

Affect in Action¹

Aaron Glasser & Zachary C. Irving

Forthcoming in Australasian Journal of Philosophy. Please cite published version.

Abstract:

Obsessive thinking is a problem case for the philosophy of mental action, insofar as it both (1) feels passive but (2) manifests our agency. Our solution to this “Puzzle of Obsessive Thinking” rests on a fundamental distinction between what we call “occurrent” and “aggregative” agency. Occurrent agency reflects the agent’s capacity to guide her current behavior and thoughts as they unfold over time. We argue that obsessive thinking is a form of occurrent mental agency, since the agent’s attention is guided at the personal level, endorsed, and resistible. Our paper’s first contribution is therefore to argue for the heterodox view that obsessive thinking is active and, thus, for a de-intellectualized, emotional, form of action. Why, then, do obsessive thoughts feel passive? We argue that this is because they undermine aggregative agency. Aggregative agency reflects the agent’s capacity to organize and distribute her actions over time. Although each *episode* of obsessive thinking is guided, the sheer frequency of those episodes undermines the agent’s ability to organize her mental actions. Obsessive thinking is therefore occurrently active but aggregatively passive. Our paper’s second contribution is therefore to use obsessive thinking as a wedge to pry these forms of agency apart.

Keywords: Mental Action, Guidance, Affect, Obsessive Thinking, Endorsement

¹ This project is indebted to conversations with and comments from Elizabeth Beckman, Sarah Buss, Jordan Bridges, Kalina Christoff, Paul de Font-Reaulx, Malte Hendrickx, Cheryl Irving, Robyn Irving, Sheisha Kulkarni, Samuel Murray, Emma Neto, Arcadio Oranday, Peter Railton, Jodie Russell, Susanna Siegel, Chandra Sripada, Noam Tiran, Evan Thompson, Rebecca Todd, Angela Sun, Anna Wang, Margot Witte, and audiences at the University of Michigan, University of Virginia, Society for Philosophy and Psychology, and European Philosophical Society for the Study of the Emotions. This research was funded by the National Endowment for the Humanities (Award #293825) and 3Cavaliers.

“I begin to feel as though my thoughts and voice here
are in some way the creative products of something outside me, not in my control,
and yet that this shaping, determining influence outside me is still me.”

David Foster Wallace, *Here and There*

Obsessive thinking takes us by storm and dominates our minds. This is especially true for patients in the throes of clinical disorders such as depression (Nolen-Hoeksema et al., 2008; Russell, 2021; Watkins, 2008). In a way, obsessive thinking feels paradigmatically passive: emotionally-loaded ideas intrude upon us, draw us in, and carry us along. As Foster Wallace depicts it, it is “as though my thoughts... are... the creative products of something outside me” (1989). Yet in another way, obsessions seem to issue from us. We tend to obsess over topics that gravely matter to us as agents. Obsessive thinking is therefore a problem case for the philosophy of mental action, insofar as it both (1) feels passive and (2) manifests our agency. Foster Wallace captures this tension when writing “this...determining influence outside me is still me.” Similarly, Frankfurt observes that “it is tempting, indeed, to suggest that [obsessional thoughts] are not thoughts that we think at all, but rather thoughts that we find occurring within us...Nonetheless I believe that there is a useful distinction to be made, however awkward its expression, between passions [obsessive thoughts] with respect to which we are active and those with respect to which we are passive” (Frankfurt, 1988; cf. O’Brien, 2013). Let’s call this the “Puzzle of Obsessive Thinking.”²

Our solution to the puzzle rests on a fundamental distinction between what we call “occurrent” and “aggregative” agency. Occurrent agency involves the agent’s capacity to guide her current behavior and thoughts as they unfold over time. Obsessive thinking often feels

² The Puzzle of Obsessive Thinking poses a different challenge for action theory than Angela Smith’s (2004) argument that we can be responsible for attitudes such as forgetting a friend’s birthday. Smith’s challenge is roughly this: (1) we do not choose our attitudes and (2) many assume that choice is necessary for action. But (1) and (2) seem to conflict with (3) the intuition that we can be responsible for attitudes. Our challenge is almost the opposite. (1*) Obsessive thinking is guided, at the personal-level, endorsed, and resistible and (2*) we assume this is jointly sufficient for occurrent action. But (1*) and (2*) seem to conflict with (3*) the intuition that obsessive thinking is passive. We discuss Smith more, and draw on her positive view, in §3.2.

active, we argue, because it involves a form of occurrent agency: guided attention (§1). During obsessive thinking, attention often meets what many action theorists assume are jointly sufficient conditions for action: it is guided (§2), at the personal level (§3.1), endorsed (§3.2), and resistible (§3.3).

Our paper's first contribution is therefore to argue for the heterodox view that obsessive thinking is active (against Cassam, 2011; Frankfurt, 1988; Martin & Tesser, 1996; Metzinger, 2013; Nolen-Hoeksema, 2000; Russell, 2021 and others). Obsessive thinking is therefore a new case study for philosophers who aim to “de-intellectualize” action theory (Dreyfus, 2005; Railton, 2006; Brownstein, 2018; Montero, 2016). Such philosophers appeal to cases of skill, expertise, or habit to argue that actions need not involve intellectually demanding mental states and processes, such as reflective decisions, deliberation, or goals. Our paper provides converging evidence for anti-intellectualism about action, from cases where *affect* guides obsessive thinking. (This way of stating our first contribution assumes a conditional that many action theorists accept: if ϕ -ing is guided at the personal-level, endorsed, and resistible, then ϕ -ing is an action. Readers who reject this conditional can rephrase our first contribution as follows: we show that obsessive thinking involves a strong form of agency.)

These arguments put pressure on the idea that affect and agency are in conflict, which has been a dominant theme in action theory since classical Greece (Plato, *Phaedrus*, secs. 256b, 253d–254) and India (*Katha Upanishad*, Verses 1.3 3–9). These traditions characterize affect as something the agent should master: the agent is a charioteer who must rein in the horses of emotion. Obsessive thinking may seem to support this view: for example, Plato appeals to Leontius' obsession with corpses to illustrate how the passions (*pathē*) undermine agency (*Republic* IV 439–441). Instead, we argue that a careful study of obsessive thinking provides a new reason to *reject* this classical picture of emotions as mere passions: emotions can result in

a relatively strong form of action, which is guided, at the personal-level, endorsed, and resistible.

Why, then, does obsessive thinking feel passive? We argue that this is because it undermines aggregative agency. Aggregative agency reflects the agent's capacity to organize and distribute her actions over time, to make her life internally coherent in the long run. Although each *episode* of obsessive thinking is active, the sheer frequency of those episodes undermines the agent's ability to effectively organize her mental actions over time. By crowding out other mental actions, obsessive thinking makes the agent's mental life less internally coherent in the long run. Obsessive thinking is therefore occurrently active but aggregatively passive. Our paper's second contribution is therefore to use obsessive thinking as a wedge to pry these forms of agency apart.

This paper has five parts. §1 offers a general account of guided attention. §2 draws on this account to argue that obsessive thinking is guided. §3 considers and rejects three explanations for why obsessive thinking is passive: that such thoughts are subpersonal (§3.1), unendorsed (§3.2), or irresistible (§3.3). We argue that obsessive thinking is often active in all three respects. §4 then presents our explanation for why obsessive thinking feels passive—it undermines aggregative agency. This explanation resolves the Puzzle of Obsessive Thinking while capturing the phenomenology of clinical disorders involving obsessive thinking.

§1: Guided Attention

We will argue that obsessive thinking involves a kind of mental action: guided attention (§2; cf. Irving, 2016; Watzl, 2017). Before we can make this argument, we need a general account of guided attention.

A theory of guided attention involves two technical concepts: attention and guidance. While our account of guidance does most of the heavy lifting, we should briefly say what we mean by “attention”. Broadly speaking, we side with philosophers who characterize attention as a *focused mode of consciousness*. Consider an example that we will return to below:

Zadie is planning a trip to Mexico. While doing so, Zadie’s focus shifts from a travel website, to a mental list of museums and beaches, to a buzzing fly, and back to the travel website.

These are shifts in Zadie’s *attention*; in what occupies the foreground of her conscious mind. We are committed to this much. Yet we remain neutral between many theories of focused consciousness (e.g. Titchener, 1910; Smithies, 2011; Watzl, 2017)³ and, therefore, attention. Philosophers who are committed to an incompatible theory can translate “guided attention” to “guided focused consciousness”.

That brings us to our second technical concept: *guidance*. We hold the Frankfurtian view that guidance is a central component of both bodily (Frankfurt, 1978; Pacherie, 2008; Railton, 2006; Setiya, 2007) and mental action (Buehler, 2022; Irving, 2021, 2016; Proust, 2013; Watzl, 2017). Whether a behavior is guided or not depends on how it is monitored and regulated as it unfolds over time. Specifically, guidance involves two forms of control, which we call “proactive” and “regulatory” control. Proactive control disposes the agent to perform an action. But no agent is perfect; unanticipated circumstances and mistakes often interfere with our best laid plans. Regulatory control mechanisms compensate for this interference, monitoring for and correcting errors. Regulatory control is counterfactual, in that it brings wayward action back on track. Through regulatory control, agents not only initiate actions, but *maintain* actions as they unfold. Regulatory control also makes action *normative*: the agent has

³ We do not assume that focused consciousness and attention are *identical*. Rather, we assume only that focused consciousness is constitutively sufficient for attention: every instance of focused consciousness is token identical to an instance of attention. We therefore leave open the possibility of both unconscious attention (Mole, 2014) and unfocused attention (e.g. gist perception).

internal standards for what counts as correct behavior and errors, and she regulates her conduct accordingly.

Agents guide their *attention* when they exert proactive and regulatory control over their focus (Irving & Glasser, 2020; Figure 1). Think of Zadie, who is deliberately planning her trip. Zadie proactively controls her attention, insofar as she directs her focus towards information relevant to the trip: the distance from hotels to beaches/museums, bus versus car rentals, etc. Yet since Zadie is only human, her attention is imperfect. She is briefly distracted by a buzzing fly. At this point, Zadie’s regulatory control kicks in. She quickly notices this “error” of attention—this distraction—and brings her attention back to the task at hand. Zadie’s attention is therefore a manifestation of her agency. Her mind does not passively flit from buzzing flies to idle thoughts. Rather, she uses proactive and regulatory control—uses guidance—to stabilize her attention for enough time to plan her trip.




Figure 1: Attention is guided when it’s subject to proactive and regulatory control. In the case of cognitive control, proactive control biases attention towards task-relevant information. Regulatory control monitors for and corrects task-irrelevant distractions.

In cognitive science, the best studied case of guided attention is cognitive control. Cognitive control involves a family of processes that are realized in the frontal–parietal and cingulo–opercular control networks, which exert proactive and regulatory control over attention (Braver, 2012; Miller & Cohen, 2001). Specifically, cognitive control mechanisms maintain attention to information in the agent’s *task set*, which represents what is relevant to the agent’s current task.⁴ Proactive cognitive control processes bias attention towards

⁴ Prospective memory (e.g. remembering to take your pills with breakfast) may seem like a counterexample to our thesis, since it operates outside attention but uses cognitive control (Murray and Henne, 2023). Yet this

information in the agent’s task set, which usually keeps the agent on task. Regulatory cognitive control processes are “late correction mechanisms” that stand ready to bring attention back on task, whenever the agent becomes distracted (Braver, 2012). During action, late correction mechanisms monitor attention: specifically, they broadcast error signals when one attends to information outside one’s task set, which bring one back on task.⁵

Admittedly, attention’s guidedness is insufficient to show that it is active, since guidance is not independently sufficient for action. There are many sub-personal forms of guidance, such as the mechanisms that dilate and contract your pupils in response to error signals (Frankfurt, 1978, p. 159). This guidance is sub-personal because neither the *error signals* nor the *response* to those signals involve the person.

In contrast, attentional guidance occurs at the personal-level. The agent experiences and responds to the normative representations at the heart of proactive and regulatory control (Irving, 2021, 630–632). When Zadie is planning her trip, she attends to travel websites because they *feel relevant* to her and ignores a buzzing fly because it *feels distracting*. Zadie’s personal-level feelings of relevance and distraction are *token-identical to* correctness and error-signals, respectively. When Zadie experiences those signals, she responds at the personal level to either maintain or refocus her attention. Those normative representations (error and correctness signals) and responses (the refocusing of attention) are *partially constitutive* of attentional guidance. So Zadie’s attention is guided at the personal level. (To clarify: attentional guidance does not merely *cause* personal-level experiences. Rather, personal-level feelings of relevance

objection conflates two kinds of prospective memory. Laboratory prospective memory tasks require cognitive control; but they *also require attention* to resolve ongoing conflicts between two tasks. In contrast, remembering to take your pills with breakfast doesn’t require sustained attention; but it also *cannot* require cognitive control. This type of prospective memory keeps the agent ready to perform a *vast* number of tasks, which far outstrip limited cognitive control resources (Murray, 2024).

⁵ Braver (2012) uses the term “reactive control” for what we call “regulatory control.” We use the term “regulatory control” to clarify that agents monitor guided behaviors *while they unfold over time*.

and distraction are *token-identical* to the correctness and error signals that *partially constitute* attentional guidance.)⁶

§2: Affective Guidance

Attention is guided during obsessive thinking. Before we present our argument for this claim, we should clarify its scope. Our aim is not to *define* obsessive thinking, but rather to trace the *implications* of obsessive thinking for action theory. To do this, we pick out obsessive thinking by ostension: we point to clear cases of obsessive thinking in thought experiments and studies of neurotypical and clinical populations. Since our aim is not to define obsessive thinking, we needn't explain what (if anything) unifies this family of cases or distinguishes them from other cases where attention is guided by affect. Rather, we intend to explain the *implications* of these cases for mental action.

Those implications begin with our thesis that obsessive thinking is guided. That is, obsessive attention is subject to both proactive and regulatory control. Consider Peter:

Peter is harshly criticized by his boss, Lumbergh. Peter fixates on this humiliation. At any free moment, his mind is drawn to the reprimand. He ruminates about it over and over again: “Is this normal? Should I quit? Is Lumbergh right?” At home, Peter feels distracted by the hobbies he typically loves. When he tries to read, for example, his attention is repeatedly drawn back to the reprimand. Frustrated, he leaves his book for another time.

Peter's attention is proactively controlled. Recall that an agent's attention is proactively controlled by a goal when she is disposed to attend to goal-relevant information. Here, Peter's attention is proactively controlled *by his affective state*. He is disposed to attend to affectively-

⁶ What to say about *unconscious* processes that filter attention is beyond the scope of our paper. Our discussion of attentional guidance is limited to cases where error and correctness signals are token-identical to personal-level metacognitive feelings of relevance. Such feelings are *sufficient* to make attentional guidance personal-level (Irving, 2021). This sufficiency claim is all we need to establish that obsessive thinking is guided at the personal level. Whether *unconscious* attentional filtering can also be personal-level (e.g. because it's reasons-responsive) is beyond our scope.

relevant information. Peter's attention is repeatedly pulled towards the harsh critique, its consequences, and his feelings of distress. His attention is drawn to this work-related content in many contexts: while sitting on the couch, attempting to read, and even when trying to sleep. Peter's affect alters what feels worthy of attention.

Empirical research on *affective biases* provides extensive evidence that affective states can proactively control attention in this way. Affectively biased attention "refers to selective attention processes by which sensory systems are tuned to favor certain categories of affectively salient stimuli before they are encountered" (Todd et al., 2012). Affective biases modulate attention in the very same way proactive control does: by probabilistically increasing an agent's tendencies to attend to (affect) relevant information. Indeed, affective neuroscientists have argued that "[t]he primary function of moods...is to modulate or bias attention" (Davidson, 1994, p. 52).

Obsessive thinking is also subject to regulatory control (Figure 2). Recall that regulatory control mechanisms monitor and compensate for interference. In the case of attention, regulatory control mechanisms detect distractions and reorient attention to relevant information. By doing so, regulatory control mechanisms keep attention stable: they allow the agent to focus on specific information over time. Affect can control attention in precisely this way. Obsessive rumination is remarkably stable: once the agent begins to ruminate about a topic, he will likely fixate on this topic for substantial periods of time (Nolen-Hoeksema et al., 2008; Russell, 2021; Watkins, 2008). When Peter begins to think about his work problems, his attention dwells there.

Rumination is stable, in part, because it feels exceptionally difficult to disengage from a ruminative stream of thoughts (Joormann et al., 2011). When Peter tries to reorient his attention away from work—say, to a novel—he feels discomfort and his attention is drawn back to the work conflict. Peter's discomfort is the affective equivalent of the *feelings of*

irrelevance that keep goal-directed thinkers on task. Specifically, Peter’s affective state functions as a kind of regulatory control or “late-correction mechanism” (Braver, 2012). When Peter’s attention begins to stray from the work conflict, an affective mechanism detects this “error” and kicks in to reorient his attention back to work. Peter’s attention is therefore subject to both proactive and regulatory control; his attention is guided.

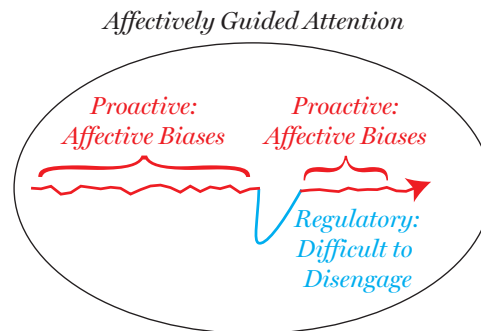


Figure 2: Obsessive thinking is a form of guided attention because it is subject to proactive and regulatory attentional control. Affect exerts proactive control by biasing attention towards affect-incongruent information. Affect exerts regulatory control by making it difficult to disengage and re-orient attention to affect-incongruent information.

We discuss the *dispositions* that characterize proactive and regulatory affective control more than the *categorical grounds* of those dispositions (Irving, 2021). That is by design: we hope to avoid contentious issues about how affective guidance is realized in the brain. We are committed, however, to these realizers including some sort of *representational architecture*. By analogy, cognitive control guides attention towards information within a subject’s *task set*, which are mental states within the executive system that encode what is *relevant to the agent’s task* (Cohen, 2017). Such representations are crucial to guidance’s normative profile: the agent—in some sense—takes certain information to be relevant and other information to be a distraction, relative to her task. Todd and colleagues propose that attention can similarly be guided towards information in an *affective control set*, which represents what is *relevant to the agent’s emotional state* (Todd et al., 2012). We assume that something like an affective control set exists and is part of the categorical grounds of the dispositions of affectively guided attention. Yet we remain neutral about which mental states encode affective control sets:

candidates include emotions and/or priority maps (Todd & Manaligod, 2018), salience (Archer, 2022; Watzl, 2017), intentions and/or task-sets (Wu, 2023), executive control (Buehler, 2022; Hendrickx, 2023). Which of these states ground affective guidance does not bear on the action-theoretic questions of this paper.

Nor does that empirical question influence how affective guidance relates to clinical conditions associated with obsessive thinking. Ratcliffe (2014) argues that depression is partly constituted by the inability to entertain certain kinds of possibilities in thought. Similarly, neurophenomenologists claim that the breadth of one's affordances—which actions-possibilities the agent is disposed to—can be restricted in both depression and OCD (de Haan et al., 2013). Fernando Pessoa echoes this in *The Book of Disquietude*: "In these times of subtle grief, it is impossible—even in dreams—to be a lover, to be a hero, to be happy. All of this is empty, even the idea that it exists" (1996, p. 115). Affective guidance can partially explain this restriction of possibility within clinical conditions associated with obsessive thinking. Someone in the grips of depression is biased towards a smaller subset of what is perceptually available to them (Liu et al., 2012). Thoughts of joy and success may be less likely to come up and feel utterly irrelevant if they do. For even if they begin to entertain these thoughts, their affect may guide their attention back to sadness and defeat, reinforcing their constricted worldview.

Before we move on, a dialectical point: our thesis that obsessive thinking is guided does not settle whether it's a form of (excessive) *inquiry* (Kampa, 2020; Hubacher-Haerle, 2023). Given attention's centrality to inquiry (Friedman, 2020; Siegel, 2025), our model could explicate action-theoretic and psychological mechanisms that make obsessive inquiry excessive. We remain neutral, though, since this depends on difficult philosophical and empirical questions, such as whether obsessive thinking always has zetetic aims and whether its excessiveness involves non-attentional zetetic activities (e.g. inferences or judgments).

We've argued that obsessive thinking is a form of affectively guided attention. By doing so, we've made headway on the Puzzle of Obsessive Thinking: namely, that obsessive thinking (1) feels passive but (2) manifests our agency. We've taken a first step towards understanding how obsessive thinking manifests agency—it is guided by affect—and thus why people who suffer from rumination are not mere patients, but rather agents with respect to their thinking. But we haven't yet fully reclaimed agency for obsessive thinkers: we haven't shown that obsessive thinking is a *mental action*, since non-actions like pupil dilation can also be guided. §3 will support that stronger conclusion by arguing that affective guidance (often) is personal level (§3.1), endorsed (§3.2), and resistible (§3.3).

§3: Affective Guidance as Mental Action

Obsessive thinking is guided by affect. But it does not yet follow that obsessive thinking is a mental action, since many non-actions are guided. We will consider, and reject, three existing arguments that obsessive thinking is passive because it is sub-personal (§3.1), unendorsed (§3.2), and irresistible (§3.3).

Our replies do not merely dispel objections but furnish a positive argument that obsessive thinking is often active. We show that obsessive thinking *is* often guided at the personal-level, endorsed, and resistible. And if ϕ -ing is guided at the personal-level, endorsed, and resistible, then ϕ -ing is an action (according to many theories). So obsessive thinking is a mental action (according to many theories).

Our argument assumes that the preceding four conditions—guidance, at the personal-level, endorsement, and resistibility—are *jointly* sufficient for action. Although this assumption is intuitive, not all action theorists will accept it. Some will reject our argument that endorsement can be non-reflective (§3.2), for example, whereas others will reject our causal

approach to action altogether. It is outside the bounds of our paper to resolve these systematic debates within action theory. But even if you reject our assumption—and thus the main conclusion of §3—you can still take something from our argument: obsessive thinking often has *far more* features of agency than philosophers and cognitive scientists assume. From this, a puzzle of Obsessive Thinking emerges: if obsessive thinking is a manifestation of our agency, why does it feel passive?

3.1: Personal Level

Not all guided behaviors are actions, since guidance is often sub-personal. Sub-personal guidance mechanisms adjust your pupil diameter, but those adjustments are not actions (Frankfurt, 1978, 159). Researchers similarly argue that obsessive thinking is not an action because it is caused by *sub-personal* affective mechanisms (Martin & Tesser, 1996; Metzinger, 2013; Treynor et al., 2003).

But the phenomenology of affective guidance suggests that it operates at the personal-level. Recall that phenomenology is relevant to whether error and correctness signals—and thus attentional guidance—are personal-level (§1). Pupils constrict in response to error signals that represent too much light hitting the retina. Yet those error signals, and the iris' response to those signals, are all sub-personal. So pupil dilation is guided at the sub-personal level.

The opposite is true of attentional guidance. When Zadie is planning a trip, for example, she experiences personal-level correctness signals (e.g. she's drawn to travel websites) and error-signals (e.g. she feels distracted by a buzzing fly). Zadie then adjusts her attention, at the personal level, in response to those normative signals. So her attention is guided at the personal level.

Affective guidance involves the person for a similar reason. When Peter is fixating on his work problems, he feels uncomfortable when he tries to distract himself with a book. This is because reading feels *irrelevant* relative to the all-consuming conflict with his boss. Peter's experience is token-identical to a personal-level error signal. Peter then responds to this feeling, at the personal-level, by bringing his attention back to his affectively-relevant thoughts about work. Peter's meta-cognitive feelings and responses to them are not merely causally related to his attentional guidance: they are *partly constitutive* of attentional guidance. So Peter's attention is guided at the personal level.

A corollary of our argument is that obsessive thinking differs from cases where affect is causally upstream, rather than partly constitutive, of guidance. Guidance regulates action *as it unfolds over time*. Guidance therefore differs from causes that are *upstream* of action, in the sense that their causal role is exhausted once an action has begun (Frankfurt, 1978). One familiar kind of upstream cause occurs during *ballistic* activities such as throwing a ball, which the agent cannot regulate after they have begun.

Another less-discussed kind of upstream cause is the *selection* of an attentional control set (Murray, 2024). Consider the feeling of mental effort, which arguably motivates us to switch from our current task to one with higher expected value (Kurzban et al., 2013; Dixon, 2012). Mental effort *influences which attentional control sets* we use to guide attention. Exhaustion can cause a student to guide her attention to her messages instead of a lecture. But this causal role is upstream of attentional guidance: after this student switches to her messages, effort's causal role is complete.

In contrast, emotions play direct causal roles in how attention is guided during obsessive thinking. Peter's metacognitive feelings of (ir)relevance do not merely motivate him to *begin* thinking about a work conflict—they are partially constitutive of how his attention is

guided *while he is thinking about the conflict*. Peter's emotions partially constitute his personal-level attentional guidance, in a way that upstream causes such as mental effort do not.

3.2: Endorsement and Identification

We've argued that obsessive thinking is guided by a personal-level affective process. Yet one could object that obsessive thinking is still passive because the thinker does not *want* to be guided by affect. Peter does not *want* to fixate on Lumburgh's criticism at the expense of his hobbies, well-being, and sleep. So his obsessive thinking is passive.

Frankfurt (1988, 59–60) argues that obsessive thinking is passive for similar reasons: we do not *identify with* or *endorse* obsessive thinking. Roughly, we identify with/endorse a thought or behaviour P when it aligns with our self-conception. Equivalently, P aligns with our higher-order desires: we *want to want* P. Peter may feel guided to think about Lumberg's criticism, for example, but he does not want to be guided in this way. Frankfurt's substantial insight is that identification/endorsement is (at least part of) what allows us to attribute guidance *to the agent*, rather than "some mechanism with which he *cannot be identified*" (Frankfurt, 1978, p. 159). Whether people identify with/endorse obsessive thinking therefore bears on whether it is a form of mental action or agency.

Do people endorse/identify with their obsessive thinking? This depends on how we interpret Frankfurt, who is not consistent throughout his oeuvre (Moran, 2002). We'll give independent reasons to reject two strong accounts of identification/endorsement and adopt a more moderate account. On this moderate reading, people often identify with/endorse their obsessive thinking because it issues from their cares. Contra Frankfurt, therefore, the most plausible interpretation of his theory suggests that obsessive thinking is often active.

Let's start with the strongest reading:

Endorsement as Second-Order Evaluative Affirmation: S actively thinks p only if S evaluatively affirms p—that is, S reflectively considers p to be aligned with her true values and cares.

Most obsessive thinking is passive on this strong reading. Peter does not reflectively decide that his fixation on Lumbergh aligns with his true values and cares. There are, however, classic reasons to reject this strong reading. For one, it generates a vicious regress (Watson, 1975). Evaluative affirmation is an action. So we must evaluatively affirm our evaluative affirmation of p. But that is an action... and so on. Further, this strong reading over-intellectualizes action, since reflective endorsement requires sophisticated cognitive capacities that many agents (e.g. children and animals) lack.

A weaker reading uses evaluative judgments only to analyze disavowal (rather than endorsement):

Disavowal as Second-Order Evaluative Denial: S passively thinks p if S evaluatively denies p—that is, S reflectively considers p to be misaligned with her true values and cares.

Many obsessive thoughts are passive on this reading. Peter may reflect on and disavow his obsessive thoughts, “I don't *actually* care about Lumbergh's comments. I wish I could ignore them and just read my book.”

We have two responses. First, thinkers—obsessive or otherwise—often do not reflectively consider, let alone disavow, their thoughts. Peter might simply not reflect on how his obsessions fit with his values or whether he'd like them to stop. He might even think these thoughts are important, revealing a genuine problem (we'll explore this possibility below).

Second, this reading is too strong because the values we reflectively endorse are not always those that guide us. Consider a modified version of Railton's Hal, a department chair who considers himself egalitarian to the bone (2006, pp. 24–25). At a college-wide function, Hal exhibits the hierarchical norms he detests: he interrupts a conversation of junior colleagues by placing his hand on their backs, then quickly pivots away once a Dean walks by. Hal's

internalized hierarchical norms guide and explain his behavior, even though he reflectively disavows those norms. So on second-order evaluative readings of Frankfurt, Hal's behavior at the function is passive.

But Hal's actions manifest something deep about his agency—about how he experiences and acts in the world.

“Hal certainly did not see himself as following such a [hierarchical] norm, and would be surprised to be told that hierarchical standing was even at issue... But we won't understand his practical framing of the situation until we recognize how it was structured by relations of hierarchy, among many others... If we wish to give a portrait of Hal 'from the inside out' as he flits from colleague to colleague that afternoon, or turns on his heel to face the Dean, we must enter more deeply into his perspective and agency than the level at which he self-narrates his actions. Indeed, to understand his own reasons, Hal himself would need to do this.” (Railton, 2006, 26–28).

Cases like Hal give us one reason not to analyze identification/endorsement in terms of reflective affirmation or denial. Frankfurt gives us another reason: someone can “become resigned to what he judges to be his defects. A person may acknowledge to himself that passions of which he disapproves are undeniably and unequivocally his...” (Frankfurt, 1998, p. 65). Hal may wish to be an egalitarian but concede that he is not. This is possible only if identification/endorsement is not reducible to reflective affirmation/denial.

Our moderate reading of identification explains how agents can endorse non-reflective actions⁷:

Endorsement as Identification-With: S actively thinks p only if S identifies with p—that is, p is aligned with S's true values and cares.

Hal *reflectively* disavows the hierarchical norms that guide his actions towards his Dean and his junior colleagues. Yet he still *identifies-with* hierarchical norms, insofar as they align with his deep values and cares.

⁷ Frankfurt comes closest to this reading when he talks of coherence/harmony between first- and higher-order desires (1971; 1998, p. 64; 164-5).

Philosophers have developed various versions of the identification-with view (Smith, 2005; Sripada, 2016; Arpaly and Schroeder, 1998; Bennet, 1974; Brownstein, 2018). According to these philosophers, an agents' true values and cares may come apart from the values they reflectively endorse. Agents like Hal therefore identify-with their actions: he may well care about or value hierarchy, even if he is a reflective egalitarian.

To be sure: we do not always identify with obsessive thinking. Consider Ernie, who discovers that his feet are unusually hairy. He might spend the afternoon Googling hairy feet for comparison, wondering, "What am I, a Hobbit?!?!", and searching for pictures of his family in sandals to see if the affliction is genetic. Nevertheless, Ernie might not care that deeply about his hairy feet. His obsession may be an emotionally intense, but idle curiosity, rather than stemming from a deep care. Ernie does not identify with/endorse his obsessive thoughts. We therefore hold neither that Ernie's attention is active, nor that it manifests his agency.

But Ernie's experience is unusual: ruminators often deeply care about—and thus identify with—the topics they dwell on. Peter obsesses about Lumbergh's reprimand because it *struck a nerve*. He was humiliated, and his dignity *matters* to him. Peter deeply cares about the comments and how they were delivered, which is why they occupy so much of his attention.

Four converging lines of evidence suggest ruminators are more often like Peter than Ernie. First, people report rumination to be self-revelatory. Morewedge and colleagues (2014) presented participants with thirteen types of thought, including rumination. Participants then rated, on a scale from 1 to 7, the extent to which each type of thought "reveals insights about myself." Participants rated rumination above the midpoint on the self-insight scale, at similar levels as three prototypical mental actions: planning, problem solving, and deliberation. Since philosophers widely hold that we identify with planning, problem solving, and deliberation, these results suggest that ruminators often identify with their thoughts. For if rumination reveals something deep about oneself, it likely reflects one's cares.

Second, diary and ecological momentary assessment studies suggest that everyday rumination is correlated with unsatisfactory progress on *personally important goals* (see Watkins & Roberts, 2020 for a review). In our terminology, such goals are “cares”.

Third, people ruminate when they encounter obstacles to accomplishing abstract goals (e.g. being treated with respect in the workplace) rather than concrete ones (e.g. responding to an email by end-of-day) (Martin & Tesser, 1996). What likely drives this effect is that people *care* about their abstract goals. To see why, consider two violin students, Ida and Bob. Ida plays the violin for its own sake but doesn’t associate music with other abstract goals. Bob plays the violin to achieve abstract goals of college admissions, money, and happiness. If the only violin teacher in town retired, Ida would likely ruminate because something she cares about is threatened. Bob would instead pivot to another shiny hobby (e.g. math olympics) that will accomplish his abstract goals. This thought experiment suggests that frustrated abstract goals cause rumination because of a third variable: we usually *care* about those goals.

Fourth, our interpretation explains why *concreteness training* is an effective treatment for rumination (Watkins et al., 2008). Concreteness training encourages agents to focus on local sensations in their immediate environment, rather than abstract topics like their personal goals and concerns. Because abstract goals are associated with cares, concreteness training likely helps distance us from thinking *too much* about our deep cares. Consider: Peter can stop ruminating on his workplace humiliation by focusing on the concrete sensations of his environment—his feet on the floor, his breath, etc. These concrete sensations are unrelated to Peter’s abstract care about his dignity. So concreteness training gives Peter respite from ruminating about his deep cares.

These lines of evidence converge on our claim that people often care deeply about—and thus identify with—what they ruminate about. Rumination is therefore often endorsed, on the plausible reading of Frankfurt: our moderate interpretation. So contra Frankfurt’s own

remarks on obsessive thought, the most plausible version of his view entails that agents often identify with their obsessions.⁸

3.3: Resistibility

Obsessive thinking is (often) guided by a personal-level, affective process that the agent endorses. Yet one could still argue that obsessive thinking is passive because it is *irresistible* and thus *uncontrollable*. Control-based theories of action are an outgrowth of the classic thesis that action requires a power for opposites (Fischer & Ravizza, 2000). Roughly, A-ing is an action only if the agent has control processes that make A-ing resistable. One could argue that the desire to obsess about a topic is often too strong to resist. Peter might complain, “I can’t help but think about Lumberg’s criticism; it’s not under my control.”

The problem is that both philosophical arguments (Watson, 1999; Arpaly, 2002) and empirical evidence (Sripada, 2018; Pickard, 2024) suggest that the vast majority of (if not all) desires are resistible. One powerful argument for this premise proceeds by analogy. A careful study of conditions including addiction and Tourette’s suggests that even our strongest desires are resistable (Sripada, 2018; Pickard, 2024; Specht et al., 2014). But the vast majority of desires are easier to resist than those. So the vast majority of desires are likely resistable.

Our argument rests on evidence that *paradigmatically* strong desires are resistable. Addicts do not experience individual drug-directed desires that are too strong to resist (Sripada, 2018; Pickard, 2024). For this reason, addicts can board planes, attend lectures, and see family— all while resisting the urge to take drugs. Agents can even resist quintessentially “uncontrollable” physical behaviors like tics in Tourette’s syndrome. Even without suppression

⁸ One could also argue that ruminators do not identify with their obsessive thinking because those thoughts are not reason-responsive (Cassam, 2011, p. 6). O’Brien (2013, 97–98) provides a convincing response to Cassam: obsessive thinkers typically form their thoughts on the basis of *some* reasons, even if those reasons are epistemically faulty. But perfect reasoning is unnecessary for identification.

training, children with Tourette's are able to resist individual spontaneous tics without subsequent build up (Specht et al., 2014). If addicts and individuals with Tourettes can resist their powerful desires, then individuals with weaker desires can likely resist those as well. And in the vast majority of cases, the desire to ruminate is weaker than the desires that underlie addiction and Tourettes. So the vast majority of obsessive thinking is active, on a simple reading of the control theory.⁹

We assume it's *easier* to resist the desire to obsess about a topic than to resist paradigmatically strong desires such as those in addiction. One might object that this premise is false, since the *aim* of control is distinct in these cases. The addicted agent's aim is to resist *acting* on their drug-directed desire, while the obsessive thinker's aim is to resist *even having* a thought about the content they tend to ruminate over. If this is true, any amount of trying to resist thinking about P will never work, because you are already thinking about P. This casts obsessive thinking as a form of intrusive thoughts, where trying *not* to think of something (e.g. a polar bear) infamously causes one to have a thought about that thing (Wegner et al., 1987; De Haan et al., 2015). So structural differences between how we resist acting on a desire and resist having obsessive thoughts may make the latter harder.

But this objection rests on a faulty assumption: namely, that having an intrusive thought is sufficient for obsessive thinking. Obsessive thinking is a form of *guided attention*: specifically, attention is guided to information one represents as relevant to one's task. But having an intrusive thought—in Wegner's sense—is insufficient for guided attention.

Attention is guided towards information that is represented as *relevant by one's task set*. Weger's "paradox" of thought suppression occurs when some information ι is *represented as a distraction* by one's task set, which causes ι -related-thoughts to become briefly active

⁹ Our response does not require the (plausible!) assumption that desires to ruminate are *never* irresistible. Perhaps obsessive thinkers experience an irresistibly strong desire once in a blue moon. Even granting this, it still follows that ruminators can resist the vast majority of their obsessive thoughts.

before they are suppressed. This explains why intrusive thoughts tend to be brief, about *1.4 seconds* on average (Wegner, 1987). Contrast this with guidance, which functions to *sustain* attention for significant periods of time. In obsessive thinking, Peter's problem is not that he thinks about work for 1.4 seconds; Peter's problem is that he keeps *dwelling* on the topic for minutes to hours, at the expense of his other projects and desires.

Wegner's paradox therefore does not apply to obsessive thinking. It is self-defeating to try and resist having any thoughts about polar bears, since the suppression of thoughts about P requires thoughts about P. In contrast, it is not self-defeating to try and resist attentional guidance: the suppression of attentional guidance to P does not require attentional guidance to P. Peter can try and stop himself from dwelling on his work conflict, for example, by trying to suppress his emotions or their salience (Glasser, m.s.). Having defused the objection, we conclude that obsessive thinking is resistible, just like acting on a drug-directed desire or a tic in Tourette's.

Of course, individuals with addiction or Tourette's are not fully in control of their actions. Sripada (2018) argues that this is because of the sheer *quantity* of desires they experience, rather than the strength of any individual desire. Consider addiction. For an addict to quit, she must resist dozens to hundreds of drug-directed desires every day. Each individual desire may well be resistible. But the control processes we use to resist desires are fallible. Even in mundane psychology experiments on executive control, such as the Stroop Task, participants make errors on 5–10% of trials (MacLeod, 1992). Given the fallibility of control, addicts who try to resist a vast number of drug-directed desires are likely to slip up at least once. People with addiction may be unable to control their desires long enough to quit, therefore, even if each desire is controllable.

One might argue that obsessive thinking is irresistible for precisely this reason. Obsessive thinkers experience *many* desires to ruminate throughout the day. Given the

fallibility of control, it's inevitable that ruminators give in *at least once*. So obsessive thinking is passive.

But this argument overlooks a key distinction between addicts and ruminators, which concerns the structure of their overall aims rather than the strength of their individual desires. When an addict tries to quit, her goal is to *never* use drugs. Similarly, when a child with Tourette's tries to suppress his tics, his goal is to *never* express them. Given the fallibility of control processes, *never* giving into a frequent urge is likely beyond one's control. But obsessive thinkers do not try to suppress *every* thought about the topic that consumes them. Peter cares deeply about his boss' reprimand (§3.2), so wants to give *some* attention to it. Unlike the cases of addiction and Tourette's, therefore, fallible control and frequent desires are not *sufficient* to explain how obsessive thinking is passive.

Let's pause to take stock of the Puzzle of Obsessive Thinking: that obsessive thinking (1) feels passive but (2) manifests our agency. We've explained how obsessive thinking is agential and thus dulled the puzzle's second horn. Obsessive thinking is guided by personal-level affective processes, which the agent identifies with and can resist.¹⁰ But now we have sharpened the first horn. If obsessive thinking is active, why does it *feel passive*?¹¹

Above, we considered whether the *quantity* of obsessive thoughts is sufficient to make them passive. Although we rejected a version of this argument, it was on the right track. Specifically, we will argue that the frequency of obsessive thinking undermines the thinker's

¹⁰ Our thesis that *attention* is active during obsessive thinking does not settle whether *other* aspects of obsessive thinking are active, such as the thinker's propositional attitudes (e.g. believing and entertaining). While a thinker actively guides her attention to an argument, for example, she may feel *compelled* to believe the argument's conclusion. Plausibly, her belief formation is passive (though see O'Brien, 2013). That one and the same instance of thinking can be active under one description (e.g. as attention) and passive under another (e.g. as belief formation) is an example of the description dependence of action (Anscombe, 1957, §23–26).

¹¹ This puzzle does not reduce to the familiar observation that there are borderline cases of action, such as those involving gradable control (Kelley, forthcoming). Instead, obsessive thinking provides us with a clear, rather than borderline, case of occurrent action that nonetheless feels passive.

ability to organize her mental actions over time. That is, obsessive thinking undermines what we call “aggregative agency” (§4).

§4: Obsessive Thinking and Aggregative Agency

Our solution to the Puzzle of Obsessive Thinking rests on a distinction between what we call “occurrent” and “aggregative” agency. Occurrent agency concerns the agent’s capacity to guide her mental and bodily actions as they unfold over time. We’ve shown that *individual episodes* of rumination involve a robust form of occurrent agency: each episode is affectively guided (§2) at the personal level (§3.1), endorsed (§3.2), and resistible (§3.3). Yet our discussion thus far has focused only on features of agency that concern relatively short timescales. Guidance, endorsement, and resistibility operate at the level of seconds or minutes to regulate ongoing (occurrent) actions (Figure 3). If we focus only on those timescales, and only on individual episodes of obsessive thinking, this thinking looks wholly active.

We can see the passivity of obsessive thinking only if we zoom out and inspect the thinker’s life over longer timescales (hours, days, or weeks). At these timescales, obsessive thinking undermines *aggregative agency*.¹² Here’s a sketch of our argument, which we develop in depth below. Aggregative agency reflects the agent’s capacity to organize and distribute actions over time. Although each *episode* of obsessive thinking is an occurrent mental action, the sheer frequency of ruminative episodes crowds out other mental actions and thus undermines the agent’s ability to organize her mental life (Figure 3). Obsessive thinking is

¹² Because we pick out obsessive thinking by ostension, we cannot say that *all* obsessive thinking undermines aggregative agency. Having a song stuck in your head may count as obsessive, for example, but not undermine your ability to organize your life in the long run. The contentious issue here concerns how to *define* obsessive thinking: can thinking count as obsessive, if it does not interfere with aggregative agency? The authors disagree about this question, so remain neutral.

therefore an occurrent mental action that interferes with aggregative mental agency. To develop this sketch, we first need to explain aggregative agency.

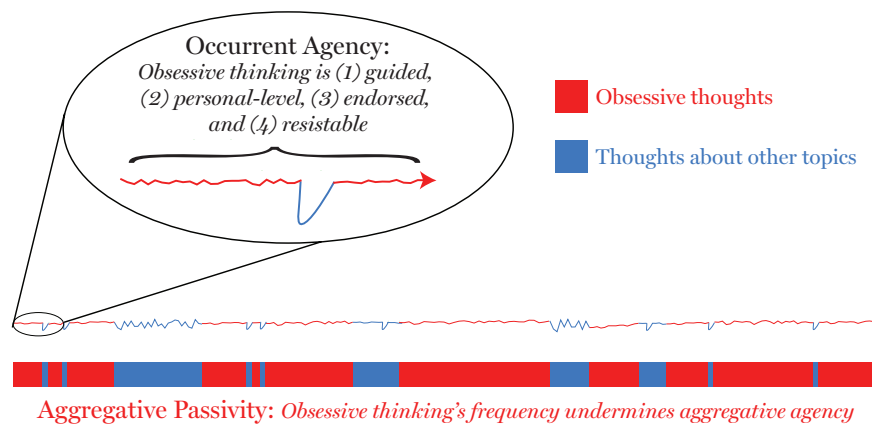


Figure 3: Obsessive thinking is an occurrent action that undermines aggregative agency. If we zoom in (above), each episode of obsessive thinking involves occurrent agency. If we zoom out (below), the frequency of obsessive thinking undermines aggregative agency because it crowds out other mental actions. This problem is illustrated by the prevalence of obsessive thinking (red) over non-obsessive thinking (blue).

Aggregative agency is necessary for creatures like us, who pursue many projects simultaneously. Right now, you may be writing a book, raising a child, teaching a class, dieting, and learning piano. You cannot pursue all those goals simultaneously, since each requires distinct bodily and mental actions. So you must distribute your actions over time, such that you spend *enough time* practicing the piano, planning your meals, prepping your classes, playing with your child, and so on. Furthermore, decisions about how to allocate your time are often underdetermined by your circumstances. In a free hour after lunch, for example, you can write a paragraph, practice chords, or research after-school programs.¹³ But this poses a fundamental problem for human agency: how should you organize your actions over time? We call this the “problem of aggregative agency,” since you must organize many actions into an aggregate that satisfies (enough of) your goals.¹⁴ Crucially, the problem of aggregative agency is distinct from

¹³ Relatedly, Tenenbaum and Raffman (2012) argue that we have latitude over when to pursue “vague projects”, which have no discrete completion conditions.

¹⁴ The closest cousin of our “problem of aggregative agency” is Murray’s “problem of plural goal pursuit” (2024). Murray is concerned with how humans overcome 1) physical and psychological constraints that limit the amount of resources we can give towards our various goals *at a moment*, and 2) cognitive and informational demands associated with allocating our resources across goals *at different times*. Our “problem of aggregative agency” strongly resembles the second issue. We call this the problem of “aggregative agency” rather than “plural goal pursuit” for three reasons. First, to indicate that we are concerned only with *half* of Murray’s problem of plural goal pursuit: how to aggregate our goals over time (issue 2) rather than regulate occurrent

the problem of how to guide individual, occurrent actions. You might successfully guide your thoughts and fingers while writing, for example, but still fail as an aggregative agent because you neglected your other projects.

We remain neutral on exactly how humans solve the problem of aggregative agency. Philosophers and cognitive scientists have proposed many mechanisms to solve this problem, including vigilance (Murray, 2024), plans (Bratman, 1987), the will (Korsgaard, 2008), willpower (Holton, 2009, Ch. 6), implementation intentions (Gollwitzer, 1999), and extended agency (Tenenbaum and Raffman, 2012). We take no stand on which (combination) of these mechanism(s) allows us to organize our actions over time. What's important for our purposes is that human agents need *some* solution to the problem of aggregative agency, *some* mechanism that organizes our actions over time. We call that mechanism “aggregative agency”.

Roughly, aggregative agency can be characterized as an input–output function. The input is your (often implicit) evaluation of what distributions of actions are acceptable, given your values, cares, and commitments (e.g. long-term projects). You might deem it unacceptable to distribute 90% of your waking hours to writing, for example, since this would leave too little time for your family, hobbies, and health. The output of aggregative agency is how you actually distribute your actions over time. Aggregative agency's function is to distribute your actions in a way you deem acceptable. This (over)simplification is sufficient to show why obsessive thinking threatens aggregative agency.

Consider Peter's obsessive thinking about his boss, Lumbergh. Each episode of rumination may be an occurrent mental action, which is guided at the personal level, endorsed, and resistible. Problems arise because Peter engages in the same mental action dozens—or

goals at a time (issue 1). Second, to indicate that we are not committed to a specific formulation of the problem or solution. Third, to highlight the contrast between aggregative and occurrent agency, rather than the contrast between agents who have one versus many goals.

hundreds—of times in a day. Peter may want to distribute *some* thought to Lumbergh: Peter genuinely cares about the reprimand, which struck a nerve. Yet Peter would deem it unacceptable to distribute, say, 70% of his thoughts in a day to Lumbergh. Doing so would crowd out other desirable mental actions like reading a novel or chatting with his neighbor Lawrence. But the affective systems that guide Peter’s attention are insensitive to how he wants to distribute his attention.¹⁵ Peter therefore guides his attention to Lumbergh’s criticism *too often*, which undermines his aggregative agency.

We now have a solution to the Puzzle of Obsessive Thinking. Obsessive thinking *manifests our agency* because it is an occurrent mental action, but still *feels passive* because it undermines aggregative agency.

Obsessive thinking is an important case study for action theory because it cleanly separates occurrent and aggregative agency. Admittedly, we are not the first philosophers to document features of agency that show up only in the aggregate. As discussed earlier, Sripada argues that many addicts can control *individual* drug-directed desires but cannot control the *quantity* of desires necessary for quitting (2018; §3.3). Yet even if individual urges are controllable, they needn’t reflect the addict’s occurrent agency. People with addiction do not (normally) endorse their drug-directed desires (Sripada, 2018). In contrast, each episode of obsessive thinking is a mental action, which the agent guides (§2) at the personal-level (§3.1), endorses (§3.2), and can resist (§3.3). Obsessive thinking therefore shows that a series of full-fledged occurrent actions can combine to make the agent aggregatively passive. Obsessive thinking therefore pries apart two kinds of agency—occurrent and aggregative—that are otherwise tempting to conflate.

¹⁵ We assume that ruminators frequently think about their worries. Why this occurs is an empirical question that is outside the scope of our paper. Our theory is compatible with (but not committed to) Russell’s (2021) thesis that rumination becomes frequent because of a positive feedback loop. Ruminators frequently experience negative affect, so come to associate many environmental and bodily cues with negative affect, which makes them experience negative affect more frequently, and so on.

§5: Clinical Implications

Our solution to the Puzzle of Obsessive Thinking helps capture the phenomenology of obsessive thinking in clinical conditions. Our model reclaims agency for people with clinical conditions involving obsessive thinking, like depression. Since obsessive thinking is pervasive in depression, it is tempting to think of people with depression as mere patients—as victims of alien thoughts. This attitude is baked into how we talk about people with depression as clinical *patients*, rather than agents. We think this is a mistake. People with depression are bona fide mental agents: even while ruminating, their thoughts are guided, at the personal level, endorsed, and resistible. Yet we explain (part of) why ruminators can feel helpless: the frequency of their obsessions interferes with aggregative agency. Our model therefore respects the agency of people with depression, while offering compassion.

One might object that we callously imply that people with clinical depression are *morally responsible* for their thoughts. This objection is too quick. We claim only that obsessive thinking is *attributable* to the agent, since it is an occurrent mental action. But attribution does not entail stronger forms of responsibility like accountability, answerability, or blame (Shoemaker, 2011; Smith, 2012; Pickard & Ward, 2013). Given the nuance necessary to adequately address questions of responsibility in cases of mental illness, answers fall beyond this paper's scope.

Another concern is how clinical populations can *resolve* the conflict between occurrent and aggregative agency. Our model motivates at least two empirically-informed strategies that may help ruminators regain aggregative agency.

The first strategy is to replace rumination with other forms of guided attention. Trait ruminators display not only more *affective* guidance than controls, but also enhanced guided

attention in *goal-maintenance* tasks (Altamirano et al., 2010). Rather than try to stop ruminators from guiding their attention (e.g. through relaxation strategies), it may be more effective to replace negative rumination with more positive/neutral forms of guided attention. To this end, Watkins and colleagues successfully reduced rumination through concreteness training aimed at repeatedly focusing on specific details. Such training generates “adaptive repetitive thought in place of maladaptive repetitive thought” (Watkins et al., 2009; Watkins & Nolen-Hoeksema, 2014). This is an example of transfer-appropriate processing in cognitive bias modification treatment, which attempts to reorient habits (such as guiding attention) instead of extinguishing them (Hertel & Mathews, 2011).

The second strategy is to schedule “worry time” (McGowan & Behar, 2013). The goal is to delay spontaneous urges to worry until a scheduled time, which is fully devoted to their concerns. One advantage of worry time is that ruminators still attend to their worries, which they care deeply about (§3.2). Furthermore, the goal of worry time is not to give ruminators more control over *individual* episodes of obsessive thinking, which they can likely already resist (§3.3). Rather, worry time addresses the core problem with rumination: that one obsesses about the same topic *too often*, which undermines aggregative agency. Worry time strategies could encourage organizational mental actions and rehabilitate aggregative agency.

Conclusion

Since ancient Greece (Plato, *Phaedrus*, secs. 256b, 253d–254) and India (*Katha Upanishad*, Verses 1.3 3–11), philosophers have cast affect as something to be mastered: an emotional horse that our rational selves must reign in. Obsessive thinking might seem case in point. Our attention is dominated by the wild horses of emotion. All we agents can do is crack the reins of cognitive control to redirect our wayward thoughts. Yet this classic picture is

false. Obsessive thinking involves a deep form of occurrent mental action: attention is guided (§2) at the personal level (§3.1), endorsed (§3.2), and resistible (§3.3). If obsessive thinking is a wild horse, occurrent agents are as much horse as rider.

We see how obsessive thinking is passive only on a longer timescale. The frequency of obsessive thinking undermines *aggregative agency*—our capacity to organize actions over time—because it leaves little time to think about anything else. The problem is not that any *individual* horse is out of control; the problem is the size of the herd. Together, their hoofbeats drown out other thoughts. Obsessive thinking therefore motivates a general distinction between occurrent and aggregative agency. What’s true of the horse may be false of the herd.

Obsessive thinking also requires a new account of how to *manage* the passions. Obsessions are best managed, not by a *charioteer* who whips passions into place from above, but by a *nomadic herder*. Nomadic herders such as Sámi “reindeer walkers” *migrate with* a herd of animals, rather than determining each animal’s movement from the top down. Yet reindeer walkers are still managers: they herd, to keep their animals on a safe and plentiful migratory path; they practice animal husbandry, to maintain a healthy herd size through breeding, culling, and so on. Likewise, we must manage obsessions at the level of the herd, not the horse. We ought to herd our obsessions into safe pastures, cull those that have begun to dominate our minds, and breed thoughts about cares that have been given less attention.

References

- Altamirano, Lee J., Akira Miyake, and Anson J. Whitmer (2010) 'When mental inflexibility facilitates executive control: Beneficial side effects of ruminative tendencies on goal maintenance', *Psychological Science* 21, no. 10 1377-1382.
- Archer, Sophie (2022) 'Salience and what matters', In S. Archer (Ed.) *Salience* (pp. 113–129). Routledge.
- Arpaly, Nomy (1999) "Praise, blame and the whole self", *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 93, no. 2: 161-188.
- Bennett, Jonathan (1974) "The conscience of huckleberry Finn", *Philosophy* 49, no. 188: 123-134.

- Bratman, Michael (1987) *Intention, plans, and practical reason*. Cambridge University Press.
- Braver, Todd S. (2012) 'The variable nature of cognitive control: a dual mechanisms framework', *Trends in Cognitive Sciences* 16, no. 2: 106-113.
- Brownstein, Michael (2018) *The implicit mind: Cognitive architecture, the self, and ethics*. Oxford University Press.
- Buehler, Denis (2022) 'Agential capacities: a capacity to guide'. *Philosophical Studies* 179, no. 1: 21-47.
- Cassam, Quassim (2011) 'Knowing What You Believe', *Proceedings of the Aristotelian society*, vol. 111, no. 1_pt_1, pp. 1-23. Oxford, UK: Oxford University Press.
- Cohen, Jonathan D. (2017) 'Cognitive control: Core constructs and current considerations', *The Wiley Handbook of Cognitive Control*: 1-28.
- Davidson, Richard J. (1994) 'On emotion, mood, and related affective constructs', *The nature of emotion: Fundamental questions*: 51-55.
- De Haan, Sanneke, et al. (2013) 'The phenomenology of deep brain stimulation-induced changes in OCD: an enactive affordance-based model', *Frontiers in Human Neuroscience* 7: 653.
- De Haan, Sanneke, Erik Rietveld, and Damiaan Denys (2015) 'Being free by losing control: What obsessive-compulsive disorder can tell us about free will', *Free will and the brain: Neuroscientific, philosophical, and legal perspectives on free will*: 83-102.
- Dixon, Matthew L., and Kalina Christoff (2012) 'The decision to engage cognitive control is driven by expected reward-value: neural and behavioral evidence', *PloS One* 7.12: e51637.
- Dreyfus, Hubert L. (2005) 'Overcoming the myth of the mental: How philosophers can profit from the phenomenology of everyday expertise', *Proceedings and addresses of the American Philosophical Association*. Vol. 79. No. 2. American Philosophical Association.
- Fischer, John Martin, and Mark Ravizza (1998) *Responsibility and control: A theory of moral responsibility*. Cambridge University Press.
- Frankfurt, Harry G. 'Freedom of the will and the concept of a person', *Journal of Philosophy*, 68, (1971).
- Frankfurt, Harry G. (1978) 'The problem of action', *American Philosophical Quarterly* 15, no. 2: 157-162.
- Frankfurt, Harry. G. (1988) 'Identification and externality'. In *The Importance of What We Care About*. Cambridge University Press.
- Frankfurt, Harry. G. (1988) 'Identification and wholeheartedness', In *The Importance of What We Care About*. Cambridge University Press.
- Friedman, Jane (2020) 'The epistemic and the zetetic', *Philosophical Review* 129.4: 501-536.
- Gollwitzer, Peter M. (1999) 'Implementation intentions: Strong effects of simple plans', *American Psychologist* 54, no. 7: 493.
- Glasser, Aaron (m.s.) 'Steering Salience'.
- Hendrickx, M. (2023) 'Agentially controlled action: causal, not counterfactual', *Philosophical Studies*, 180(10), 3121-3139.
- Hertel, Paula T., and Andrew Mathews (2011) 'Cognitive bias modification: Past perspectives, current findings, and future applications', *Perspectives on Psychological Science* 6, no. 6: 521-536.
- Holton, Richard (2009) *Willing, wanting, waiting*. Oxford University Press.
- Irving, Zachary C. (2021) 'Drifting and directed minds: The significance of mind-wandering for mental agency', *The Journal of Philosophy* 118, no. 11: 614-644.
- Irving, Zachary C. (2016) 'Mind-wandering is unguided attention: accounting for the "purposeful" wanderer', *Philosophical Studies* 173, no. 2: 547-571.

- Irving, Zachary C., and Aaron Glasser (2020) 'Mind-wandering: A philosophical guide', *Philosophy Compass* 15, no. 1.
- Joormann, Jutta, Sara M. Levens, and Ian H. Gotlib (2011) 'Sticky thoughts: Depression and rumination are associated with difficulties manipulating emotional material in working memory', *Psychological science* 22, no. 8: 979-983.
- Kelley, Mikayla (2024) 'Separating action and knowledge', *The Philosophical Quarterly*: 75(3), 956-977.
- Korsgaard, Christine M. (2008) 'The normativity of instrumental reason', In *The constitution of agency: Essays on practical reason and moral psychology*. Oxford University Press.
- Kurzban, Robert, et al. (2013) 'An opportunity cost model of subjective effort and task performance', *Behavioral and Brain Sciences* 36.6: 661-679.
- Liu, Wen-hua, et al. (2012) 'Facial perception bias in patients with major depression', *Psychiatry Research* 197.3: 217-220.
- MacLeod, Colin M. (1992) 'The Stroop task: the "gold standard" of attentional measures', *Journal of Experimental Psychology: General* 121, no. 1: 12.
- Martin, Leonard L., & Abraham Tesser (1996) 'Some ruminative thoughts' In *Advances in Social Cognition*, 9: 1-47.
- McGowan, Sarah Kate, and Evelyn Behar (2013) 'A preliminary investigation of stimulus control training for worry: Effects on anxiety and insomnia', *Behavior Modification* 37, no. 1: 90-112.
- McIntosh, William D., and Leonard L. Martin (1992) 'The cybernetics of happiness: The relation of goal attainment, rumination, and affect', In M. S. Clark (Ed.), *Emotion and Social Behavior* (pp. 222-246). Sage Publications, Inc.
- Metzinger, Thomas (2013) 'The myth of cognitive agency: subpersonal thinking as a cyclically recurring loss of mental autonomy', *Frontiers in Psychology* 4: 931.
- Miller, Earl K., and Jonathan D. Cohen (2001) 'An integrative theory of prefrontal cortex function', *Annual Review of Neuroscience* 24, no. 1: 167-202.
- Mole, Christopher (2014) 'Attention to unseen objects', *Journal of Consciousness Studies* 21, no. 11-12: 41-56.
- Montero, Barbara Gail (2016) *Thought in action: Expertise and the conscious mind*. Oxford University Press.
- Moran, Richard. (2002) 'Frankfurt on identification: Ambiguities of activity in mental life', In S. Buss & L. Overton (Eds.), *The Contours of Agency: Essays on Themes from Harry Frankfurt*: 189-217.
- Morewedge, Carey K., Colleen E. Giblein, and Michael I. Norton (2014) 'The (perceived) meaning of spontaneous thoughts', *Journal of Experimental Psychology: General* 143, no. 4: 1742.
- Murray, Samuel, and Paul Henne (2023) 'Bringing Self-control into the Future', *Advances in Experimental Philosophy of Action*: 51.
- Murray, Samuel (2024) 'The Nature and Norms of Vigilance' *American Philosophical Quarterly*, 61(3), 265-278.
- Nolen-Hoeksema, Susan (2000) 'The role of rumination in depressive disorders and mixed anxiety/depressive symptoms', *Journal of Abnormal Psychology* 109, no. 3: 504.
- Nolen-Hoeksema, Susan, Blair E. Wisco, and Sonja Lyubomirsky (2008) 'Rethinking rumination', *Perspectives on Psychological Science* 3, no. 5: 400-424.
- O'Brien, Lucy (2013) 'Obsessive thoughts and inner voices', *Philosophical Issues* 23: 93-108.
- Pacherie, Elisabeth (2008) 'The phenomenology of action: A conceptual framework', *Cognition* 107, no. 1: 179-217.
- Pessoa, Fernando (1996) *The Book of Disquietude* (B. Soares (ed.)). Trans. Richard Zenith.

- Sheep Meadow Press.
- Pickard, Hanna (2024) 'Craving for drugs', *Mind & Language*, 39(5), 598-626.
- Pickard, Hanna, and Lisa Ward (2013) 'Responsibility without blame: Philosophical reflections on clinical practice', *Oxford Handbook of Philosophy of Psychiatry*: 1134-1154.
- Proust, Joëlle (2013) *The Philosophy of Metacognition: Mental agency and self-awareness*. Oxford University Press.
- Railton, Peter (2006) 'Normative guidance', *Oxford Studies in Metaethics* 1, no. 3: 2-33.
- Railton, Peter (2012) 'That obscure object, desire', *Proceedings and Addresses of the American Philosophical Association*. Vol. 86. No. 2. American Philosophical Association.
- Ratcliffe, Matthew (2014) *Experiences of depression: A study in phenomenology*. Oxford University Press.
- Russell, Jodie Louise (2021) 'Stuck on repeat: Why do we continue to ruminate?', *Synthese* 199, no. 5: 13143-13162.
- Setiya, Kieran 2009 'Reasons without rationalism', In *Reasons without Rationalism*. Princeton University Press.
- Shoemaker, David (2011) 'Attributability, answerability, and accountability: Toward a wider theory of moral responsibility', *Ethics* 121, no. 3: 602-632.
- Siegel, Susanna (forthcoming) 'How do lines of inquiry unfold? Insights from journalism', *Oxford Studies in Epistemology: Special Issue on Applied Epistemology*.
- Smith, Angela M. (2005) 'Responsibility for attitudes: Activity and passivity in mental life', *Ethics* 115, no. 2: 236-271.
- Smith, Angela M. (2012) 'Attributability, answerability, and accountability: In defense of a unified account', *Ethics* 122, no. 3: 575-589.
- Smithies, Declan (2011) 'Attention is rational-access consciousness', In Mole, Smithies, and Wu (Eds.) *Attention: Philosophical and psychological essays*: 247-273.
- Specht, Matt W., Cassandra M. Nicotra, Laura M. Kelly, Douglas W. Woods, Emily J. Ricketts, Carisa Perry-Parrish, Elizabeth Reynolds et al. (2014) 'A comparison of urge intensity and the probability of tic completion during tic freely and tic suppression conditions', *Behavior Modification* 38, no. 2: 297-318.
- Sripada, Chandra (2016) 'Self-expression: A deep self theory of moral responsibility', *Philosophical Studies* 173, no. 5: 1203-1232.
- Sripada, Chandra (2018) 'Addiction and fallibility', *Journal of Philosophy* 115, no. 11: 569-587.
- Tenenbaum, Sergio, and Diana Raffman (2012) 'Vague projects and the puzzle of the self-torturer', *Ethics* 123, no. 1: 86-112.
- Titchener, Edward Bradford (1910) 'Attention as sensory clearness', *Journal of Philosophy, Psychology and Scientific Methods* 7, no. 7: 180-182.
- Todd, Rebecca M., William A. Cunningham, Adam K. Anderson, and Evan Thompson (2012) 'Affect-biased attention as emotion regulation', *Trends in Cognitive Sciences* 16, no. 7: 365-372.
- Todd, R. M., & Manaligod, M. G. M. (2018). 'Implicit guidance of attention: The priority state space framework', *Cortex*, 102, 121–138.
- Treynor, Wendy, Richard Gonzalez, and Susan Nolen-Hoeksema (2003) 'Rumination reconsidered: A psychometric analysis', *Cognitive Therapy and Research* 27, no. 3: 247-259.
- Wallace, David Foster (1989) 'Here and There', In *Girl with Curious Hair* (pp. 149–172). WW Norton & Company.
- Watkins, Edward R (2008) "Constructive and unconstructive repetitive thought." *Psychological Bulletin* 134, no. 2: 163.

- Watkins, Ed R., Celine B. Baeyens, and Rebecca Read (2009) 'Concreteness training reduces dysphoria: proof-of-principle for repeated cognitive bias modification in depression', *Journal of Abnormal Psychology* 118, no. 1: 55.
- Watkins, Edward R., and Susan Nolen-Hoeksema (2014) 'A habit-goal framework of depressive rumination', *Journal of Abnormal Psychology* 123, no. 1: 24.
- Watkins, Edward R., and Henrietta Roberts (2020) 'Reflecting on rumination: Consequences, causes, mechanisms and treatment of rumination', *Behaviour Research and Therapy* 127: 103573.
- Watson, Gary (1975) 'Free agency' *Journal of Philosophy* 72, no. 8: 205-220.
- Watson, Gary (1999) 'Disordered appetites: Addiction, compulsion, and dependence', In J. Elster (Ed.) *Addiction: Entries and exits*, 3-28.
- Watzl, Sebastian (2017) *Structuring mind: The nature of attention and how it shapes consciousness*. Oxford University Press.
- Wegner, Daniel M., et al. (1987) 'Paradoxical effects of thought suppression', *Journal of Personality and Social Psychology* 53.1: 5
- Wu, Wayne (2023) *Movements of the Mind: A theory of attention, intention and action*. Oxford University Press.