reaction misses the deeper underlying lesson. Suppose the arguments adduced so far fail to convince that mental acts cannot be neatly distinguished from bodily acts. That is, contrary to the argument of section *II*, suppose there is in fact a viable mark of mental acts that tells them apart from bodily acts; and moreover, that this mark is theory-neutral in the sense that identifying it does not require one to subscribe to any particular theory of action. And suppose further that mental acts do not in fact extend beyond brain and body in the way argued for in secs. *III-IV*; perhaps the question of any such purported extension is no more than a verbal dispute over 'mental', or over the labels for specific act-types standardly thought of as mental, e.g. 'deliberating' and 'calculating'. Still, at least the foregoing arguments show that the distinction between mental and bodily acts *is rather shallow and that consequently, attempts to formulate it properly and put it to use are unlikely to bear constructive fruits.* (To be sure, it may be instructive to learn that the mental act/bodily act distinction is in fact untenable, and why; to that limited extent, efforts devoted to studying the distinction need not be pointless. More on this in Sec. *VI* below).

It is a commonplace, after all, that mental states and events are individuated *partly* at least by their characteristic contributions to explanations of behaviour. But the attempt to demonstrate that some mental acts extend, even if it is ultimately unsuccessful, at least reveals that mental and bodily acts exhibit key functional/explanatory similarities. The difference between reading with and reading without eye saccades, or between a silent monologue that does and a silent monologue that does not involve lip motions is very thin indeed. It is only in contrived circumstances that such differences make for a different explanation of a subject's behaviour. It seems, then, that teasing apart such acts and classifying them differently holds little if any theoretical value and should play no major role in attempts to understand action and agency, even if it can be achieved. The conclusion of the paper can hence be captured by the following disjunctive claim: Either the distinction between mental and bodily acts is specious or it is unimportant.

Of course, at this point in the dialectic, we are assuming that, contrary Sec. II, there is in fact a viable way to draw the distinction between mental and bodily acts. (Therefore, incidentally, any lessons drawn here are entirely independent of the discussion in Sec. II – and so for one thing, are neutral on the question of physicalism). And whatever else soliloquizing with and without lip motions have in common, at least this much trivially tells them apart: Only the former involves lip motions. However, this difference cannot be *assumed* to mark a meaningful basis for classification. We do not treat writing with one's left hand as belonging in a different class of acts from writing with one's right hand. Nor do we think that calculating using one's fingers belongs in a different class from calculating using pen and paper. Why then think that calculating in one's head *is* importantly different? At this stage in the dialectic, the mere difference in the body's involvement cannot in itself be treated as meaningful. Things would seem otherwise if this difference could be shown to facilitate meaningful debates or frame meaningful questions. But a brief examination of a recent influential discussion of the possibility of mental action only serves to strengthen the contrary suspicion.

On the face of it, the debate Galen Strawson engages with in his paper, 'Mental Ballistics or the Involuntariness of Spontaneity' (2003), is precisely the sort of place one might expect to find the category

of mental action doing important work. In that paper, Strawson famously defends a sceptical view of mental agency. He aims to show that (perhaps with the exception of a few negligible instances), mental action “is entirely *prefatory*, it is essentially — merely — *catalytic*” (Ibid, 231). By this Strawson means that mental action is largely limited to preliminary steps of priming, triggering, and facilitating the occurrence of mental processes, which themselves merely *happen*. Instances of the varieties cited above, widely seen as genuine mental acts in their own right, are in fact, according to Strawson, events in which the subject is largely passive. Prefatory mental ballistics aside, “[t]he rest is waiting, seeing if anything happens, waiting for content to come to mind, for the ‘natural causality of reason’ to operate in one” (233). The present aim is not to assess Strawson’s position, but rather to show that fruitful discussion of it is actually ill-served by focusing narrowly on acts performed ‘inside one’s head’, as he and his critics tend to do. This focus is not surprising, in view of the received identification of mental action with action that supervenes on one’s brain and bodily state. Nevertheless, it remains true that the important questions Strawson raises have broader application, beyond body and brain. Disregarding this fact can only obscure the deeper underlying issues.

Briefly, Strawson’s main argument takes off from the familiar point that for a mental act-type  $V$ , one typically (perhaps ever) does not and cannot have an intention with the content  $\langle \text{to } V \{ \textit{proposition} \} \Phi \rangle$ . For such an intention would plainly obviate the need to  $V$ . Thus for example, one does not intend to judge that  $\Phi$ , since this intention already seems to involve affirming that  $\Phi$ , which would obviate the corresponding judgment. Now given that the content of one’s intentions is non-spontaneous in this sense, the content must already be there, “available for consideration and adoption for intentional production”, in which case “it must already have ‘just come’ at some previous time in order to be so available” (236). This highlights for Strawson the *automaticity* of (almost all) mental processes, which conflicts with thinking of them as genuine acts. The mental processes that comprise the bulk of our mental life, and which many confuse for mental action, turn out to occur automatically, without the agent having any control or involvement in their unfolding: “Contents occur, spring up — the process is largely automatic” (229). Something like this diminished form of control over mental processes and events seems to drive many to be suspicious of mental agency.

Notice however, that the central idea Strawson aims to resist, as he himself explicitly acknowledges, is that *cognition* is ever agential, “that reasoning, thought, judgement and belief-formation are ... a matter of action” (the list should presumably be extended to include at least also *attention* and *memory*).<sup>14</sup> But if it is cognitive processes whose agential nature is at issue, an exclusive focus on their inner manifestations is unwarranted and counter-productive. After all, exactly the same issues of automaticity and control that drive Strawson to his skepticism seem to arise when we examine the outer variety, too. That is, exactly the same “springing up” of e.g. certain numbers to mind which occurs when one calculates some product in one’s head, also occurs when one does the calculation on a piece of paper. If the former is taken to undercut the agential status of inner calculations, then the latter should do the same for outer calculations. In response, one might claim that this overlooks a crucial disanalogy which speaks directly to the question of one’s control and hence one’s agency, viz. the involvement of one’s arms. The outer calculation is not automatic and is intentionally controlled, the thought goes, insofar as the agent can stop and continue her hand from writing down numbers on the page. However, a closer look reveals this disanalogy to be specious. For what the agent controls when she decides to stop or continue writing on the page is *whether or not* she is calculating. And this is a type of control she enjoys over her inner cognitive activities just as well: She can simply decide to do something else instead of calculating. While the means of exercising this kind of control are certainly different, the domain it is exercised over is exactly the same — viz. the occurrence or otherwise of the cognitive process in question. Once this point is acknowledged, any temptation to see a disanalogy with the outer case should evaporate.

In fact, the point just made could be used by Strawson’s opponents to counter his argument against mental agency. Recall Clark & Chalmers’ Parity Principle, which states that, if, when imagining the internalization of some outer process, we have no hesitation to accept it as a genuine cognitive process, then we should accept it as a genuine cognitive process. Now one could appeal to a ‘reverse parity principle’ of sorts, which says that if we have no hesitation viewing some outer process as agential, then

---

<sup>14</sup> Strawson 2003: 239. Strawson does discuss more briefly the case of decisions, but there he seems happy to concede that at least some decisions are acts, for example when deciding in scenarios of mere picking, where there are no more weighty reasons to prefer one alternative over the other. See Strawson 2003: 243-4.

(absent some salient difference), we should accept the inner process as agential, too. And given that the sort of control an agent enjoys over her outer calculation is sufficient to convince that she is genuinely acting, the corresponding control she has over her inner calculation should be similarly sufficient. This point is bound to be missed if the battleground for the debate over mental agency is drawn around the boundary of skin and skull.

\* \* \*

Various reasons have been provided above to doubt the serviceability of the mental/bodily act distinction. But perhaps the deep division is not one that contraposes the mental with the bodily at all. Perhaps a different distinction, drawn in different terms, actually *tracks* the mental-act/bodily-act, and can thus explain and vindicate the motivation for keeping the two categories firmly apart. In closing, we explore briefly a salient suggestion for a surrogate distinction along these lines — the distinction between covert and overt acts.

The proposed distinction turns on the issue of observability or third-person access. As a first pass, to say that an act is overt is to say that it is in principle observable by a suitably placed bystander; otherwise it is covert. The thought here is that when one performs an overt act, but not when one performs a covert act, the observer can see what one is doing. Covertiness does indeed seem to be a feature of the act-types that typically get cited as paradigmatic mental acts: A calculation done in one's head or a silent monologue are not as such observable. Nevertheless, the distinction is not problem-free. Notice first that, while the agent's calculating may not be observable as such, *that* she is performing a mental act *is* often observable: One can often see that an agent is absorbed in thought. Hence, for the distinction to be viable, it cannot be taken to refer to the mere observability of action taking place, but rather just to the observability of the particular act-type taking place. However, it is then no longer obvious that overtiness in this specific sense will be a feature of all bodily action.

Consider Donald, who is hard at work in his garden, sawing wood planks, hammering nails, cutting and sewing together large sheets of polytarp. An observer looks and (thinks she) sees him building a boat. In fact, however, crafty Donald is building something that is designed to look every bit *like* a boat. But

he cares not at all whether his creation sinks or floats, or how fast it would go on water etc., and does not bother insulating it properly or using the right kind of paint. For his sole aim in building it is to irritate his next door neighbour, who (Donald knows) longs to sail in a canoe of her own and would be devastated were she to believe that Donald is fulfilling the dream she is unable to fulfil herself. The example seems to demonstrate that some bodily acts can be covert, too. While there is perhaps a sense in which Donald is overtly building a boat, under such descriptions as ‘building an object that would irritate his neighbour’, his behaviour certainly seems covert. How can the proposed distinction accommodate this point?

Soteriou makes a promising proposal. An action is overt, he suggests, “if one can ... identify the bodily events that constitute it as agential without having to identify them with mental events that cannot, as such, be observed” (2009b: 234); otherwise, it is covert.<sup>15</sup> Soteriou’s suggestion seems to comfortably handle Donald’s case. It classifies Donald’s acts as overt, since the observer can identify Donald’s bodily movements as agential without having to identify them with any (unobservable) mental acts of his. The observer may not be able to discern exactly *what* Donald is doing under certain descriptions, but that is not required by the definition. However, doubts may arise about the applicability of Soteriou’s proposal when one recalls the phenomenon often dubbed the ‘broadness’ of the progressive (Anscombe 1957: 39; Falvey 2000; Thompson 2008: part II). Broadness consists in the independence of intentional action ascriptions from any particular bodily movements taking place in the narrow, more localized sense. Thus it may be true of Sarah that she is making bread, even though she is currently reading a newspaper on the sofa; she may simply be waiting for the dough to rise. The potential problem raised by the phenomenon of broadness should be clear enough: Sarah’s act is overt but the observer may not be able to identify her bodily movements as agential while she is lying supine on the sofa, as she waits for the dough.

A natural suggestion for getting around the problem posed by broadness specifies that the observer must witness the *entire* duration of the act. This may help with the sufficiency of the condition but at the

---

<sup>15</sup> It should be emphasized that Soteriou does *not* offer this as a means of distinguishing mental from bodily action across the board, but rather by way of zeroing in on the particular variety of mental acts he goes on to discuss in his chapter.

cost of sacrificing necessity, as can be seen from Donald's case. In any event, there is a further potential problem with the proposed definition, viz. the suspicion that it marks a deeply contingent feature. The range of observable phenomena is after all constantly broadening with scientific and technological advents. Specifically, brain imaging techniques such as functional MRI (fMRI) reveal an increasingly more nuanced and accurate picture of the brain mechanisms that underlie human thought and action. And in fact, recent research associates distinct neural pathways with voluntary and intentional action. It thus certainly seems possible that future work will provide means to determine that some agent's brain motions are agential without having to first identify them with some unobservable mental act of hers, contrary to Soteriou's proposal. Indeed, we are arguably already in a position to do so with respect to e.g. *decisions* (Jeannerod, 1997; Haagard 2008; Brass & Haagard 2008; Zapparoli et. al. 2017).

The issues here deserve a fuller and more careful examination than can be offered at the tail-end of this paper. But the brief discussion here seems to suggest that the covert/overt distinction is much more than a notational variant on the inner/outer theme, and consequently on the time-honoured distinction between mental and bodily acts. Indeed, the covert/overt act distinction represents a significant departure from the standard framing of the debate. Importantly, it prescribes a somewhat different direction for discussions that focus on the interplay between *these* two categories, compared with the sort of questions that tend to be at the fore in discussions of mental vs. bodily acts. Such topics as the nature of covert acts, the possible significance of their being covert, and the implications their covertness may have for understanding agency and related notions are currently much under-explored. Two exceptions to this oversight are, besides Soteriou 2009b, also Sterelny 2004, who highlights how overt information processing is susceptible to manipulation in a way to which the covert is immune; and Dennett (1991: 197), who speculates that talking to oneself in *sotto voce* may have had the evolutionary virtue of creating a useful privacy. In highlighting considerations of biological fitness, Sterelny's and Dennett's respective discussions help to bring out a further possible reason for regarding the covert/overt distinction as potentially more fruitful than the mental/bodily: The former seems to be backed up by a much more straightforward and compelling evolutionary rationale.

## VI

This concludes the central argument of the paper. Where does it leave us?

Sec. *V* argued that the mental act/bodily act distinction is shallow. One possible upshot of this claim is in the spirit of a point made by Peacocke 2007. In the course of arguing for what he calls ‘The Principal Thesis’ – viz., that “[a] thinker’s awareness of those of his mental events that are mental actions is a species of action-awareness” (2007: 358) – Peacocke suggests that, as he puts it, “[m]ental actions and bodily actions are actions in exactly the same sense”. Both types may in fact be subsumed under a common principle. (For Peacocke, this principle is the constitutive involvement of an attempt [362].) Taken in this vein, the shallowness claim of Sec. *V* could be seen as complementing and augmenting Peacocke’s general observation (while remaining neutral on the particular principle he endorses). Not only do the differences between mental and bodily action fail to expose them as fundamentally distinct kinds of event, as he maintains; indeed, those differences are explanatorily negligible – and the underlying commonality, whatever it may be, consequently much weightier.

But the argument can be taken a step further. If the differences between mental and bodily acts are merely superficial, it is unclear that fleshing out and understanding them is a project that merits close philosophical attention. Indeed, in some central cases, as illustrated above with the discussion of Strawson, such attention constitutes a potentially damaging distraction. The covert act/overt act distinction may be a much more fruitful surrogate distinction to explore.

Finally, recall that Sec. *II* argued, on independent grounds, that the mental act/bodily act distinction is not merely shallow but in fact illusory. There does not seem to be a principled way to mark the distinction – at least if one accepts some form of physicalism about the mental, as most philosophers tend to do. This conclusion seems to underscore the futility of sustained work on the distinction. It may certainly be important to learn that, and understand why, the distinction does not hold water. But the point remains that conceiving of human agency as manifested in essentially two different domains – the mental and the bodily – is a wrongheaded paradigm that should be discarded.



## References

- Adams, F. & Aizawa, K. 2008. *The Bounds of Cognition* (Wiley-Blackwell).
- Adams, F. & Aizawa, K. 2009. 'Why the Mind is Still in the Head', in Aydede & Robbins 2009: 78-95.
- Anscombe, E. 1957. *Intention* (Basil Blackwell).
- Bishop, J. 1983. 'Agent Causation'. *Mind* 92: 61-79.
- Bratman, M. 1987. *Intention, Plans, and Practical Reason* (Cambridge University Press).
- Brass, M., & Haggard, P. 2008. The What, When, Whether Model of Intentional Action. *Neuroscientist* 14: 319–325.
- Brent, M. 2017. 'Agent Causation as a Solution to The Problem of Action'. *Canadian Journal of Philosophy* 47: 656-673.
- Carter, A., Gordon, E. & Palermos, O. 2016. 'Extended Emotion', *Philosophical Psychology* 29: 198-217.
- Chalmers, D. 2008. 'Forward', in: Clark 2008: ix-xvi.
- Clark, A. (2010). 'Memento's Revenge: The Extended Mind Extended', in R. Menary (ed.), *The Extended Mind* (MIT Press): 43–66.
- Clark, A. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension* (Oxford University Press).
- Clark, A., & Chalmers, D. 1998. 'The Extended Mind', *Analysis* 58: 7–19.
- Davidson, Donald 1980. *Essays on Actions and Events* (Clarendon Press).
- Dawson, M. 2014. 'Embedded and Situated Cognition', in Shapiro 2014: 59-67.
- Dennett, D. 1991. *Consciousness Explained* (Penguin Books).
- D'Oro, J. (ed.) 2012. *Reasons and Causes: Causalism and anti-Causalism in the Philosophy of Action* (Palgrave MacMillan).
- Falvey, K. 2000. 'Knowledge in Intention', *Philosophical Studies* 99: 21-44.
- Gertler, B. 2007. 'Overextending the Mind?', in Gertler & Shapiro (eds.), *Arguing About the Mind* (Routledge).
- Gibbons, J. 2001. 'Knowledge in Action', *Philosophy and Phenomenological Research* 62: 579-600.
- Haggard, P. 2008. 'Human Volition: Towards a Neuroscience of Will'. *Nature Reviews Neuroscience* 9: 934–946
- Hurley, S. 1998. *Consciousness in Action*. Harvard University Press.
- Jeannerod, M. 1997. *The Cognitive Neuroscience of Action* (Blackwell).
- Markosian, N. 2012. 'Agent Causation as the Solution to All the Compatibilist's Problems'. *Philosophical Studies* 157: 383-398.

- Markosian, N. 1999. 'A Compatibilist Version of The Theory of Agent Causation'. *Pacific Philosophical Quarterly* 80: 257-277.
- Mayr, E. 2011. *Understanding Human Agency* (Oxford University Press).
- Menary, R. 2007. *Cognitive Integration: Attacking the Bounds of Cognition* (Palgrave Macmillan).
- O'Brien, L. & Soteriou, M. (eds.) 2009. *Mental Actions* (Oxford University Press).
- O'Connor, T. 2000. *Persons and Causes: The Metaphysics of Free Will* (Oxford University Press).
- Peacocke, C. 2007. 'Mental Action and Self-Awareness (I)', in J. Cohen & B. McLaughlin (eds.), *Contemporary Debates in the Philosophy of Mind* (Blackwell): 358-376.
- Robbins, P., & Aydede, M. (eds.) 2008. *The Cambridge Handbook of Situated Cognition* (Cambridge University Press).
- Rowlands, M. 1999. *The Body in Mind: Understanding Cognitive Processes* (Cambridge University Press).
- Rupert, R. (2004). 'Challenges to the Hypothesis of Extended Cognition'. *Journal of Philosophy* 101: 389-428.
- Searle, J. 1983. *Intentionality* (Cambridge University Press).
- Setiya, K. 2007. *Reasons without Rationalism* (Princeton University Press).
- Soteriou, M. 2009a. 'Introduction', in O'Brien & Soteriou 2009: 1-16.
- Soteriou, M. 2009b. 'Mental Agency, Conscious Thinking, and Phenomenal Character', in O'Brien & Soteriou 2009: 231-252.
- Shapiro, L. (ed.) 2014. *The Routledge Handbook of Embodied Cognition* (Routledge).
- Shapiro, L. 2011. *Embodied Cognition*. Routledge.
- Strawson, G. 2003. 'Mental Ballistics or the Involuntariness of Spontaneity', *Proceedings of the Aristotelian Society* 103: 227-257.
- Sterelny, K. 2004. 'Externalism, Epistemic Artefacts and the Extended Mind', in R. Schantz (ed.), *The Externalist Challenge* (De Gruyter): 239-254.
- Thompson, M. 2008. *Life and Action* (Harvard University Press).
- Wilson, R. 2004. *Boundaries of the Mind. The Individual in the Fragile Sciences: Cognition* (Cambridge University Press).
- Wilson, R. & Clark, A. 2008, 'How to Situate Cognition: Letting Nature Take its Course', in Robbins & Aydede 2008: 55-77.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford University Press.
- Zapparoli, L., Seghezzi, S. and Paulesu, E., 2017. "The What, the When, and the Whether of Intentional Action in the Brain: a Meta-Analytical Review, *Frontiers in Human Neuroscience* 11: 238.