

Imagination, Endogenous Attention, and Mental Agency

Tom Cochrane

Forthcoming at *Phenomenology and the Cognitive Sciences*

Abstract

This paper develops a mechanistic account of basic mental agency by identifying similarities between two of its major exemplars: endogenous attention and imagination. Five key similarities are identified: i) that both capacities are driven by currently prioritised goals that are either person-level or apt to become person-level. ii) that both deliver their outputs to the working memory iii) that both range across all and only conceptual contents; iv) that both proceed under the guidance of norms and/or habits; and v) that both directly activate rather than inhibit content. These five features are consolidated by proposing that basic mental agency is essentially the power to call conceptual content to mind and hold it within our working memories.

Keywords: mental agency, attention, imagination

This paper develops a mechanistic account of mental agency. Philosophers have mostly been content to discuss particular aspects of mental agency without specifying what, in general, it involves (exceptions include Peacocke 2007 and Levy 2022). I find this odd since a general account of mental agency will give us a clear overview of the issue. It should also helpfully contribute to investigations of autonomy, personhood, well-being, and disorders where mental agency is compromised such as OCD, ADHD, and PTSD.

Mental agency is, roughly speaking, the deliberate production of mental events such as thoughts and sensations. Examples include forming plans, imagining how one's partner may react to a gift one is picking out, or focusing one's attention upon the background ambient noise. Such mental acts can be distinguished from mental events that we do not deliberately bring about such as suddenly feeling a pain in one's back or being surprised by a loud bang. There are also mental events that may be described as failures of mental agency such as aiming to visualise a certain actor's face but finding that another actor's face keeps coming to mind

instead. The account I develop in this paper will clarify what makes certain mental events deliberate.

We might distinguish mental acts from bodily acts such as driving a car, climbing some stairs, or chewing food. Yet it is arguably the case that all bodily acts involve the deliberate production of mental events, i.e. the intentions and sensations guiding those bodily acts. Moreover, much deliberate mental activity is achieved by means of bodily activity (as accounts of embodied and extended cognition emphasise). Consider for instance fluent speech or calculating with pen and paper. To mark out the target of our analysis then it helps to make three points: First, I wish to include mental acts that are not reliant on observable bodily motion. For instance, a locked-in patient, incapable of any outward signs of bodily motion could still deliberately and successfully form thoughts.¹ Second, even if certain thoughts or feelings are involved in our bodily actions, we need not be aiming specifically at those thoughts and feelings, whereas I'm interested in cases where we specifically aim to bring about a mental event. This is what people care about when suffering disorders of mental agency such as OCD in which certain worries become intrusive, or when we please ourselves by deliberately recalling a pleasant experience. Third, there will be cases in which we aim at a mental event as the outcome of a sequence of bodily actions. For instance, we might aim at the feeling of orgasm by means of having sex. However, I am interested in those mental events that we can bring about directly, without any intervening actions. That is, I am interested in direct or basic mental agency. This, it seems to me, is foundational to accounts of agency, mental or otherwise, and so it is of considerable value to outline its mechanism. Note in particular that given the widespread claim that actions are guided by intentions (e.g. Schlosser 2019; Paul 2021) if we also mentally act in bringing about intentions then this mental act will be implicated in all other acts.

The strategy I use to explore the mechanism of mental agency is to see what two of our basic mental capacities—imagination and endogenous attention—have in common. So the main body of this essay will involve systematically examining some key features shared by these capacities, which I will then consolidate under a common mechanism. As such, exactly what is meant by these two faculties should become clearer as we go along. But for the sake of starting on the same page, here is an initial characterisation: By 'imagination' I mean that activity or process in which we bring ideas and sensory images to mind or creatively recombine

¹ With reference to this sort of point, Levy (2019) prefers to distinguish between overt and covert agency.

them. For instance, I imagine meeting my friend later and mentally rehearse some things I might say. Note that I do not identify imagination simply with imagery. Doing so would include the phenomena of involuntary hallucinations and after-images which I exclude and exclude cases of propositional or conceptual imagination which I include (I share this approach with philosophers such as Currie & Ravenscroft 2002 and Langland-Hassan 2020). Note also that I do not distinguish imagination from supposition. Some philosophers have argued that supposition has voluntary aspects that imagination lacks (e.g. Balcerak Jackson 2016), but I regard supposition as a sub-category of imagining in which certain norms of vividness are absent (more on norms in section 4). Thus, the concept of imagination that I'm working with here is a very broad one.

Meanwhile, I focus on the sub-category of attention that is most clearly under our direct voluntary control. Endogenous attention is the capacity we employ when we try to 'pay attention' to something. For example, if I pay attention to the lecture I'm listening to, I concentrate on the speaker's words and try to tune out irrelevant distractions such as the street noise outside. Endogenous attention is ordinarily distinguished from exogenous attention, in which our attention is drawn to something like a loud bang (the distinction was originally made by James 1890/1950). Thus endogenous attention is more clearly self-motivated than exogenous attention (though this distinction will be refined below).²

Thus it is a crucial assumption of my approach is that imagination and endogenous attention are exemplary instances of direct or basic mental agency. We aim to pay attention to the road and we simply do it. We aim to imagine something blue and we simply do it. I should note here that I am not very concerned with whether imagination and attention are entirely distinct and unified capacities.³ Indeed, I don't believe that they are entirely distinct from each other, and are better understood as conventionally distinguished sub-types of mental agency. Yair Levy's

² Watzl (2011a: 846) suggests that mind-wandering is a case of involuntary endogenous attention, but offers no argument in defense, and the source he cites in support (Smallwood and Schooler 2009) does not actually make this claim. If one's mind wanders while say, reading, we can say that the reader has strayed off-task, but this is hardly involuntary in the way that beliefs or pains are involuntary, and fully aligns with the points I make about motivational control in section 1 of this paper.

³ It is worth noting recent scepticism in psychology and philosophy that the attention is a unified capacity (see Taylor 2015; 2018 for reviews). Psychologists operationally distinguish three types of attention in experimental studies: 'selection' (which covers increased processing of selected information), 'executive attention' (a mechanism for managing conflicts between thoughts, feelings, and responses) and 'alerting' (in which we maintain high sensitivity to incoming stimuli). The distinction between selection and executive attention aligns with the distinction between exogenous and endogenous attention, though I think alerting attention can have both exogenous and endogenous varieties.

recent account of mental agency similarly blurs the distinction. His main claim is that “attention is the most general act” (Levy, 2022). That is, particular mental acts (such as observing or thinking or, most importantly for our purposes, imagining) stand as determinates to the general determinable of being attentive. To make this claim plausible, Levy excludes cases of exogenous attention as merely ‘pre-attentive’, precisely because we lack voluntary control over such cases. So Levy’s claim turns out to be that endogenous attention is the most general act and imagination is a determinate of endogenous attention.

I’m not entirely satisfied with Levy’s taxonomy, and I think philosophers and psychologists will dispute Levy’s characterisation of exogenous attention as merely pre-attentive and not attention proper. However, I will not debate this. Whether imagination is a determinate of endogenous attention, or they are both sub-types of a third thing, both Levy and I should eventually end up with comparable general accounts of mental agency. This is because both of us will only attribute to mental agency features that imagination and endogenous attention have in common and not, say, features that are exclusive to imagining. Meanwhile, my approach has an advantage that in surveying several features shared by imagination and endogenous attention, I can develop a more fleshed out account of the characteristics of mental agency than Levy offers.

The point is that the shared features of these two exemplary cases of mental agency should be a good guide to the nature of mental agency. What we must watch out for are features that imagination and endogenous attention share but which do not apply to mental agency. Another danger is that there are major features of mental agency which my comparison neglects. However, so long as our comparison cases are both paradigmatic of mental agency and also reasonably different from each other, we can be confident to find key components of mental agency by including features that they share and excluding features that they do not share. Naturally, more cases of mental agency would help to refine this process further, but space permits a detailed comparison of only two capacities here.

Finally, I should address any concern that my strategy is circular. Am I developing an analysis of what counts as mental agency on the basis of cases already defined as mental acts? I am not, because my goal is to develop a mechanistic account of what mental agency involves, and not a definition or analysis of what counts as mental agency. Starting with some accepted cases and then identifying their common characteristics is an entirely standard practice for

mechanistic accounts in the philosophy of mind. Let us, then, begin to identify these common features.

1. Motivational Structure

Since our aim is an account of mental agency, it makes sense to begin with the motivational features of our paradigm cases. It seems that imagination and endogenous attention share a common motivational structure. In a nutshell, both processes are driven by currently prioritised goals that can shift between the sub-personal and personal level.

I'll start with imagination. Introspection indicates that we can form person-level goals to imagine which we then satisfy. For instance, I wonder 'what shall I do this weekend?' and bring various possibilities to mind. This behaviour can be treated like any other action in the sense that it has satisfaction conditions which will continue to drive the activity until they are met or interrupted by some other motivation.

Although we can explicitly intend to imagine, it is not necessary for the driving motivation to be articulated in imagery or inner speech. Consider imaginings driven by hunger. Here one's hunger drives the imaginative activity until one comes up with something suitable to eat. Or consider an act of imaginative empathy stimulated by anxiety about someone's erratic behaviour. Meanwhile, there are cases where imaginings arise unbidden, or even in conflict with our reflectively endorsed desires. For instance, one finds oneself imagining that something terrible has happened to one's child despite desiring not to. However, a prioritised motivational driver is still clearly at work here; anxiety about one's child and the aim to anticipate possible harms that could be managed.

Considerations of parsimony, as well as mechanical respectability, imply that there isn't a fundamental difference between the motivational drivers here, so far as their basic causal power to draw forth ideas or imagery is concerned. As such, I believe we should regard goals as mental states that can be either personal or sub-personal.⁴ That is, goals can drive our behaviour (including mental behaviour) without our reflective or person-level awareness, but the very same mental states can also become conscious and reflectively endorsed. For instance, we can catch ourselves imagining food, or possible harms, or what bird corresponds to that burst of

⁴ Cf. the approach to perceptual representation outlined by Burge (2005: 49-50).

birdsong, but we can also be consciously and reflectively involved in the goals driving these behaviours.

We can tie this point to Wittgenstein's well-known declaration that imagination is 'subject to the will' (1980: §80). Here Wittgenstein seems to suggest that while some imaginings arise unbidden, we can nevertheless manipulate them, if we try. Gerrans (2014: 142) similarly argues that, unlike perception and belief, we are still free to imagine the opposite of whatever contents have come to mind. Hopkins (2018) meanwhile argues for the claim that episodic memory is a species of imagining by observing that it is possible to take an episodic memory and alter details at will until it becomes a pure creative imagining.⁵

I propose the following way to sharpen up Wittgenstein's claim: What is meant by 'the will' is just those cases in which the driving motivation is person-level or conscious. All we need add is that a motivation may start out sub-personal and become person-level (if say, it becomes urgent). Or else the sub-personal motivation to form an image may be interrupted by a person-level motivation to inhibit or adjust the imaginative activity.

Can similar claims about motivational structure be justified with respect to endogenous attention? It seems that they can. During any act of the attention it is always possible to alter the content to which one attends, if only temporarily. Similarly, we can discern the goals driving the attentional behaviour as entirely analogous to the imaginative case, e.g. in feeling hungry, or threatened by someone's behaviour. What is less obvious is how the motivational state driving endogenous attention may be sub-personal and yet still count as endogenous (rather than exogenous) attention. Again we can point to cases like hunger in which one may not consciously realise that hunger is making one more attentive to sources of food. The distinction from exogenous attention is still viable because one does not simply notice food because it is striking. One is disposed by one's motivational state to notice even subtle food sources.

In the sub-personal hunger case, we should say that the goal to eat is relatively high among the agent's current priorities, even if it is not yet high enough to impinge on the consciousness of

⁵ Ichikawa (2009: 106-108) also makes use of the Wittgensteinian formulation, and helpfully reviews some confusions about its usage.

the agent. The appeal to currently prioritised goals is important for distinguishing endogenous attention from exogenous attention. This is because the way a sudden noise or pain grabs our attention can be understood within a broader framework in which such stimuli are relevant to our dispositional goals to be aware of what is happening around us or to avoid bodily harm.⁶ To accommodate such cases of exogenous attention, we can specify that content that is already minimally processed triggers the prioritisation of long-standing goals. In other words, our goals to avoid pain or identify certain features of the background environment are relatively dormant until they are ‘woken up’ by highly relevant stimuli. In contrast, in acts of endogenous attention currently prioritised goals seek out stimuli which need not (though can be) already be present to the agent. The same is also true of acts of imagination.

A further similarity in the motivational structure of imagination and endogenous attention is that both can be experienced as effortful (Preston & Wegner 2009).⁷ It can often be difficult to maintain one’s attention towards a target, either due to distractions or boredom. Similarly, it can often take effort to imagine certain complex contents or to hold onto an imagined scene over time. The sense of effort seems partly to be a product of one’s intentions (or satisfaction conditions) failing to match up with the detected results, and partly a product of motivational conflict as our various goals jostle for priority. We want to maintain attention or to mentally construct the scene, but the discomfort or tiredness we experience motivates us to stop. Particularly where the task is not intrinsically rewarding, more enticing goals can interfere.

Overall, imagining and endogenous attending are capacities that are driven by currently prioritised (and sometimes effortful) goal states, which given their high priority are either person-level or continuously susceptible to person-level awareness. This motivational feature is plausibly central to a general account of mental agency because it begins to characterise the nature of our mental control. This is also not a feature of exogenous attention and low-level perceptual processing more generally. At the lower stages of perception, content is drawn into

⁶ It is worth noting that different views of perception emphasise its link with agency. Ecological views regard perception as for agency (e.g. Gibson 1979). Predictive processing views also treat perception under the goal to predict the world accurately (e.g. Hohwy 2013). These views will however accept the intuitive difference between being surprised by a sudden bang, and paying attention to an anticipated bang. Even though the surprising bang will certainly feed into ongoing action, its coming to mind is not because a currently prioritised goal seeks that content. Similarly, even though predictive processes will seek to accommodate the bang, they do not predict it. This is why the surprising bang counts as case where low-level perceptual processing brings content to mind and not a case of basic mental agency. Thanks to an anonymous referee for pressing me on this issue.

⁷ Worth noting here is that Jennings regards effort as a primary marker of voluntary attention (2020: 209).

the cognitive system prior to goals seeking out that content. Whether that content is then taken up depends on its relevance to the agent's goals. There is a difference between motivational drives actively seeking out content and motivational drives being awoken by content.

2. Consciousness and working memory

Following up on the observation that the goals driving imagination and endogenous attention are continuously susceptible to person-level awareness is the observation that the *products* of these capacities are typically conscious. Most likely this is due to the person-level at which imaginative and attentional goals tend to operate. However, it may be possible for the products of these capacities to be unconscious.

For instance, with regards to imagination, Church (2008) describes a case in which a mother sneaks into her son's room while he is away. Her behaviour indicates that she is imagining that someone else is in the house, though she is not consciously aware of doing so (and Church argues that her behaviour would be rather different if she unconsciously *believed* there to be someone else in the house).⁸ Meanwhile Mole (2020) argues that attention can be unconscious, e.g. there are various experiments in which the selective processing of visually unconscious details can be discerned, evidenced by the facilitated performance of the individual on tasks that follow.

Now it may be argued that attention bears a closer connection to consciousness than imagination. It is hard to judge to what extent we can have unconscious imaginings.⁹ However, it does not seem that the purpose of imagining is so directly tied to becoming conscious of something. That is, attention seems necessary for consciousness in a way that imagination is not. The connection is so close that Jesse Prinz (2012) has developed a theory of consciousness based upon this connection. He claims that consciousness is essentially a process whereby attention makes contents available to the working memory. We need not debate the soundness of Prinz's model to recognize that his model would not be at all plausible were attention not so closely allied to consciousness.¹⁰

⁸ Gaut (2003) also argues that imagining can be unconscious, while Nanay (2021) argues that *mental imagery* can be unconscious, though it is not clear if this also applies to imagination.

⁹ E.g. Kind (2021) is unconvinced by Church's case, arguing that the mother's behaviour could alternatively be explained by beliefs such as that "she should be stealthy just to be safe" (2021: 6).

¹⁰ Note that the evidence that attention is not necessary for consciousness, based on individuals' capacity to report the 'gist' of peripheral contents while concentrating on something else, is not very compelling. As De Brigard &

Related to this point, we can observe that it is a central or paradigmatic function of pains and pleasures to stimulate attention. In contrast, it is not central or paradigmatic for these affective states to stimulate imaginings. Surprise similarly has a direct connection with the arousal of attention that it does not have with imagination. Of course, we frequently do engage in imaginings as a result of affective states. The point is just that this connection is not so tight. The close connections between attention and affect seems to relate to the widely recognized role that attention has in prioritising content that is urgent or particularly significant for our well-being (cf. Mole 2011, Watzl 2011b, Wu 2014).

Yet, when we look closer, we see that these distinctive features of attention can be attributed to the *exogenous* form of attention. Pain, pleasure, and surprise do not distinctly occupy our attention because we are paying attention to whether or not they will occur. On the contrary, these affective states intrude or force us to attend to them; they stimulate attention exogenously. Similarly, maybe some or other form of attention is necessary for consciousness, but it need not be endogenous attention. So if we focus exclusively on endogenous attention, it appears that the way it tends to bring contents to consciousness is similar to the way that imaginative projects tend to bring contents to consciousness.

Given the current status of the evidence (and the difficulties surrounding consciousness) claims about the connections between imagination, endogenous attention, and consciousness must, I think, be tentative. However, a more definite claim is that both endogenous attention and imagination deliver their outputs to the working memory. Working memory is understood as the faculty in which we hold and manipulate information in service of ongoing tasks (Baddeley 1986). This is precisely what both imagination and endogenous attention achieve, and indeed, it is typical for models of working memory to reference capacities that are explicitly linked to both imagining and endogenous attention such as planning, problem solving, and the ‘visuospatial sketchpad’. More recently, neuroscientists Miller and Cohen have influentially characterised working memory as serving to maintain “the goals and rules of a task” (2001:

Prinz (2010: 57-58) have argued, attention could merely be divided in such cases. The evidence that attention is not sufficient for consciousness is stronger, though it seems to be a kind of attention involving merely the increased selective processing of information, rather than alerting attention or executive attention which are driven by motivational salience. I discuss this in some detail, with particular reference to the influence of our motivational systems in Cochrane (forthcoming).

186). Given the goal-driven nature of imagination and endogenous attention, and their guidance by norms (see section 4 below), there is a good fit between these capacities and Miller and Cohen's account of working memory.

Currently there is suggestive, though not decisive evidence that the contents of working memory can be unconscious (Velichkovsky 2017; Persuh et al. 2018). Thus I believe our best stance for now is to correlate the answer of whether the products of imagination and endogenous attention are conscious with the ongoing empirical investigations into the consciousness of working memory.

While this somewhat modest conclusion may not be considered especially illuminating of the nature of mental agency, it does however serve to mark out mental agency from cognitive activities that do not deliver their products to working memory. For instance, there is good evidence that inference processes can sometimes be unconscious and fail to deliver their outputs to working memory (see e.g. Johnson-Laird 2006: Ch.5; Ludwig & Munroe 2019 for reviews). In many cases we are not aware of having performed these inferences until perhaps we draw on their results later on. It seems clear that automatic, non-agential, processes are at work here. Note that I'm not claiming that inference is never a case of mental agency, only that there are definite instances that are not cases of mental agency. Thus reasoning should not be regarded as a central case of mental agency. Instead, it is more likely to be a function that we deploy as part of agentic processes (more on this in section 4 below).

3. Concepts

To further fill out the nature of mental agency, it is important to observe that both the imagination and endogenous attention are domain-general abilities. That is, there is no subject matter or practical context from which these capacities are excluded. We might suppose that endogenous attention is 'outwardly' focused where the imagination is more 'inwardly' focused. However, there are cases of outwardly focused imagination, such as seeing-as (or in the case of artworks, what Wollheim (1980) called seeing-in). There are also cases of inwardly focused endogenous attention, as when we pay attention to the thoughts flowing through our minds. Thus both imagination and endogenous attention are extraordinarily flexible. This, alongside the person-level of motivation, seems to be the major factor behind the intuition that these capacities are central rather than peripheral to our mental economies. Yet despite this flexibility,

it seems that we cannot imagine, or endogenously attend to something for which we lack a concept.

By ‘concept’, I mean a sub-propositional mental structure which specifies correctness conditions for a type of content. This characterisation is deliberately neutral between classical theories of concepts, according to which concepts encode necessary and sufficient conditions for their application (Laurence & Margolis (1999: 9), and prototype theories, according to which concepts encode statistical analyses of the properties their members tend to have (Laurence & Margolis 1999: 27). Meanwhile, a standard definition of non-conceptual content is “where a mental state represents the world while lacking the concepts required to specify its content” (Bermúdez & Arnon 2020). Many philosophers believe that perceptual content is at least partially non-conceptual because it is more fine-grained than our conceptual repertoire (see Toribio 2007; Wright 2015 for discussion). That is, it is possible to perceptually discriminate something while lacking the concepts that specify the kind of thing one is perceiving.

The reliance on concepts serves to further distinguish endogenous and exogenous attention. With exogenous attention, I can certainly be struck by the appearance of some new thing and gaze in wonder as I try to make sense of it.¹¹ But this is not true of endogenous attention. I cannot start to pay attention to anything without at least some correctness conditions what I’m seeking. For instance, if I had no concept of rabbits or ducks, I would not be able to attend to either the rabbit aspect or the duck aspect of the duck-rabbit picture. Lacking these concepts would however not prevent me from being struck by the specific shapes and colours that make up this picture.

Yet note that there can be vagueness regarding the exact point at which exogenous attention becomes endogenous attention. As I maintain attention towards something over time, the active quality of my behaviour begins to imply endogenous attention. For instance, suppose I find myself in a weird mood for which I lack a concept. Before long I may start to form an understanding of the mood. I may note to myself that it’s not merely tiredness or anxiety, and begin to identify some of its distinctive felt characteristics. Accordingly, I suggest that in cases

¹¹ Consider for instance Tim Crane’s (1988) discussion of the waterfall illusion, the contradictory appearance of which Crane argues must outstrip our conceptual resources.

of attention-maintenance, something like a demonstrative concept THIS STRANGE MOOD may suffice for endogenous attention. However given that concepts encode correctness conditions, it is a minimal requirement that one start to have some idea about what counts as this strange mood; to be able to note where the mood begins or ends, or distinguish it from other moods, for instance. Without such correctness conditions, there is no way to determine if one's attentive goal is satisfied or not.¹² Then in cases where the target is *not* currently present, a more substantive concept is demanded.

Imagination has a similar connection to concepts. Obviously cases of inner speech and propositional imagining are structured by concepts. But I also claim that we cannot imagine sensory qualities for which we lack a concept. Consider the auditory image of Barack Obama's voice that may come to mind when we are presented with a mute video of him speaking (Nanay 2018). Perhaps we lack the terminology to describe this particular vocal quality, but that doesn't mean we lack a concept of 'Barack Obama's voice' or have the kind of general grasp of this quality that would allow us to correctly identify it across more or less similar sensory instances and distinguish it from close neighbours.¹³

Margherita Arcangeli (2021) has recently argued that all imaginings, including mental imagery, recruit a concept "for each constituent of its content" (2021: 3198).¹⁴ It might be thought that mental imagery is non-conceptual in so far as it is a counterpart to the non-conceptual nature of perception (e.g. McGinn 2004: 36-39). However Arcangeli observes that we need to possess some understanding of how things of different kinds look, in order to be able to visually imagine things of that kind (whereas this not required for seeing those things). For instance, I cannot visually imagine a birch tree (as a birch tree) unless I have some understanding of what birches are supposed to look like. Furthermore, imagining is very plausibly less fine-grained than perceptual experiences. Though we may be able to visually imagine quite subtle shades

¹² The same argument would apply to a meditator endogenously attending to their breathing. It is an interesting issue to what extent meditation can bring about conceptual attenuation, and I expect we shall see more research into this in the future (Coseru 2022 has some initial discussion). But note that the silencing of inner speech need not imply the loss of concepts, especially since concepts need not be linguistic in format (see Williams 2021 for an extended discussion). Thanks to a referee for raising this case.

¹³ There are perhaps 'inner' visual experiences for which we lack concepts. For instance, it seems theoretically possible to experience a colour after-image for which we lack a concept. However as noted in the introduction, we should not confuse imagination with mental imagery. Just like percepts and hallucinations, after-images are 'bottom-up' triggered sensations (cf. Kind 2001; Currie & Ravenscroft 2002; McGinn 2004). In contrast, imagination is a 'top-down' faculty to the extent that it relies on concepts.

¹⁴ Arcangeli allows that mental imagery need not be conceptual at the attitudinal level such that it can play an inferential role in the same manner as a belief.

of red, concepts of these different shades would be required to be capable of specifying them in our imaginations.¹⁵ Meanwhile cases of imaginative seeing (such as seeing-as) may involve non-conceptual content, but only in so far as they blend in independent perceptual elements (2021: 3193-3195).

Supporting this point is the widely accepted stipulative nature of the imagination (McGinn 2004: 31; Kung 2010: 626; Langland-Hassan 2016: 61-62; Balcerak Jackson 2018: 212-214; Kind 2019: 167). Wittgenstein remarked that when he imagines King's College on fire, it hardly matters whether the image he forms actually resembles King's College (1958: 39). In any case, that image stands for King's College (or at least what he personally means by King's College). A plausible explanation for the irrelevance of resemblance here is that the imagining of King's College is at least partly constituted by the concept KING'S COLLEGE. If it did not, the image would not automatically refer to King's College.

Note that my claim is not that an explicit intention to imagine must exactly specify the content which we bring to mind prior to imagining it. Various details, such as the orientation of some imagined object, are liable to be specified only in the act of imagining. Nevertheless, that orientation will be one over which we have a conceptual grasp, again in the sense that we could correctly identify it across more or less similar instances- a prototype concept perhaps, rather than a linguistic tag.

But can't we have mental images that are obscure to us; where we puzzle over what we are imagining? I suppose there can be cases where we are uncertain what a certain pattern of shapes and colours resembles. But this does not mean that we lack a concept of the pattern as a certain kind of pattern. All this means is that we have yet to apply an additional concept to the overall gestalt that would pick out an additional object, as in the duck-rabbit case I mentioned above. Similarly, if I imagine the face of an actor in a film I know, and then come to realise that the same actor played a role in a different film, I am now linking up two concepts I previously possessed, and forming an additional, more general, concept of that actor.

¹⁵ All of this is compatible with the Humean claim that imagery originates in perception, because we can say that perceptual experience allows us to acquire concepts (Arcangeli 2021: 3193). Indeed, Arcangeli distinguishes between sensory imagining and propositional imagining by arguing that the former, but not the latter, exclusively relies on observational concepts; concepts that "can be directly applied on the basis of perceptual experience alone" (2021: 3201). That is, to form an observational concept, we need to have perceptually experienced its constituent contents.

This example illustrates the further point that imagination can play a role in the formation of new concepts and learning (more on which below). Endogenous attention equally plays a role in new concepts. Indeed, Smithies (2011) argues that attention is necessary for demonstrative thought, and demonstrative thought is required to form new justified beliefs about the target (while this justification is missing in cases of blindsight, where conscious attention is absent). Smithies' argument does not mark out endogenous attention in particular, but since endogenous attention will bring about demonstrative thoughts as we search for and then maintain attention towards things, endogenous attention thereby allows us to form new concepts.

Overall, it is a significant feature of both imagination and endogenous attention that they operate by means of concepts. Like the motivational features surveyed in section 1, this conceptual feature helps to pick out mental agency as a person-level capacity, because the person-level tends to be identified with high-level features of cognition (i.e. abstract thought) in contrast to low-level (aspects of) sensory processes. However, there are also non-voluntary capacities (such as the mere possession of a belief and automatic inference processes) that operate with concepts. Thus the conceptual feature cannot sufficiently distinguish mental agency from non-agentive capacities. Instead, this feature helps to set the limits on mental agency—we go this far and no further.

4. Norms and Habits

The next common feature I identify is again concerned with setting limits on mental agency. Although both imagination and endogenous attention are characterised by their spontaneity, the precise manner in which an act of imagination or endogenous attention proceeds is structured by norms and/or habits outside of our direct control. This condition is disjunctive. Sometimes we are guided by explicit or implicit mental representations of norms. At least as commonly however, and particularly for repeated activities, our mental activity is scaffolded by our neural, bodily and environmental structures, some of which have been intentionally designed in accordance with a norm for what mental events are legitimate, some of which have accumulated entirely unintentionally (i.e. mere habits) and some of which are the by-product of intentional behaviour.¹⁶

¹⁶ I use the term 'scaffolded' in reference to Sterelny's (2010) notion. More generally, my appeal to norms is intended to embrace the possibilities of embodied, extended, embedded and enactive (i.e. 4e) cognition.

It is here that a significant contemporary issue regarding the imagination is raised. As Amy Kind has argued (2013) it is hard to reconcile the imagination as the paradigmatic creative capacity with the use of imagination to tell us interesting truths about the world. Kind and Kung (2016) call this ‘the puzzle of imaginative use’. Kind’s solution to this puzzle (2016b) is to make the imagination something that can be subject to different psychological constraints, depending on the task at hand.

Thus in projects where the imagination aims to track the truth, there are likely to be relatively strict representations of norms in play determining relevance and permissible format. For instance, in imagining another person’s mental state, I rely on norms about how minds work, e.g. that an expression of rage is typically accompanied by an impulse to strike or destroy. Similarly, in the case of scientific thought experiments (like Einstein’s tram), I am constrained by what I know about light. Meanwhile, in more creative imaginative projects, the guiding factors may be relatively loose, but they are still present. Even when I challenge myself in a very open-ended way to come up with something as weird and wonderful as possible, I can in retrospect discern habits and norms at work. For instance, there may be an implicitly represented norm to select contents that are as unrelated as possible; to actively violate categories, or to make broad analogies. I may have further acquired tropes for what counts as weird and wonderful from science fiction and fantasy. I moreover find that the more I repeat this sort of challenge, the more that habits start to form.

Overall, different sub-types of imagining (and endogenous attention) are distinguished in large part by the different norms and habits that they operate under. Much of the important detail and richness of our specific imaginative capacities will be revealed by investigating how exactly these norms and habits work. I can only offer a few generalisations here.

So far as mental agency is concerned, the important point here is that although we can be immediately aware of the motivational drive triggering the imaginative activity, and although the referential content of what we imagine is stipulated, the norms or habits that structure our thoughts need not be under our control at all. Our conceptual repertoire can develop structure entirely independently of our will. We may have no control over the way our minds associate some ideas more closely than others, or the way that an idea may inhibit its direct contradiction. Besides this, we acquire from culture various standards for properly formed thoughts (e.g. in

mathematical or linguistic thought). Our lack of direct control over these factors potentially allows the imagination to be instructive.

However, it should be borne in mind that because norms are learned, we can choose to learn them by means of deliberate practice (this is also the main thrust of Wu's 2013 response to Strawson's 2003 sceptical attack on mental agency). It can also be up to us whether a certain norm is brought into play in a given imaginative project, since the relevant norms seem to be dictated by one's goals and these goals can be person-level. Finally, even if strict norms are in play within certain conceptual domains, the use of these domains can be embedded within wider imaginative projects that treat outputs more freely.

Similar points can be made with respect to endogenous attention. Obviously when we pay attention to the environment, the structure of that environment will significantly constrain the particular things we notice (and in what order). But more than this, acts of endogenous attention are also structured by habits and internally represented norms. For instance, if one is searching a room, paying attention to where one's book may be located, there are definite patterns that our search routines habitually display or which one feels it ought to display, such as proceeding systematically across one's shelves (cf. Wu 2014: 35-36). Consider also more cognitively demanding cases of endogenous attention. Suppose that I am paying attention to a piece of music in order to memorise it. My act of attention will be structured by implicitly represented norms concerning the features of the music that are most relevant to its construction such as how closely it fits a certain metrical or tonal structure.

In general, paying attention is permeated by expectations regarding the way the environment is likely to unfold, and these expectations are generated by learned habits and norms. These expectations are moreover specific to the particular domain—be it a room's layout, a piece of music, or the arrangement of pieces on a chess board. Just as in acts of the imagination, we can deliberately learn the various norms and habits generating these expectations and thus maintain some indirect control over them, even if in the moment, they are not directly controlled.

Related to the role of norms in imagination and endogenous attention is a similar way in which these activities have a degree of focus. At one end of the scale there is a very specific kind of content that we are focused upon generating. At the other end of the scale we allow a broad range of contents to strike us. We can deliberately defocus our attention or allow it to wander

(Irving 2016) just as we can defocus our imaginative activities and allow our imaginations to wander.

The degree of focus is partly determined by the norms in play. Some norms permit a greater range of satisfiers than others. In addition, more than one norm may operate simultaneously, resulting in a restricted range of suitable items. For instance, when trying to come up with a good joke I may require acceptable items to be original and amusing and morally appropriate. When I'm paying attention to an article, searching for a suitable quotation, I may require items to be both informative and succinct.

Besides norms, the other factor that determines focus is the driving goal or motivation. If there is a strong or urgent motivation in play, drawing upon the full person-level resources of the individual, then the mental activity will accordingly be tightly focused on the delivery of a certain kind of content. If, on the other hand, there are no strong or dominating motivational factors, then a plurality of motivations can operate simultaneously to encourage a broad range of content to come to mind.

Given the role of background motivations, the degree of focus is subject to the will in the sense articulated earlier. We can choose to concentrate fiercely upon a mental task, or to allow our minds to wander. Mind wandering is particularly associated with creative thought. Picciuto and Carruthers (2016) connect this with the defocusing on attention, and a corresponding reduction in 'latent inhibition', that allows unobvious ideas to 'bubble up' to consciousness.

Overall, the role that norms and habits play in our capacities for imagining and endogenous attending help to further fill out the way that mental agency is not an entirely spontaneous matter. Our control over our minds is significant, but it is subject to important constraints. These constraints moreover underpin the usefulness of our mental activities in helping us to respond intelligently to our circumstances. The fact that norms (and the associated degree of focus) are relative to the task we aim to complete does however retain a sense in which the person is ultimately able to direct the content of their minds. So while the bare fact that imagination and endogenous attention are structured by norms and habits is not distinctive to mental agentive projects, the influence of our variable goals over which *specific* norms or habits are imposed allows us to see how the mental agentive process can be voluntarily modulated.

5. Activation not inhibition

The final similarity between endogenous attention and imagination that I identify is that both directly bring content to mind and do not directly inhibit or banish content.

In a typical case of imagining we may form the goal to think of something blue, and instantly possible candidates come to mind: the sky, blue ink, blueberries and so on. Suppose we want an unusual blue thing. We ignore these initial ideas and pursue the task a little longer. Before long, more unusual cases come up, say... a blue whale! Now we latch onto this idea; we actively maintain its presence in our working memories. All this happens very fast and very fluently. But it must be noted that we have lifelong practice in initiating searches, passing by unsuitable contents, concentrating on items that bring us closer to our aims, and reiterating the search until satisfactory content is reached.

Now it may be observed that there is a difference between searching content and grasping content in one's working memory. Indeed, the mechanisms involved in searching through one's vast repertoire of concepts are not discernible to introspection. I want to allow that there are two distinguishable processes going on here, only one of which is directly controllable. In section 4, I described how the transitioning or flow of thoughts is guided by various norms and habits. We do not directly control the functioning of these constraints. Nevertheless the selection of mental content is up to us. If I might use a metaphor; it is as if there is a conveyor belt or touch screen menu before us upon which potential items rush past. The manner in which items scroll by, or the order they come in, is not up to us, but we can start the conveyor going, and we reach and grasp for suitable items. That power of starting the search, of reaching and grasping for items is what I mean by bringing content to mind.

A fundamentally similar process seems to be involved in endogenous attention. When one searches a room for a mislaid book, one is paying attention to the room as one moves around it, making oneself receptive to signs of the book's presence. This is structurally similar to cases of difficult mental search, such as when one struggles to recall a famous actor's name. In both cases, one is guided by a goal that makes one receptive to a particular kind of content. Hints of the desired content (e.g. phonemes that fit the actor's name, colours that match that of the book) then become held or prioritised in one's consciousness when they appear, enabling more focused search that may bring to light the desired content.

So both imagination and endogenous attention involve actively bringing content to mind. But some may argue that these capacities also involve inhibiting content. In particular, simulation theories of imaginative empathy describe a process of ‘mental quarantining’. The idea is that when we use our own psychological mechanisms to simulate the mental state of another person, this must be done ‘off-line’ so as to avoid the overrule of one’s pretend beliefs and desires by one’s actual beliefs and desires, or the contamination of one’s actual set of beliefs and desires by the pretend set (Currie & Ravenscroft 2002). This is potentially a kind of inhibition performed by the imagination.

In response, I argue that mental quarantining can be reduced to the active search for content, at least with respect to the part we directly control. We lack the direct ability to turn off mental content, just as we cannot directly stop ourselves from thinking about a white bear once we have been told to do so (Wenzlaff and Wegner 2000). What we have to do instead, just as in the white bear example, is to think positively about something else, and trust that this process of itself will exclude other content.¹⁷ So there is admittedly an extra element here. It seems to be a background feature of our cognitive economy that the activation of content of one kind inhibits the activation of directly conflicting content (cf. Nichols 2006; Weinberg & Meskin 2006: 188). But this conflict inhibition is not something we directly control.

In support of my claim, I note that appealing to the deliberate activation of content rather than deliberate inhibition of content helps to resolve a puzzle facing simulation theories of imagination. The puzzle is that the quarantines cannot be a simple matter of adopting some pretend beliefs that conflict with what we really believe, and then sticking only with those pretend beliefs for the duration of the imaginative project. This is because all kinds of supporting background beliefs are required to sustain simulative projects (cf. Stock 2011). For instance, when one engages with a fiction and imagines the point of view of the character being described, one makes use of background assumptions that the character is human and has ordinary human needs. Such background assumptions are implicit and open-ended, and so it is unclear how the relevant background contents get incorporated within the quarantine zone. Yet if quarantining actually involves only the activation of the pretend states, then background

¹⁷ The same should be true of creative thinking in which we ignore obvious ideas to get at more creative ones (Beatty & Silvia 2012).

states (so long as they do not directly contradict) need not be walled-off at all and are available to contribute to the simulative project.

Similar considerations apply to putative counter-examples regarding endogenous attention. In their analysis of the mental agency involved in meditative practices, Upton and Brent identify a capacity for ‘mental silencing’, in which the meditating agent prevents “the litany of automatically generated thoughts” (2019: 59). On the face of it, this looks like the inhibition of content. However, Upton and Brent claim that mental silencing is achieved by means of shifting one’s attention towards one’s breath “to the preclusion of the automatically generated thoughts.” They are explicit that mental silencing is not directly achieved (2019: 68 ft.28) but is the indirect result of focusing one’s attention upon a certain sensation, and shepherding one’s attention back to that sensation when one’s mind begins to wander. Thus the direct act of mental agency here is still one of bringing a certain content to mind.

Overall, our attempts to screen off distractors in the case of attention are comparable to mental quarantining in cases of mental simulation. We need to focus on what is relevant and ignore or neglect contents that are irrelevant or conflicting. We can bodily remove ourselves from the distractors. We can close our eyes or stop up our ears. But we cannot simply stop paying attention to something without deliberately turning our attention to something else instead.

One last point on this feature; since there may be effort involved in both imagining and endogenous attention, we might interpret the ‘giving up’ of a task as a case in which we inhibit or banish imaginative or attentive content. That is, the experience feels like we have turned off our intense focus. But all that has happened here is that we have replaced a focused mental task with one or more less focused tasks with broader satisfaction conditions. Fundamentally we don’t shift from attention to the mere absence of attention. We replace one kind of attention with another. Consider also the process of falling asleep. This process causes us to stop attending. But we do not actively fall asleep. It is not a process we directly control at the person-level. Rather some sub-personal mechanism quietly shuts off our attentions.

6. The call for content

In summary, I have identified five key features shared by acts of imagining and endogenous attention. Given that imagination and endogenous attention are paradigmatic forms of mental

agency I propose to use their shared features to specify the mechanism of mental agency. Thus according to my model, mental agency:

- i) Is driven by currently prioritised goals that are either person-level or apt to become person-level.
- ii) Delivers its outputs to the working memory.
- iii) Ranges across all and only conceptual contents.
- iv) Proceeds under the guidance of norms and/or habits.
- v) Directly activates rather than inhibits content.

The first of these features seems most decisive in distinguishing mental agency from non-agentive faculties such as pain and pleasure, after-images, or beliefs. Yet the other features are vital for clarifying how exactly this motivational power functions, the field in which it operates and its constraints. Indeed, features ii, iii, and v seem to be necessary for all instances of mental agency. I cannot identify any clear counter-examples. Meanwhile feature iv is only unnecessary in the sense that one could imagine or attend to a single idea which immediately satisfies one's goal, and thus brings to a close one's mental act. The guidance by norms and habits will certainly operate within extended acts of imagining or attending.

Given that the five features all help to flesh out the nature of mental agency, a further question is whether they can be combined into a definite and unified mechanism of mental agency. It looks like this is indeed possible. We can characterise mental agency as, essentially, the power to call conceptual content to mind and to hold it within our working memories. Again, consider the 'menu-scrolling' analogy that I described in section 5. This characterises the basic nature of the activity. We may then say that imagining and endogenous attention are sub-types of this general and person-level power that vary (roughly) with regards to the location of the items we search and the specific guiding norms or habits in play.

Are there other sub-types of mental agency that we can incorporate into this model? Given the broad conception of imagination that I've been working with in this paper, I believe that the vast majority of our mental acts are already covered. It already includes supposition and inner speech, for instance. However, another possible case is decision-making. Not all decisions are voluntary or deliberate, but in those that are, we can observe a process in which the options to be considered are deliberately brought before the working memory and their outcomes are

projected. Norms and/or habits structure this process both with respect to the options that come to mind and the determination of outcomes, which is at least sometimes an imaginative simulation task. Overall, there seems to be a basic fit with the model, though of course, the specifics of decision-making require a far larger analysis.¹⁸

Do the features I identify justify the claim that mental agency is a distinctive and unified capacity in its own right? I suppose the ‘call for content’ to be a definite psychological mechanism; it has its typical inputs and outputs and functions linking the two. However I must emphasise that if it is a definite mechanism, it must be thoroughly distributed rather than modular. That is, to call for content must be a power possessed by a wide range of motivational drives that can become person-level or conscious, and which can draw upon norms or habits in the transitions between content. The call for content is not to be conceived as a central executive, except perhaps in a ‘virtual’ sense that the interaction and coordination between multiple drives emerges as a coherent domain-general power, and that the working memory in which it operates is experientially central to a person’s mental life.

It’s worth noting that this account of mental agency converges somewhat with Strawson’s (2003) sceptical position. The overall thrust of Strawson’s argument is that cognition is largely automatic or ballistic in nature. However, he regularly says that we can set our minds to a problem, or focus our concentration upon certain ideas. Thus Strawson’s descriptions allow for at least one person-level mental power, and one that moreover resembles the call for mental content. Yet I do not dismiss the power I have described as merely ‘catalytic’ (as Strawson puts it). In particular, Strawson is wrong to say that the imagination operates independently of person-level control (2003: 239-242). Our person-level goals are able to directly call for content. These motivations link up with various norms. The end result is the specification of desired content, often at a fine-grained level.

Overall, I believe the comparison between imagination and endogenous attention allows us to formulate a substantive general account of the mechanism of basic mental agency. This should hopefully contribute to wider discussions about the nature of agency, particularly in identifying

¹⁸ I take the final selection of an option to be an affective mechanism in which our goals are set depending on the strength of our affective responses to projected outcomes. These goals are either person level or susceptible to person-level awareness in the manner I discuss in section 1 and will themselves drive various mental agentive projects. I discuss this sort of mechanism and its controversies in Cochrane (2018: section 6.4).

a mental act that contributes to indirect or non-basic acts. It should also contribute to various philosophical and psychological issues where mental agency is relevant such as our capacities to regulate well-being and our understanding of disorders where agency is compromised.¹⁹

References

Arcangeli, M. (2021). The conceptual nature of imaginative content. *Synthese*, 199(1-2), 3189-3205.

Baddeley, A. D. 1986. *Working Memory*. Oxford: Oxford University Press.

Balcerak Jackson, M. 2016. On the epistemic value of imagining, supposing, and conceiving. In A. Kind & P. Kung (eds.) *Knowledge through Imagination*. Oxford: Oxford University Press, 41-60.

Balcerak Jackson, M. 2018. Justification by imagination. In F. Macpherson and F. Dorsch (eds.), *Perceptual Imagination and Perceptual Memory*. Oxford: Oxford University Press, 209-26.

Beaty, R. E. & Silvia, P. J. (2012). Why do ideas get more creative across time? An executive interpretation of the serial order effect in divergent thinking tasks. *Psychology of Aesthetics, Creativity, and the Arts*, 6(4), 309-319.

Bermúdez, J. & Arnon C. (2020). Nonconceptual Mental Content. *The Stanford Encyclopedia of Philosophy* (Summer 2020 Edition), Edward N. Zalta (ed.) <<https://plato.stanford.edu/archives/sum2020/entries/content-nonconceptual/>>

Burge, T. 2005. Disjunctivism and perceptual psychology. *Philosophical Topics*, 33(1), 1-78.

¹⁹ This paper has had a rocky path to publication and I have received a considerable amount of feedback over the last 5 years. Thanks to Ian Ravenscroft, Magdalena Balcerak Jackson, Eric Peterson, and Shen-yi Liao for providing detailed comments on drafts. Thanks also to workshop participants at Yang-Ming University Taiwan, Adelaide University, and the online ‘Second Annual Covid gathering’ organized by Amy Kind. I am also grateful to the various anonymous referees who reviewed my paper.

Church, J. 2008. The hidden image: A defense of unconscious imagining and its importance. *American Imago*, 65(3), 379-404.

Cochrane, T. (2018). *The Emotional Mind: A control theory of affective states*. Cambridge: Cambridge University Press.

Cochrane, T. (forthcoming). Consciousness, attention, and the motivation-affect system. *Journal of Consciousness Studies*.

Coseru, C. (2022). Consciousness, content, and cognitive attenuation: A neurophenomenological perspective. In R. Repetti (ed.), *Routledge Handbook on the Philosophy of Meditation*. London: Routledge, 354-368.

Crane, T. (1988). The waterfall illusion. *Analysis*, 48: 142–147.

Currie, G. and Ravenscroft, I. (2002). *Recreative Minds. Imagination in Philosophy and Psychology*. Clarendon Press: Oxford.

De Brigard, F. & Prinz, J. (2010) Attention and consciousness. *Wiley interdisciplinary reviews: Cognitive science* 1 (1), 51-59.

Gaut, B. 2003. *Creativity and imagination*. In B. Gaut and P. Livingston (eds.) *The Creation of Art*. Cambridge: Cambridge University Press, 148-173.

Gerrans, P. 2014. *The measure of madness: Philosophy of mind, cognitive neuroscience, and delusional thought*. Cambridge, MA: MIT Press.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.

Hopkins, R. 2018. Imagining the Past: on the nature of episodic memory. In F. MacPherson & F. Dorsch (eds.), *Memory and Imagination*. Oxford: Oxford University Press.

- Ichikawa, J. 2009. Dreaming and imagination. *Mind & Language*, 24(1), 103-121.
- Irving, Z. C. (2016). Mind-wandering is unguided attention: accounting for the “purposeful” wanderer. *Philosophical Studies*, 173(2), 547-571.
- James, W. A. 1890/1950. *The Principles of Psychology*. New York: Dover.
- Jennings, C. D. (2020). *The Attending Mind*. Cambridge: Cambridge University Press.
- Johnson-Laird, P. (2008). *How we reason*. Oxford: Oxford University Press.
- Kentridge, R. W. 2011. Attention without awareness. In C. Mole, D. Smithies & W. Wu (eds.), *Attention: Philosophical and Psychological Essays*. Oxford: Oxford University Press, 228-246.
- Kind, A., & Kung, P. (Eds.). 2016. *Knowledge through imagination*. Oxford: Oxford University Press.
- Kind, A. 2001. Putting the image back in imagination. *Philosophy and Phenomenological Research*, 62(1), 85-109.
- Kind, A. 2013. The heterogeneity of the imagination. *Erkenntnis*, 78(1), 141-159.
- Kind, A. 2016. Imagining under constraints. In A. Kind & P. Kung (eds.) *Knowledge through imagination*. Oxford: Oxford University Press, 145-59.
- Kind, A. 2019. Mary’s powers of imagination. In S. Coleman (ed.), *The Knowledge Argument*. Cambridge: Cambridge University Press, 161-179.
- Kind, A. 2021. Can imagination be unconscious? *Synthese* topical collection ‘Imagination and its Limits’, 1-21. <https://doi.org/10.1007/s11229-021-03369-0>
- Kung, P. 2010. Imagining as a guide to possibility. *Philosophy and Phenomenological Research*, 81(3), 620-663.

Langland-Hassan, P. 2016. On choosing what to imagine. In A. Kind & P. Kung (eds.) *Knowledge through imagination*. Oxford: Oxford University Press, 61–84.

Langland-Hassan, P. 2020. *Explaining Imagination*. Oxford: Oxford University Press.

Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In *Concepts: Core readings*. Cambridge, Massachusetts, MIT Press, 3-81.

Levy, Y. (2019). What is ‘mental action’? *Philosophical Psychology*, 32(6), 969-991.

Levy, Y. (2022). The most general mental act. In M. Brent & L. M. Titus (eds.), *Mental Action and the Conscious Mind*. New York: Routledge, 79-99.

Ludwig, K. & Munroe, W. (2019). Unconscious Inference Theories of Cognitive Achievement. In T. Chan & A. Nes (eds.) *Inference and consciousness*. London: Routledge, 15-39.

McGinn, C. 2004. *Mindsight*. Cambridge, MA: Harvard University Press.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual review of neuroscience*, 24(1), 167-202.

Mole, C. 2011. *Attention is cognitive unison: An essay in philosophical psychology*. Oxford: Oxford University Press.

Mole, C. (2020). Consciousness and Attention. In U. Kriegel (ed.) *The Oxford Handbook of the Philosophy of Consciousness*, Oxford: Oxford University Press, 498–519.

Nichols, S. 2006. Imaginative blocks and impossibility: An essay in modal psychology. In S. Nichols (ed.) *The architecture of the imagination*. Oxford: Oxford University Press, 237-255.

Nanay, B. 2018. Multimodal mental imagery. *Cortex*, 105, 125-134.

Nanay, B. 2021. Unconscious mental imagery. *Philosophical Transactions of the Royal Society B*, 376(1817), 20190689.

Paul, S. K. 2021. *Philosophy of Action: A Contemporary Introduction*. London: Routledge.

Peacocke, C. 2007. Mental action and self-awareness. In B. P. McLaughlin & J. Cohen (eds.) *Contemporary debates in philosophy of mind*. London: John Wiley & Sons, 358-376.

Persuh, M., LaRock, E., & Berger, J. 2018. Working memory and consciousness: the current state of play. *Frontiers in human neuroscience*, 12(78), 1-11.

Picciuto, E., & Carruthers, P. 2016. Imagination and pretense. In A. Kind (ed.) *The Routledge Handbook of Philosophy of Imagination*. New York: Routledge, 334-345.

Preston, J. & Wegner 2009. Elbow grease: The experience of effort in action. In E. Morsella, J. A. Bargh & P. M. Gollwitzer (eds.), *Oxford Handbook of Human Action*. Oxford: Oxford University Press, 569-586.

Prinz, J. 2012. *The Conscious Brain: How attention engenders experience*. Oxford: Oxford University Press.

Sterelny, Kim (2010). Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences* 9(4), 465-481.

Schlosser, M. 2019. Agency. In *The Stanford Encyclopedia of Philosophy*. E. N. Zalta (ed.) <<https://plato.stanford.edu/archives/win2019/entries/agency/>>

Smallwood, J. & Schooler, J. W. 2009. Mind-Wandering. In *The Oxford Companion to Consciousness*, T. Bayne, A. Cleermans & P. Wilken (eds.) Oxford: Oxford University Press, 143-5.

Smithies, D. 2011. What Is the Role of Consciousness in Demonstrative Thought? *Journal of Philosophy*, 108(1), 5-34.

Stock, K. 2011. Unpacking the boxes: the cognitive theory of imagination and aesthetics. In E. Schellekens & P. Goldie (eds.) *The Aesthetic Mind: Philosophy and Psychology*. Oxford: Oxford University Press, 268-282.

Strawson, G. 2003. Mental ballistics or the involuntariness of spontaneity. *Proceedings of the Aristotelian Society*, 103(3), 227–257.

Taylor, H. 2015. Against unifying accounts of attention. *Erkenntnis*, 80(1), 39-56.

Taylor, H. 2018. Attention, psychology, and pluralism. *The British Journal for the Philosophy of Science*, 69(4), 935-956.

Toribio, J. 2007. Nonconceptual content. *Philosophy Compass*, 2(3), 445-460.

Upton, C. L., & Brent, M. 2019. Meditation and the scope of mental action. *Philosophical Psychology*, 32(1), 52-71.

Velichkovsky, B. B. 2017. Consciousness and working memory: Current trends and research perspectives. *Consciousness and cognition*, 55, 35-45.

Watzl, S. 2011a. The nature of attention. *Philosophy compass*, 6(11), 842-853.

Watzl, S. 2011b. Attention as Structuring of the Stream of Consciousness. In C. Mole, D. Smithies & W. Wu (eds.), *Attention: Philosophical and Psychological Essays*, Oxford: Oxford University Press, 145-173.

Weinberg, J., & Meskin, A. 2006. Puzzling over the imagination: Philosophical problems, architectural solutions. In S. Nichols (ed.), *The Architecture of the Imagination: New essays on pretence, possibility, and fiction* (Oxford: Oxford University Press, 175-202.

Wenzlaff, R. M. & Wegner, D. M. 2000. Thought suppression. *Annual Review of Psychology*, 51, 59-91.

Williams, I. R. (2021): *Concepts and Representational Format: The Structural Limits of Conceptual Representation*. Monash University. Thesis. <https://doi.org/10.26180/17021705.v1>

Wittgenstein, L. 1958. *The Blue and Brown Books: Preliminary Studies for the 'Philosophical Investigation'*. Oxford: Blackwell.

Wittgenstein, L. 1980. *Remarks on the philosophy of psychology, Volume I*. G. E. M. Anscombe. (trans.), G. E. M. Anscombe & G. H. von Wright (eds) Oxford: Blackwell.

Wollheim, R. 1980. *Art and its Objects*. Cambridge: Cambridge University Press.

Wright, W. (2015). Nonconceptual Content. In M. Matthen, M. (ed.) *The Oxford handbook of philosophy of perception*. Oxford: Oxford University Press, 181-197.

Wu, W. 2013. Mental action and the threat of automaticity. In A. Clark, J. Kiverstein & T. Vierkant (eds.) *Decomposing the Will*. Oxford: Oxford University Press, 244-261.

Wu, W. 2014. *Attention*. New York: Routledge.