

Agency and Observation in Knowledge of One's Own Thinking

Forthcoming in *European Journal of Philosophy*, please cite published version

Casey Doyle  
casey.doyle@st-hildas.ox.ac.uk

St Hilda's College, University of Oxford

*Abstract.* This essay addresses the question how we know our conscious thinking. Conscious thinking typically takes the form of a series of discrete episodes that constitute a complex cognitive activity. We must distinguish the discrete episodes of thinking in which a particular content is represented in phenomenal consciousness and is present “before the mind’s eye” from the extended activities of which these episodes form a part. The extended activities are themselves contentful and we have first-person access to them. But because their content is not represented in phenomenal consciousness, this access cannot be broadly observational. Instead, I argue, it is agential. Furthermore, that extended activities are intentional explains the possibility of a nonobservational form of introspective access to the discrete episodes in consciousness.

*When we form an image of something we are not observing. The coming and going of the pictures is not something that happens to us. We are not surprised by these pictures, saying ‘Look!’ (Contrast with e.g. after-images)*

Wittgenstein 1967: §632

*...when you think a particular thought, there is of course no intention in advance to think that particular thought.*

Peacocke 1999: 209

## 1. Introduction

Much of our conscious thinking is expressed in episodes of inner speech.<sup>1</sup> These are processes of auditory imagining that unfold over time, that occupy one's attention, and that are explicitly *there* before the mind's eye (or ear). As Hulburt et. al. put it, "the speakings are generally apprehended to be in the person's own naturally inflected voice, in the same rhythm, pacing, expressivity, tone, hesitations, and style as external speaking" (2013, 1482). Wittgenstein is surely right that, at least in non-pathological cases, when we become aware of such episodes we are not surprised. Suppose you are in the middle of a department colloquium and notice that your phone isn't in your pocket. You can't go looking for it, so you are left to sit and wonder where it might be. Wondering about the phone is an extended activity that takes place, let's say, over a minute. In the course of this minute various specific thoughts will run through your mind, some explicitly about the phone and others not. These include thoughts like, "I haven't been in the office since this morning, but I had it with me at lunch." You aren't surprised to "hear" yourself think this, any more than you would be to hear yourself say it out loud. It's not surprising because you can tell, straightaway, what you're up to when you're doing it. Ordinarily, you don't have to interpret your thoughts to tell whether you're providing yourself instruction, trying to answer a question, or imagining an interlocutor. That gives us reason to think that our relation to our own thoughts is agential. Yet, on the other hand, Peacocke is surely right to point out that, typically, one does not form a prospective intention to think any particular thought.<sup>2</sup> You may

---

<sup>1</sup> Psychologists estimate that inner speech occupies around 25% of our waking conscious lives (Heavey and Hulburt 2008). I do not assume that all conscious thinking involves inner speech, though I focus on it here as a paradigm case.

<sup>2</sup> One might decide to rehearse a memorized poem in one's head, for example, but this is atypical. And even in this case, it's doubtful that your intention picks out the specific words. You might intend to rehearse "Ozymandias", but not "I met a traveller..." Thanks to Maarten Steenhagen for pointing this out.

not be surprised by your thought, but you didn't exactly see it coming, either. Since it is news to you that you are thinking that particular thought at that particular moment you might suspect that you accessed your thinking by eavesdropping on silent soliloquy, as Ryle put it (Ryle 1949: 148).

The contemporary literature on the epistemology of conscious thinking finds itself pushed in both directions. Most philosophers assume that each of us possesses a uniquely first-personal and epistemically privileged way of knowing our acts of conscious thinking. Call this "first-person knowledge." On what I'll call "the Agency Model", this knowledge is so-called "practical knowledge" or "agent's awareness" (O'Brien 2007, 2013, Roessler 2009, 2013, Soteriou 2013). Many philosophers hold that if an agent is doing A intentionally then she possesses or is in a position to possess nonobservational knowledge of her action. Defenders of this model conceive of our thoughts as actions and accordingly hold that our epistemic relation to them is practical. On an alternative, which I'll call the Observational Model, we know what we are thinking by observing events in the auditory imagination (Byrne 2010, Cassam 2011, Carruthers 2011, Hulburt et. al. 2013). The following from Quassim Cassam is representative:

The sense in which one is aware of inwardly saying to oneself that P is that one 'hears' oneself saying to oneself that P. This is hearing with the mind's ear rather than the ears attached to one's skull.

Cassam 2011: 10

Our intuitions point in both directions, as well. On the one hand, it is intuitive that thinking is something a subject does rather than something that merely happens to her. Accordingly, one might worry that by assimilating the epistemology of thinking to that of pain and other sensations, the Observational Model misses out on a crucial aspect of how we relate to

our own thinking. On the other hand, it is far from clear that the concepts of intentional action, such as choice, decision, or plan, have any place in our conscious thinking. Furthermore, by rooting our knowledge of our thoughts in our intentions or practical thought more generally, the Agency Model fails to make use of something “phenomenologically real” in the epistemology of thought (Cassam *ibid.*). Inner speech is a real phenomenon; we really can “hear” ourselves think. It is peculiar to ignore this in one’s account of our knowledge of our thoughts. It is reasonable to expect that an adequate epistemology of thinking will make sense of both its passive and active dimensions.<sup>3</sup>

My goal in this paper is to propose such an account. In the literature philosophers tend to restrict their focus to inner speech. But conscious thinking is quite a heterogeneous category and extends beyond that case, as our example makes plain. The restricted focus is problematic because it ignores structural features of conscious thinking that are relevant to properly understanding its epistemology. Accordingly, I’ll suggest that getting clear on the epistemology of conscious thinking will require getting clear on its structure.

A complaint against restricting our focus to inner speech is forcefully made by Ryle in the following passage:

It is often supposed by philosophers and psychologists that thinking is saying things to oneself, so that what *le Penseur* is doing on his rock is saying things to himself... Very likely *le Penseur* was just murmuring something under his breath or saying it in his head...

---

<sup>3</sup> The proponent of the Agency Model can of course allow that there is a rich phenomenology of conscious thinking but deny that we gain knowledge of thinking by attending to it. My point is twofold. First, I aim to take seriously the thought that phenomenology does play some kind of role in the epistemology of conscious thinking. Second, intuitively, we don’t have intentions that pick out the specific contents of inner speech. So it looks like the model ignores something epistemologically significant and that the phenomena don’t neatly fit it.

The thin description “murmuring syllables under his breath”, though true, is the thinnest possible description of what he was engaged in. The important question is “But what is the correct and thickest possible description of what *le Penseur* was trying for in murmuring those syllables?”

Ryle 2009: 501

A theme throughout Ryle’s late papers on thinking is that we have to characterize it in these thick terms. Our example is useful here. At any given moment you may be mumbling certain words to yourself, but as should be familiar, an adequate answer to the question, “What are you thinking?” usually cites the extended activity: you are wondering where your phone is. Conscious thinking typically takes the form of a series of discrete episodes that constitute a complex cognitive activity. We must distinguish the discrete episodes of thinking in which a particular content is represented in phenomenal consciousness<sup>2</sup> and is present “before the mind’s eye” from the extended activities of which these episodes form a part. This distinction is important for the epistemology of thinking for the following reason. The extended activities are themselves contentful and we have first-person access to them. But because their content is not represented in phenomenal consciousness, this access cannot be broadly observational. Instead, I argue, it is agential. Furthermore, this agential awareness of the extended activity makes possible a non-observational form of introspective access to the episodes in consciousness. This can explain why we are not surprised by our thoughts although our epistemic relation to them is, in an important sense, passive.

## 2. First-Person Knowledge, Observational Knowledge, and Practical Knowledge

In this section I make some preliminary remarks about the three forms of knowledge considered here. It is a matter of controversy how precisely to understand each of them. I won't defend a particular view of any here. Our task is to consider whether our epistemic relation to conscious thinking is broadly observational or agential. Given this, I assume that we are permitted to operate with a fairly crude version of each model and leave many controversial questions open.

First-personal knowledge of our thoughts has the following three features. First, this knowledge is uniquely first-personal: it is achieved by means of a capacity that is only available to the subject. Second, this self-knowledge is epistemically privileged: my beliefs about my mind are more likely to count as knowledge than another's beliefs about my mind. Third, self-knowledge is epistemically immediate in that it is not based on other beliefs one has, say about one's behavior (Pryor 2005).<sup>4</sup> I will call this knowledge "first-person knowledge." I assume only that we sometimes possess it and that it has the features listed.

Let us use the term "introspection" to refer to the form of epistemic access involved in awareness of the phenomenal character of phenomenally conscious episodes.<sup>5</sup> The Observational Model holds that first-person knowledge of thinking is achieved by introspection in this sense. Introspection can be understood a number of different ways: as acquaintance (Gertler 2001), as

---

<sup>4</sup> That first-person knowledge is immediate is, to my mind, a datum that any adequate theory must respect. However, for present purposes it won't matter if one thinks instead that it intuitively or subjectively strikes us as immediate.

<sup>5</sup> A few clarifications about how I will understand talk of introspection and consciousness. I assume that in introspection one is aware of phenomenal character, or what it's like to be in some conscious state, and that this involves having features presented or appeared to one. Thus, in introspection one is aware of representational contents of experience. That makes it plausible to suppose that if a content is not represented in phenomenology then one cannot become aware of it by introspection alone, a point to which I return below. I assume that introspection is restricted to phenomenal properties. But this claim is stipulative. Given the stipulation, it is intelligible to say, as I will, that there are sources of first-person knowledge other than introspection. Finally, I remain neutral on the question whether conscious states necessarily have phenomenal properties (Pitt 2013). However, later, following Ryle, I will argue that there are occurrent acts of thinking the contents of which are not represented in phenomenal consciousness. Whether these acts count as "conscious" or not strikes me as merely a terminological matter. See section 3 for further discussion.

the interpretation of sensory information (Carruthers 2011), as constitutive and implicit self-awareness (Zahavi 2005), or as mental state detection (Nichols and Stich 2003). The differences between these views won't concern us. They all accept that conscious thinking involves the production of auditory imagery (or perhaps cognitive phenomenology) which is known by means of some form of higher-order receptive awareness.

If there is a useful contrast to be drawn between the two models we are considering, then we must assume, with Anscombe and others, that practical knowledge of one's own intentional actions is non-observational (Anscombe 1957, Velleman 1987, Moran 2004, Setiya 2007, O'Brien 2007). If it is first-person knowledge, in the sense outlined above, then it is also noninferential.<sup>6</sup> However, these are only negative characterizations of practical knowledge. Here, again, I remain neutral on many controversial issues. As I understand it, practical knowledge is "knowledge in intention" in the very weak sense that intending to do A can itself provide a subject with an entitlement to self-ascribe the action of doing A, now or in the future. This is consistent with the possibility that other evidence may defeat the entitlement provided by the intention, such as evidence that one is unlikely to succeed. And it leaves room for different ways of understanding the entitlement and more besides. For example, some philosophers, taking the phrase "knowledge in intention" literally, hold that an intention can itself constitute knowledge either because intention is constituted by belief and subject to the standards of theoretical rationality (Velleman 1987) or because practical knowledge is a *sui generis* form of knowledge subject only to the standards of practical rationality (Anscombe 1957). This contrasts with the view that intending to do A or doing A intentionally can themselves be reasons for the self-ascription of action, though one can intend to do A without possessing practical knowledge

---

<sup>6</sup> Thus, I assume that knowledge of what one is doing is not based on knowledge of what one intends or is trying to do. See Setiya 2007 and Gibbons 2010.

(O'Brien 2007). As I understand it, to endorse the Agential Model is to hold that an intention to think a particular thought is either the basis of one's self-knowledge or itself constitutes that knowledge.<sup>7</sup>

Distinguishing these forms of knowledge allows us to capture our puzzle more clearly. We would like an account of the epistemology of thinking that is compatible with two intuitions: first, that phenomenology plays a role in our knowledge of what we're thinking, and, second, that our relation to our acts of thinking is in some way agential or practical. The first intuition requires that our knowledge of our inner speech involves attention to items in consciousness and that seems to undermine the second intuition. Furthermore, we do not seem to have intentions to think the particular thoughts that run through our minds (a point to which I will return.) While it is intuitive that our relation to our thinking is practical, this doesn't seem possible given other facts about that relation.

We can begin to see our way to an answer of how agential awareness of conscious thinking is possible by distinguishing an episode of thinking from the mode or way it is presented to the subject. Although our access to our acts of inner speech is broadly receptive, we are not surprised by them because we recognize them as our doings. What needs explaining is not how an intention gives rise to specific thought contents, but rather how introspective awareness of them could take a practical or agential form. If not an intention, in virtue of what are acts of conscious thinking agential? My proposal is that discrete acts of inner speech are agential in virtue of the extended activities of which they form a part rather than intentions that

---

<sup>7</sup> I assume that the weaker view is true and thus that you can intend to do A without possessing practical knowledge or belief. Everything I say here is consistent with the stronger view, but it requires some wrinkles about particular cases that I won't consider.

lie behind them. In turn, those activities are intentional in the standard way: the agent has an intention to perform them. I discuss these extended activities in the next section.

### 3. Extended Cognitive Activities

You're wondering where the phone is. That's an activity that unfolds over a minute before you're interrupted or complete it. Over the course of that minute various thoughts come and go in inner speech. There's a part-whole relation here. Inner speech is part of the activity of wondering, just as the act of chopping onions is part of the extended action of making an omelette, or the sentences you utter are part of a talk you give. Reflection on our example is, I think, enough to make this point vivid, but it is worth spelling out in more detail.

Ryle provides the following example:

... I ask a tired visitor from London what he has been thinking about. He says, 'Just about the extraordinary peacefulness of your garden.' If asked, 'Than what do you find it so much more peaceful?' he replies, 'Oh, London, of course.' So in a way he was thinking not only of my garden but of London, though he would not, without special prompting, have said for himself that he had had London in mind at all.

Ryle 2009, 413-14

Ryle's example of an extended activity involves a form of thinking the content of which is not given, or not entirely given, by what is represented in phenomenal consciousness. There is an obvious sense in which the Londoner is thinking of London even though LONDON doesn't show up before the mind's eye. Regarding a related case, thinking-as, Ryle writes:

A person who thinks of something as something is, *ipso facto*, primed to think and do some particular further things; and this particular possible future that his thinking paves the way for needs to be mentioned in the description of the particular content of that thinking... Roughly, a thought comprises what it is incipiently, namely what it is the natural vanguard of. Its burthen embodies its natural or easy sequel.

Ryle 2009, 416

Because he focuses on these cases, Ryle settles on the label “incipient thought”. But this is misleading, since it suggests that what is additional to an episode in inner speech is an episode of the same kind in the near future. But it is clear that the Londoner is thinking of London in or by thinking about how peaceful the garden is and that you are wondering where your phone is by or in telling yourself that you haven’t been to the office since the morning. The additional content is a feature of one’s occurrent thinking, it is just not a part of what is represented in inner speech, or phenomenal consciousness more broadly. Again, Ryle’s discussion is instructive. In a case of incipient thought

It was not that [he] just *forgot* and had to be reminded...but that his task of telling just what he had had in mind was in some important ways totally unlike the task of trying to recall, say, just how many telephone calls he had made during the morning...

ibid. 414

In the same spirit he remarks that a question about the content of extended cognitive activities is “not one that he could settle by racking his brains to recollect a bygone fleeting something” (ibid.)

A few observations about this phenomenon. First, it is not rare or *recherché*. Indeed, Ryle claims that thought, as such, is “constitutionally incipient”, a point to which we will return. Second, one must not be misled by a familiar if cartoonish picture of Ryle as someone hostile to the “inner”. Nowhere in these discussions does he deny the existence of inner speech. His point is, rather, that such occurrences do not exhaust the contents of our conscious thinking. Third, the content of incipient thought needn’t be represented in phenomenal consciousness. The Londoner needn’t have the idea of London before the mind’s eye in order to be thinking of London. Fourth, incipient thought is not an additional thought of the same kind as the underlying episodes in inner speech. The man is not thinking of the garden and London, as it were, side by side, he is thinking of the former *by* or *in* thinking of the latter. Fifth, one should not conclude from the fact that the content is not represented in the stream of consciousness that it is the content of a dispositional thought<sup>8</sup>. As I read him, Ryle’s point is that some of these contents are aspects of one’s occurrent thinking and cannot be captured either as occurring in phenomenal consciousness (“a bygone fleeting something”) or a disposition to think something under other circumstances.<sup>9</sup> Whatever words are before the mind’s eye, you are thinking of the phone, here and now, the Londoner is thinking of London here and now.<sup>10</sup>

---

<sup>8</sup> Thanks to an anonymous referee for raising this issue.

<sup>9</sup> There are features of Ryle’s presentation, most obviously the term “incipient thought”, which suggest the dispositional reading. And it is surely true that in these cases the subject is disposed to think “mutter to himself” various thoughts under other circumstances. But focusing on this obscures an important feature of the phenomenon, namely that both thoughts, that realized now and that which is to come, are aspects of a process underway— in this case, comparing—a process which itself has intentional content.

<sup>10</sup> Are the contents of incipient thought conscious? That depends on whether we think conscious thoughts necessarily involve phenomenal properties. Reflecting on Ryle’s example, I find it natural to say that comparing the garden and London is a conscious activity performed by the Londoner. However, so long as one recognizes that this

A natural question is whether the mode or attitude type (wondering, comparing, counting) is represented in phenomenal consciousness. If it is then the Observation Model might seem to have the resources to explain self-knowledge of these activities.

I am skeptical that there is such a thing as the distinctive phenomenology of the mode of an act of thinking, any more than there is a distinctive *sound* of assertion or interrogation. But I am also far from clear on how to go about establishing a negative phenomenological claim like that in a satisfying way. Doing so would require getting clear on which parallels between inner and outer speech should be taken seriously, a task beyond the remit of this paper. Nevertheless, even if there is such a distinctive phenomenology, it is not up to the task for two reasons. The first reason has been canvassed: a complete epistemology must explain our access to the contents of extended cognitive activities and I have claimed these aren't represented in phenomenal consciousness. There may be a distinctive phenomenology of *comparing* but not *comparing something with London*. Second, we want to explain the practical dimension of our access to our thinking, and following others (Moran 2001, Rödl 2009, Boyle 2011), I am impressed by the thought that the passive awareness involved in introspecting phenomenally conscious episodes cannot do this<sup>11</sup>.

---

activity is *occurrent*, it is neither here nor there whether one calls it conscious. In discussion, some philosophers insist that if the Londoner really is thinking about London here and now, then the content must be represented in phenomenal consciousness somehow or other, perhaps in cognitive phenomenology. I find this description strained and theory driven rather than rooted in reflection on the phenomenology, but am also unsure how to argue against it in a decisive way. I grant, of course, that the fact that the Londoner is comparing the garden with London makes a difference to the character of his conscious experience. For example, various emotions he feels may depend on the fact that he is comparing the garden with London rather than Edinburgh. But that doesn't mean that London must be represented somewhere in the stream of consciousness at every moment while he is comparing the two. Thanks to Katalin Farkas and Laura Gow for pushing me on this.

<sup>11</sup> Later I will propose that there is an active form of introspection, but this is explained by background intentions, rather than phenomenal character.

#### 4. Self-Knowledge of Extended Cognitive Activities

In this section I will argue for four claims about our knowledge of extended cognitive activities.

First, we can possess first-person knowledge of these activities. Your access to these thoughts seems to be made possible by a uniquely first-personal capacity, which yields beliefs that are epistemically privileged. Reflection on examples should make this plain. If, mid-wonder, I ask you what you are doing, you can answer straight away that “I am wondering where my phone is.” You do not need to observe your behavior or rely on any other form of evidence, including attention to your inner speech. You do not notice that the thought “it can’t be in the office” is running through your head and conclude, on that basis, that you must be wondering about the phone.

Second, when in our pretheoretical innocence we ask questions like, “what are you thinking about?” we usually have in mind extended cognitive activities rather than episodes of inner speech. Under ordinary circumstances, it would be inappropriate for you to answer my question by saying “I was thinking that I haven’t been in the office since this morning but I had it with me at lunch.” If it is appropriate, it is not a final answer but only elicits another: “why were you thinking that?” where this question should be heard as analogous to, “why are you chopping onions?” When we want to know what another is thinking we want to know what activity they are engaged in. In Ryle’s terms, we want the thick description.

Third, our discussion provides the resources to respond to skeptical worries about our epistemic access to our thinking. Consider the following passage from an interview with a subject in one of Hulburt’s experiments on introspection. (Hulburt equips subjects with a device that beeps randomly throughout the day and when the thing beeps the subject is supposed to attend to the character of her conscious experience.)

P1: I was saying to myself that I should start making dinner.

I2: What do you mean by “saying to myself”?

P3: It was 5:00 and I was hungry.

I4: I understand that, but our question is about your experience at the moment of the beep. What was in your experience then.

P5: I was thinking that I should make dinner.

I6: Originally you said that you were “saying to myself.” Now you use “thinking that.” Are those the same thing? Either way is OK – our task is to describe the characteristics of the experience, whatever they were.

P7: Ummm. I guess I was saying it to myself.

I8: What exactly do you mean by “saying it to myself?” Did you use words, inflection, and so on?

P9: I don’t remember the words.

Hulburt et. al. 2013: 1481

Now Hulburt is no skeptic about introspection. But he does think that noticing such things as what precisely you are saying in inner speech is an acquired skill that takes his subjects some time to master. Be that as it may, here’s what’s significant about this, apparently wholly representative, passage. The subject is quite unsure about the character of her inner speech, but is in no doubt about what she is thinking. She’s thinking it’s time to make dinner. You might take passages like this to motivate skepticism about self-knowledge (Schwitzgebel 2011), but only skepticism about our ability to introspect the intrinsic features of inner speech. Nothing here should give us pause over our pretheoretical confidence in our ability to know what we are

thinking, so long as we keep in mind the distinction between inner speech and extended cognitive activities.

Fourth, first-person knowledge of extended cognitive activities is not observational. Intuitively, it does not seem that I need to introspect to know that I am wondering about the phone. I don't need to look "inward" at all. Aside from reflection on cases, though, this is the natural conclusion to draw if one takes seriously the idea that the content of extended activities needn't be represented in phenomenal consciousness. It is plausible that one can only introspect those features of her mental activity that are represented in consciousness. If the content of an incipient thinking is not represented in a "bygone fleeting something" then it is not available to introspection. While you may directly observe episodes of inner speech, it would be wrong to say that you have unmediated introspective awareness of the speaking *as part of an unfolding event*.

##### 5. Practical Knowledge of Extended Cognitive Activities

The natural alternative is that our awareness of these activities is practical knowledge, knowledge in intention. The natural source of resistance to this is Peacocke's observation: we rarely plan to engage in such activities. That might make you skeptical that we perform them intentionally at all, which would rule out the Agency Model. In this section I'll argue that we can make sense of a role for intention here and thus that the Agency Model has considerable plausibility. As O'Shaughnessy puts it, though we do not have control over the contents of our thoughts, we do decide on "the content of the governing enterprise" (2000: 89).

An intention sets a standard on performance. The content of the intention specifies how things ought to go with the action.<sup>12</sup> Whereas the content of a belief ought to reflect the facts, one's action ought to reflect the content of one's intention. Where there is a mismatch between one's belief and the facts, the mistake lies with the belief, but where there is a mismatch between one's intention and one's action, the mistake lies with the performance, as Anscombe puts it:

the mistake is not one of judgment but of performance. That is, we do not say: What you said was a mistake, because it was supposed to describe what you did and did not describe it, but: What you did was a mistake, because it was not in accordance with what you said.

Anscombe 1957: 57

One reason to think that extended cognitive activities are intentional is that they are subject to this kind of standard. This is most obvious for cognitive activities that are goal directed, but it applies to others such as contemplating a work of art. At any given moment this activity is realized in a discrete thought and we can evaluate these thoughts against the standard given by the activity. So if you are trying to come up with an answer in your head, then we can evaluate the thoughts that cross your mind (had we access to them) by considering whether they further the goal of arriving at an answer. Likewise, if in the course of wondering where the phone is you get stuck meditating on the color of the phone and regretting your choice to get a phone that color, we can say that you are making a mistake in performance. Wondering about the phone's location is something you can do better or worse.

---

<sup>12</sup> This point is often made by appeal to the idea of "direction of fit" (Velleman 1987), but it needn't (Frost 2014).

Of course, it is not true that every process that is governed by such a standard is thereby intentional. Intentions are personal-level attitudes, the standard set by an intention is one that is set by the subject herself. Since the standard constrains and governs the process, we can say that an intention involves a self-imposed constraint (Soteriou 2014: 264). If extended cognitive activities plausibly involve such constraints, we would have good reason to conceive of them as intentional.

A compelling reason for thinking that the standards governing cognitive activities are self-imposed is that they are attributable to the subject. When you come under criticism, by yourself or another, for deviating from the standard governing one's activity, that criticism is directed at you rather than some subpersonal part of you. And it is attributable to you because the standard you fail by is given by a constraint that you have taken on for reasons. Your thinking should move in the general direction of answering the question or figuring out where the phone is because what you're up to is trying to answer the question or wondering where the phone is. And you are doing these things for reasons that are in a familiar sense your own. You are wondering where the phone is because you don't know where it is and you want it. You are trying to answer the question because you are giving a talk and don't want to look foolish. These are rationalizing explanations of a familiar kind.<sup>13</sup>

Furthermore, like other intentional actions, extended cognitive activities are subject to Anscombe's famous "why?" question, the sense where an answer gives the subject's reasons (1957). In providing an answer to Anscombe's question the subject manifests awareness not only of what she is up to and why, but of the standard by which the activity is constrained. She herself expresses commitment to that standard, which is why the standard is attributable to her. Thus,

---

<sup>13</sup> This may not apply to the case of an isolated thought running through one's head. I discuss this example in section 8.

extended cognitive activities bear many of the marks of intentional actions. It is extremely plausible to conceive of them as intentional actions. That, in turn, makes it plausible to conceive of our epistemic access to them as knowledge in intention.

Still, typically we don't plan extended activities. While I agree with the point, it could only be taken to show that cognitive activities aren't intentional actions if all intentional action requires prospective intention. But it doesn't. Many of our actions, like spontaneously pacing, are manifestations of proximal intentions rather than intentions for the future (Mele 1992: 173).<sup>14</sup> If by "choice" or "decision" we have in mind something like a considered selection from a menu of available options, then, of course, no such thing precedes our thinking, at least in the standard case. But no such thing precedes many of our intentional bodily actions either. Still, in both cases there is an ongoing process that is constrained by a standard which is imposed by the subject herself in light of her desires and goals.

Still, granting this, one might think there are extended activities that don't involve intention. Indeed, you might think that Ryle's own example is a case where an extended cognitive activity is not intentional. Does the Londoner *intend* to compare London and the garden?

It strikes me as plausible to hold that, insofar as it involves extended contemplation, comparing is just the sort of thing you can't really do without intending to. There may be a weak sense of compare on which you can compare two objects unintentionally, for example, merely by recognizing their differences and similarities. But the Londoner is doing quite a bit more than that. He isn't just forming one representation, but unifying a series of them in the pursuit of some goal, whether he reaches it or not.

---

<sup>14</sup> Thanks to an anonymous referee for suggesting this.

One reason to be suspicious of this is that it is possible, perhaps even common, to engage in the kind of activities that these two are engaged in without being aware of it. The Londoner might realize that, perhaps despite himself, he's still got London on the brain. However, that the subject isn't aware of her action while she's in the course of it is compatible with the action's being intentional. I have only assumed a weak view on which an intention provides a subject with an entitlement to self-ascribe action.<sup>15</sup> My claim is not that the Londoner is intending to compare London and the garden, and therefore he must know this, but rather that its being an intentional action of his explains why he is in a position to gain first-person access to it even though the content isn't represented in consciousness.

#### 6. Knowledge in Intention of Inner Speech

After articulating the distinction between episodes of inner speech and extended cognitive activities, I argued that we have first-person knowledge of the latter, that this knowledge isn't observational, and that we have good reason to think that it is practical knowledge. That brings us back to the question of how we know the former. It would be natural to suppose that the Agency Model provides an attractive answer. In this section I'll provide reasons for resisting extending the picture in this way.

One reason to think of extended activities as intentional actions is that the subject who engages in them is subject to Anscombe's "why?" question. But it seems that the same is true of episodes of inner speech. If I tell you that I just thought "it can't be in the office", then you might ask me why I am thinking that. One answer would appeal to the extended activity: I'm thinking it

---

<sup>15</sup> Even those who think that an intention itself can constitute practical knowledge can accept that it needn't always, or they may work with a distinction between implicit and explicit self-knowledge. The Londoner may need to make his knowledge in intention explicit in order to fully appreciate what he is up to.

can't be in the office because I'm wondering where the phone is. If that's right then you might think that the lessons of the last section apply here *mutatis mutandis*.

However, this would be too quick, for it ignores two ways in which something might be said to be done intentionally. In the first case, which we have considered, doing A is intentional when it satisfies an intention. The other way is when something is done in the service of such an intention, in an immediately recognizable way, but where the description of the thing done doesn't satisfy the content of an intention. An example will help. Suppose that I am walking to the river. As I walk my joints and feet move about in various ways. Take one particular subtle movement of my right foot as it secures itself on a rock. Suppose too that I do not attend to this particular movement, it's not deliberate, just an unthinking part of walking. It is perfectly natural to say, of the foot moving, that it was intentional or that I "meant" to do it. That's because the movement of the foot is something done in the service of executing my intention to walk to the river, and it is (or should be) immediately recognizable to me as such. But it would be wrong, I think, to say that I had a further intention which that particular motion satisfied (or failed to satisfy.) I am moving the foot intentionally because I intend to walk to the river, not because I intend that I move my foot in that particular way.<sup>16</sup> It helps here to keep in mind the idea that an intention sets a standard for performance. It should be clear that, at least in the ordinary case, the only standard that bears on the movement of the foot is whether it aids in the achievement of what I intend to do: walk to the river. But in order to do *that* it needn't move in any one particular way, we may suppose. Were the foot to move in a slightly different way, we wouldn't have grounds for claiming there was a failure in performance, unless the subtle change prevented me from walking to the river.

---

<sup>16</sup> See McDowell 2011 and Bratman's rejection of the "simple view" in his 1987.

The suggestion, then, is that inner speech is analogous to fine-grained bodily movements. There's no motivation for postulating additional intentions because there's no reason to think that there's an additional standard for performance. The only standard is the one set by the overarching intention, that which governs the extended activity. Suppose that at the time in question, instead of murmuring to yourself, "It can't be in the office", you murmur, "it must be somewhere else", or "where did I last see it?" Differences here do not seem to involve any deviation from a standard, given that all of these thoughts are recognizably in the service of the governing intention to wonder where the phone is. And when deviations like this are permissible, we should say that there aren't intentions that pick out the specific thoughts, though we may still say that, in a derivative way, the thinking is intentional.<sup>1718</sup>

Wittgenstein remarked that we aren't surprised by what we imagine, a point which applies to the auditory imaginings that comprise inner speech. That's the beginning of wisdom, but only the beginning. I've argued that there's some subtlety to the idea of a lack of surprise here. When a specific thought is uttered in inner speech it is news to the subject. It is not something she planned in advance and it does not satisfy an intention to think that specific thought. Although there is a sense in which inner speaking is intentional, it is not something

---

<sup>17</sup> Extended cognitive activities often follow rational steps, of course, in which case what counts as a success may be more constrained than in our example. But that doesn't mean that we have standards of performance for what one specifically says in inner speech. We can grant what Peacocke calls an "intention to think a thought which stands in a certain relation to other thoughts or contents" (1999: 210), without postulating an intention to think the specific thought.

<sup>18</sup> It has to be said that we must operate with a fairly capacious understanding of what is "in the service of" here. In the midst of wondering about the phone you try to bring to mind an image of it, and this might lead you to think about how you shouldn't have bought a phone that color. I don't suppose this moves you forward in your pursuit, but it may be part of the overarching activity. This doesn't strike me as implausible, though, as many walks have their detours. Cases where the thought seems too far afield are plausibly cases where a separate, perhaps concurrent, activity intervenes. You might get distracted from your mental search for the phone to indulge in some buyer's remorse, for example. I don't suppose, either, that the lines between these will be any more sharp in the mental case than in the physical one.

about which we typically possess knowledge in intention. The lack of surprise must be explained otherwise.

## 7. Non-Observational Knowledge of Inner Speech

We can possess knowledge in intention of extended cognitive activities but not inner speech. Doesn't that simply keep in place the Observational Model, so long as we remind ourselves that it is restricted to certain forms of thinking? I want to suggest a picture that would allow us to answer "no." The picture is an application of the idea that not all perceptual knowledge is observational.

Anscombe argued that one can know what one is doing without observation. Naturally enough, most philosophers have taken this to mean that first-personal knowledge of intentional action isn't perceptual. But others have insisted that we distinguish observation from perception (O'Shaughnessy 1963, Pickard 2004, Gibbons 2010, Schwenkler 2011). One clue here is O'Shaughnessy's claim that it is *impossible* to observe your intentional action while performing it (Ibid.) This would be ridiculous if it meant that one could not watch oneself walk to the river. Instead, the suggestion goes, the form of perception involved here doesn't merit the name "observation" because it isn't passive or receptive. It is, instead, an active form of perceptual awareness. Suppose that you are walking to the river. As you walk you perceive the river, the path, and the motion of your body. This is perceptual knowledge, but it can be practical in two ways. First, plausibly, your intention plays a role in shaping the character of your perceptual experience. It is because you intend to go to the river that what you perceive looks like yourself heading to the river, as opposed to your body simply moving in space (Gibbons 2010). Your ability to recognize yourself as acting is explained, at least in part, by the intention. This

distinguishes your perceptual experience of your action from mine, which means that your way of knowing what you're doing is uniquely first-personal. Second, your perceptual awareness of the action shapes the way in which you act as the action unfolds. It is in virtue of perceiving yourself and your environment that you are able to guide your behavior in a way that conforms to the intention (Roessler 2003, Schwenkler 2011, Shepherd 2016). This knowledge is practical because it plays a causal role in the production of action and is made possible by intention. Call this the "Active Perception" account of practical knowledge.<sup>19</sup>

Defenders of the Active Perception account think that it offers a complete account of first-person knowledge of intentional action. But just as there are aspects of our thinking that can't be observed in inner speech, so too are there descriptions of our actions that go beyond what can be given in perception (Moran 2004). Regardless, the Active Perception account provides us with the resources to explain how we can possess practical awareness of inner speech.

I am walking to the river, but the river is not in sight. I assume that I can have knowledge in intention that I am doing this but not perceptual knowledge. There is good reason to think that practical knowledge is restricted to fairly coarse-grained descriptions of our actions, the kind that figure in practical thought. Knowledge of the movement of one's limbs is likely not to be knowledge in intention, at least in the standard case. But we can still grant that it is practical

---

<sup>19</sup> The Active Perception Account holds that an intention can make a difference to what one can see and the form this perception takes, in much the way that background knowledge informs perceptual experience. I do not assume any particular view about how this works in either case. One view is that background knowledge changes the phenomenal character of one's experience. Another view is that it makes a difference to what one can see, or what one can know by perception alone, even though it doesn't change the phenomenal character. Either view is fine for present purposes. What matters is that other mental states make a difference to what you can see, in the sense of what is available to you as an object of immediate perceptual knowledge. (For further discussion see Pickard 2004). So it would be wrong to construe the background knowledge as a premise in an inference you perform, perhaps unconsciously (Gibbons 2011). Likewise, in line with the view of knowledge in intention sketched earlier, the Active Perception model denies that the fact that one intends to do A serves as an inferential base of one's self-knowledge.

knowledge, in the sense given by the Active Perception account. When I perceive my body moving a particular way, I am not simply observing it, but I also don't know the movement in virtue of intending to make that specific movement. Rather, my intention to walk to the river enables me to have a practical perceptual awareness of the movement. It is immediately perceptually recognizable to me as my action and as in the service of my intention. That is why I am not surprised to see my limbs moving about this way or that, though I lacked knowledge of the movement beforehand.

The same can be said of inner speech. Our awareness of episodes of inner speech is the same, in broad outline, as our awareness of our sensations: we introspect them. But we are not surprised by their contents because this awareness takes a practical form. We recognize what is presented to us in inner speech as episodes of an unfolding cognitive activity which we know through intention. Our intentions render these episodes immediately recognizable as part of the unfolding activity. When you introspect, you become aware of what you're thinking *as an action*, not because you see it against the background of your knowledge of certain facts about yourself (as you might perceive a car as a Buick against the background of some knowledge about cars) but against the background of your intention. While it is news to you that you think that particular thought, you are not a mere spectator of it. That is why you are not surprised.

Our epistemic relation to our inner speech is puzzling because it lies somewhere between the way we relate to our actions and the way we relate to our sensations. The present account allows us to find a home for both the observational and agential dimensions of our relation to our own conscious thinking and thereby succeeds where the two more familiar models fail. Unlike the Observational Model, our view can explain why we are not surprised by our thoughts: we possess practical awareness of them. But unlike the Agential Model, this view finds a place for

the phenomenology of inner speech in its epistemology: our practical awareness is introspective. Furthermore, it does not require holding that we have knowledge in intention of every fine-grained episode in inner speech. Rather, we have knowledge in intention of extended cognitive activities, the thoughts we typically cite when someone asks us, “what are you thinking?”, the thoughts we often know even when we struggle to introspect our inner speech.

## 8. Objections

I will conclude by considering objections to the account.

The account has two components. First, knowledge of extended cognitive activities is practical in the familiar way. These activities execute intentions and the intention entitles the subject to self-ascribe them. Second, unlike extended activities, we do not intend to think particular discrete conscious thoughts. But these episodes, often, and perhaps most of the time, are parts of extended cognitive activities. We lack knowledge in intention of these episodes, instead, this knowledge is broadly perceptual, involving attention to phenomenal character. Nevertheless, this knowledge is not observational but is instead active (quasi)perceptual awareness. Intending to engage in the extended activity enables one to immediately recognize one’s thoughts as one’s own deeds in the way that an intention enables one to immediately recognize fine-grained bodily movement as one’s deeds. We have located a practical dimension to our awareness of our thoughts while respecting the facts that this awareness is based in phenomenal consciousness and it is not based on an intention.

One might worry about the relation between the components. It might seem that, on the account, one must know that one is engaging in an extended cognitive activity in order to gain knowledge of one’s discrete conscious thoughts. This would be problematic for a number of

reasons. First, it doesn't square with the phenomenology. Second, it makes our knowledge of our conscious thoughts look inferential, since it would be based on a prior awareness. Third, there are obvious counterexamples, cases where the subject is unaware of the intention motivating the extended activity yet she is aware of her conscious thinking.

However, the view is not committed to this. A practical awareness of one's conscious thinking need not proceed by way of prior knowledge in intention of the extended activity. Rather, the intention causes one's awareness to take a practical form. One immediately recognizes one's thinking *as one's doing* in virtue of one's holding the intention to perform the activity of which it forms a part. Again, the intention functions like background knowledge in enabling higher-order perception. It is only in virtue of knowing certain things that one is able to recognize a Buick on the lawn. But to immediately recognize the Buick you don't need to reflect upon that knowledge. Similarly, the intention itself enables you to recognize the movements of your body as actions, but you don't need to have a prior awareness of the intention to gain this knowledge. And one is able to recognize the thinking as one's own deed in virtue of holding an intention. The intention itself, not your awareness of it, causes your perceptual awareness to take a practical form.

One might object that this view can only capture cases where inner speech is active or, as Peacocke puts it, "directed" (1999). In addition to directed thinking, such as the cases we have considered, there is undirected thinking, such as when one allows one's thoughts to drift. It's natural to think that we know our directed and undirected thoughts in the same way, but agency, perceptual or otherwise, could only explain one.

But on the current proposal we do use the same capacity in both cases. However exactly we understand introspecting conscious episodes, that is what you do when you know your inner

speech. It just turns out that such awareness can take either an active or passive form, just as perception can. Arguably, this is reflected in the phenomenology of introspection in the two cases. When you know what you're on about in having a certain sentence run through your mind, then you are not surprised by it and recognize it as a part of an unfolding activity. But sudden undirected thoughts are not like that. (The thoughts that run through one's head first thing in the morning or right before sleep might be an example.) This worry would only be apt if one held the Agency Model as a complete epistemology for thinking. But we have good reason for rejecting that.

Still, one might think that the account cannot make sense of isolated thoughts that run through one's head. Suppose you walk into a room and suddenly think, "I'm hot!" This may not be part of something that you were up to ahead of time; it just occurs to you. This seems like a garden-variety thought, one which any epistemology of thinking should explain. Furthermore, it seems like this is something one *does*. It's not like the stray, bizarre ideas that bounce around as you're waking up. It is intelligible and agential. Can the account make sense of such cases?

I agree that the account should make sense of such examples, and I think it can. But the point here is subtle, for I also think that we need to leave room for thoughts toward which we are wholly passive. The existence of passive thoughts that we must observe isn't a counterexample to the account. It would be a counterexample if such thoughts represent a *paradigm* of the thoughts we know and how we know them. But I don't think this is correct.

The idea of a thought that just "comes to one" is common, but a thought that is wholly isolated from one's ongoing intentional activities is a rarer bird than this suggests. We should take seriously Ryle's claim that thought as such is constitutionally incipient. This is easier to see once we recognize that the line between cognitive and physical actions is often blurry; many of

the activities of which a conscious thought is a part are realized in thought and action. It's obvious that we can realize our cognitive actions, in part, physically, such as when you count with your fingers. But mental actions can also be a part of physical actions. Thinking, "I'm hot!" might be a part of noticing yourself sweating, distracting yourself from an unpleasant conversation, ensuring one's guests are comfortable, and so on.<sup>20</sup>

Furthermore, from the fact that a single thought races through one's head it doesn't follow that it is not itself a part of another activity, though that activity may not be allowed to develop. A person who takes a single step before getting struck by a car was, before the strike, taking a step and going for a walk. What matters here is not the realization of a series, but the proper application of a "thick" description that invokes a self-imposed constraint. It is plausible that many isolated thoughts can be described in thicker terms. Quite generally, the way we relate to our thoughts is practical. They are recognizably our doings, even if we do not intend them.<sup>21</sup>

#### Works Cited

Anscombe, G.E.M. 1957. *Intention* (Oxford: Blackwell).

Bayne, T. and Montague, M. (eds.) 2013. *Cognitive Phenomenology*. Oxford: OUP.

Bratman, M. 1987. *Intentions, Plans, and Practical Reason*. Cambridge, MA: Harvard

University Press.

Byrne, A. 2005. "Introspection", *Philosophical Topics*, 33 (1): 79-104.

— 2011. "Knowing that I am Thinking", in A. Hatzymiosis (ed.), *Self-Knowledge*. Oxford: OUP.

---

<sup>20</sup> Again, the point isn't that such a thought *couldn't* be isolated from extended activities. The point is that this is plausibly quite rare and thus these examples aren't paradigms for the epistemology of thought. The intuitions we set out to accommodate plausibly don't concern wholly isolated thoughts, given that they are rare.

<sup>21</sup> Thanks to Ori Beck, Lucy Campbell, Tim Crane, Katalin Farkas, Nicholas Georgalis, Laura Gow, Alexander Greenberg, Maarten Steenhagen, audiences at the University of Braga and University of Cambridge, and an anonymous referee.

- Carruthers, P. 2011. *The Opacity of Mind*. Oxford: Oxford University Press.
- Cassam, Q. 2011. “Knowing what you Believe”, *Proceedings of the Aristotelian Society*, 111(1): 1–23.
- Falvey, K. 2000. “Knowledge in Intention”, *Philosophical Studies*, 99(1): 21–44.
- Frost, K. 2014. “On the Very Idea of Direction of Fit”, *Philosophical Review*, 123(4): 429–484.
- Gertler, B. 2001. “Introspecting Phenomenal States”, *Philosophy and Phenomenological Research*, 63: 305–328.
- Gibbons, J. 2010. “Seeing what you’re Doing”, *Oxford Studies in Epistemology*. Oxford: Oxford University Press.
- Heavey, C. L. and Hurlburt, R. T. 2008. “The Phenomena of Inner Experience”, *Consciousness and Cognition*, 17: 798–810.
- Hurlburt, R.T., Heavey, C, Kelsey, J. 2013. “Toward a Phenomenology of Inner Speaking”. *Consciousness and Cognition*, 22: 1477–1494.
- McDowell, J. 2011. “Some Remarks on Intention in Action.” *The Amherst Lecture in Philosophy* 6: 1–18. <<http://www.amherstlecture.org/mcdowell2011/>>
- Mele, A. 1991. *Springs of Action*. Oxford: OUP.
- Moran, R. 2004. “Anscombe on 'Practical Knowledge'”. In *Agency and Action*, Royal Institute of Philosophy Supplement 55: 43–68. Cambridge University Press.
- Nichols, S. and Stich, S. 2003. *Mindreading*. Oxford: Clarendon Press.
- O'Brien, L. 2007. *Self-Knowing Agents*. Oxford University Press.
- 2013. “Obsessive Thoughts and Inner Voices”, *Philosophical Issues*, 23: 93–108.
- O’Shaughnessy, B. 1963. “Observation and The Will,” *Journal of Philosophy* 60: 367–392.
- 2000. *Consciousness and the World*. Oxford: OUP.

- Peacocke, C. 1999. *Being Known*. Oxford: OUP.
- Pickard, H. 2004. "Knowledge of Action without Observation," *Proceedings of the Aristotelian Society*. 104: 205-230.
- Pitt, D. 2013. "Introspection, Phenomenality and the Availability of Intentional Content", in T. Bayne and M. Montague (eds.) *Cognitive Phenomenology*. Oxford: OUP.
- Pryor, J. 2005. "There is Immediate Justification", in M. Steup and E. Sosa (eds.) *Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Roessler, J. 2003. "Intentional Action and Self-Awareness", In *Agency and Self-Awareness*, N. Eilan and J. Roessler (eds.) Oxford: OUP.
- Ryle, G. 1949 *The Concept of Mind*. London: Hutchinson.
- 2009. *Collected Essays 1929-1968*. London: Routledge.
- Schwenkler, J. 2011. "Perception and Practical Knowledge", *Philosophical Explorations* 14 (2): 137-152.
- Schwitzgebel, E. 2011. *Perplexities of Consciousness*, Oxford: OUP.
- Setiya, K. 2007. *Reasons without Rationalism*. Princeton University Press.
- Shepherd, J. 2016. "Conscious Action/Zombie Action", *Nous*. 50/2: 419-444.
- Soteriou, M. 2013. *The Mind's Construction*. Oxford: OUP.
- Velleman, J.D. 1989. *Practical Reflection*. Princeton, NJ: Princeton University Press.
- Vendler, Z. 1972. *Res Cogitans*. London: Cornell University Press.
- Wittgenstein, L. 1967. *Zettel*. G.E.M. Anscombe (trans.) Oxford: Blackwell.
- Zahavi, D. 2005. *Subjectivity and Selfhood: Investigating the first-person perspective*, Cambridge, MA: The MIT Press.