

## Thinking, inner speech, and self-awareness

JOHANNES ROESSLER

(forthcoming in a special issue of *Review of Philosophy and Psychology*)

### Abstract

This paper has two themes. One is the question of how to understand the relation between inner speech and knowledge of one's own thoughts. My aim here is to probe and challenge the popular neo-Rylean suggestion that we know our own thoughts by 'overhearing our own silent monologues', and to sketch an alternative suggestion, inspired by Ryle's lesser-known discussion of thinking as a 'serial operation'. The second theme is the question whether, as Ryle apparently thought, we need two different accounts of the epistemology of thinking, corresponding to the distinction between thoughts with respect to which we are active vs passive. I suggest we should be skeptical about the assumption that there is a single distinction here. There are a number of interesting ways in which thinking can involve passivity, but they provide no support for a 'bifurcationist' approach to the epistemology of thinking.

A number of authors have recently advocated what might be called a neo-Rylean account of the way we know our own thoughts.<sup>1</sup> The suggestion they elaborate and defend is that we know what we think by 'overhearing', or 'eavesdropping on', 'our own (...) silent monologues.' (Ryle 1949: 176). Yet Ryle left us not one but two accounts of the way we know our thoughts. The passages that have provided the inspiration for neo-Rylean work are to be found in section 5 of chapter VI of *The Concept of Mind*, entitled 'Disclosure by Unstudied Talk'. Ryle pursues an apparently quite different suggestion in the preceding section, entitled 'Self-knowledge without privileged access'. The idea here is that there is a distinctive sense in which a person engaged in a 'serial operation' is 'alive to', and knows, what he is doing, and that a person can be said to know, in that sense, what he is 'at this moment (...) thinking'.

---

<sup>1</sup> Langland-Hassan 2008, Byrne 2011, Cassam 2011, Caruthers, 2012.

(1949: 166, 171) We know our thoughts, on this account (which I'll call Ryle's first account), insofar as thinking can be an example of a 'serial operation'. We know our thoughts, on the second, more familiar account, by 'overhearing' our own inner speech.

This paper has two themes. One is the question of how to understand the relation between inner speech and knowledge of one's own thoughts. My aim here will be to scrutinize and criticize recent answers inspired by Ryle's second account, and to sketch an alternative suggestion, inspired by Ryle's first account. The second theme is the question whether Ryle was right that we need some kind of 'bifurcationist' approach to the epistemology of thinking. On one view, pace Ryle, an account in terms of 'overheard' inner speech provides a complete, unitary explanation of the way we know our own thoughts. (This view is implicit in Byrne 2011.) On another view, a convincing rationale for adopting a 'bifurcationist' approach is provided by the difference between activity and passivity with respect to one's own thoughts: it is specifically thoughts with respect to which we are passive that we know by 'overhearing' our silent monologues. (This view is advocated by Cassam 2011.) I will make a case for a third view, on which, pace Ryle, an account in terms a thinker's 'being alive to what she is doing' provides a complete, unitary explanation of the way we know our own thoughts. I will suggest that we should be skeptical, not about the idea that there are interesting senses in which we may be said to be active or passive in relation to episodes of thinking, but about interpreting that distinction (or those distinctions) in terms that would support a 'bifurcationist' approach.

I begin by giving a more detailed exposition of Ryle's suggestions, and offering reasons to think that we are really dealing with two different accounts, not two formulations of the same account. In section 2 I look at Byrne's elaboration of Ryle's second account. In section 3 I propose a way to develop Ryle's first account. In section 4 I argue that this account is not impugned by the various ways in which thinking can involve passivity. I conclude with some brief remarks on how this discussion bears on the interpretation of 'thought insertion'.

There is a significant omission I'd like to make explicit before I start. I'll assume that thinking frequently involves inner speech (I take it few would deny this), but I won't discuss the question whether it does so invariably. Nor will I consider the puzzling issue of how to account for what seem to be conflicting introspective reports on the prevalence of wordless thoughts.<sup>2</sup> Whether Ryle's first account can be brought to bear on such thoughts, assuming there are convincing examples, is another question I won't be able to address in this paper. In that respect, my case for a 'unitary' explanation of how we know our own thoughts will remain incomplete.

### **1. Ryle's two theories**

In his treatment of 'serial operations' (in section 4 of chapter VI) Ryle offers no account of the basis on which we come to know our current thoughts. The discussion revolves around what might be described as the significance or importance, rather than the basis, of knowledge of what one is doing. The suggestion is this. There are

---

<sup>2</sup> Compare Carruthers's statement 'I can report that most of my thoughts occur in the form of imaged conversations' (1996: 51) and O'Shaughnessy's contention that 'ideas and questions are continually entering one's mind instantaneously and wordlessly' (2000: 247).

certain kinds of tasks the execution of which falls into recognizable (though sometimes only ‘artificially divisible’) steps or stages. One step may relate to another step as a means to an end (laying the table in order to have a meal) or as several parts of a whole (eating consecutive courses of a meal). Carrying out such tasks demands being ‘alive to what one is doing’, where this requires ‘having in mind, in some sense, what is to be done next and what has already been done.’ (168-9) Ryle appears to suggest not only that knowledge of what one is doing is indispensable to the pursuit of complex tasks but also that it represents a special kind of knowledge, or even a special sense of ‘know’. ‘Aliveness’, in this sense, is to be distinguished from the propositional knowledge one expresses when answering the question ‘What are you doing?’. But the two things are not unconnected. I think there are two ways in which Ryle’s account of ‘aliveness’ bears on an explanation of propositional knowledge of one’s current thoughts.

First, Ryle claims that an agent’s ability to (knowledgeably) answer questions as to what she is doing is an integral *part* of her being alive to what she is doing. (1949: 170-1) Suppose we ask: ‘How is it that she knows (has propositional knowledge) that she is multiplying 79 by 45?’ On Ryle’s account, at least a moderately illuminating answer to this can be given by pointing out that she is engaged in a serial operation, hence alive to what she is doing, where this in turn involves the ability to articulate what she is doing, and thus to express knowledge of the fact that she is multiplying 70 by 45. Roughly put, her propositional knowledge is rendered intelligible in terms of her exercising certain interconnected abilities.

Second, there is reason to think that Ryle would not have approved of the demand for what many would regard as the only genuinely illuminating explanation, one that would provide an answer to the question ‘*How does she know she is multiplying 79 by 45?*’ This question is plausibly interpreted as a request for some kind of basis or source or grounds by the use of which she may be said to have acquired her knowledge. There are several passages that imply Ryle would not have regarded this as a sensible demand. He cautions against the idea that announcements of what one is doing reflect ‘any second order performance or process of monitoring the first order performance.’ (170) He also ridicules the idea that serial operations are accompanied by any kind of ‘occurrences’ (‘flashings or dawnings’) in virtue of which we might be said to know what we are doing. (See p. 167)

In the light of this, compare and contrast Ryle’s second account of knowledge of one’s own thoughts:

One of the things often signified by ‘self-consciousness’ is the notice we take of our own unstudied utterances, including our explicit avowals, whether these are spoken aloud, muttered, or said in our heads. We eavesdrop on our own voiced utterances and our own silent monologues. In noticing these we are preparing ourselves to do something new, namely to describe the frames of mind which these utterances disclose. But there is nothing intrinsically proprietary about this activity. I can pay heed to what I overhear you saying as well as to what I overhear myself saying. (1949: 176)

It is not easy to tell exactly what Ryle means by ‘eavesdropping’ and ‘overhearing’. Is he referring to an auditory experience? Or should these expressions be treated as metaphorical? If so, how should they be cashed? Even without attempting to answer these questions, though, we can discern some salient structural differences between the two accounts. Unlike the first account, the second account identifies a source by the use of which we come to know what we think. We are said to ‘notice’ our own unstudied utterances, utterances that, at least to someone who has the relevant recognitional capacities, can ‘disclose’ the speaker’s states of mind. There would thus appear to be a straightforward answer to the question ‘How do we know what we think?’ Whatever ‘eavesdropping’ is supposed to involve, it is by engaging in it that we supposedly come to know our own thoughts. In brief, the second account offers a ‘source-based’ explanation of self-knowledge, in contradistinction to the first account’s ‘ability-based’ explanation.

Another structural difference is this. The second account endorses a certain kind of realism. The facts we ascertain by noticing and describing our unstudied utterances are facts that obtain completely independently of our ability to acquire such knowledge. The capacity for ‘unstudied talk’ is independent of the capacity to pay heed to such talk.<sup>3</sup> In contrast, on Ryle’s analysis, the capacity to engage in ‘serial operations’ is inseparable from the capacity to know what one is doing in being so engaged.<sup>4</sup> This is not to say that the facts in this area are in some sense ‘response-

---

<sup>3</sup> Ryle suggests that we ‘learn to make this study of our own talk from first taking part in the public discussion of anyone’s talk’ (149: 176).

<sup>4</sup> This may need to be qualified in light of Ryle’s remarks about the development of ‘aliveness’ (‘the boy who is just capable of working out a simple sum is probably not yet able to state precisely what he is doing (...)’ (1949: 172)). Still, the capacity for

dependent', or constituted by one's knowledge. Still, the putative interdependence between the relevant abilities amounts to the rejection of a strong form of realism. One's knowledge and the known facts in this area are in an important sense 'made for one another': the latter are not independent of the capacity for the former.

A third structural difference is implicit in Ryle's denial that paying heed to one's utterances is 'intrinsically proprietary'. The capacity to know what someone is thinking by 'overhearing' their unstudied utterances has no special relationship to any particular personal pronoun. In contrast, there does seem to be an intrinsic connection between 'aliveness' to a 'serial operation' and the first-person pronoun. Reflection on the nature of such operations may illuminate knowledge of what I am doing, possibly also, in cases of joint 'serial operations', knowledge of what we are doing; but it can shed no light on our knowledge of activities of which we are not agents. I cannot be 'alive' to an activity, in the relevant sense, unless I myself am engaged in it.

What is the rationale for adopting a 'bifurcationist' approach to the epistemology of thought? Ryle's emphasis throughout is on the contrast between premeditated thoughts and thoughts that are spontaneous and may come as a surprise to the thinker. Properly articulating the contrast, though, is not a straightforward matter. For one thing, as Ryle himself points out, one may be surprised by the conclusion of an argument, even if the conclusion is reached by a well-executed 'serial operation'. (See 1949: 170) When you multiply 79 by 45 you'll presumably expect to reach a conclusion at a certain point, but you may be taken aback when you discover, and find

---

self-knowledge does seem to be part of a *properly developed* capacity for engaging in 'serial operations'.

yourself thinking, ‘equals 3555’. Maybe you thought the result would be a smaller number. But it is in any case not self-evident that (in some sense) ‘premeditated’ thoughts and (in some sense) ‘spontaneous’ thoughts demand different epistemologies. An interesting way to strengthen the case for a ‘bifurcationist’ approach would be to argue that it is the *passivity* of certain kinds of thoughts that demands special treatment. (Cassam 2011) I’ll return to that suggestion in section 4. More immediately, I want to consider the prospects for developing and defending either of the two accounts. I begin with the better-known suggestion that we know our thoughts by ‘overhearing’ our inner speech.

## **2. Acts of inner speech and inner speech acts**

Vygotsky distinguishes two traditions in psychological discussion of inner speech. In one tradition, inner speech is conceived as ‘truncated external speech’, ‘speech minus sound’, or ‘subvocal speech’. According to the other tradition, ‘inner speech differs from vocal speech only as the idea or image of an object differs from the real object.’ (1986: 224) A basic question on which the two traditions disagree is whether ‘saying something in inner speech’ entails actually *saying* anything. The first conception affirms this. Inner speech is real speech, albeit inaudible or incomplete. The second conception denies it. Inner speech is merely imagined speech.

Alex Byrne’s neo-Rylean theory is firmly rooted in the second tradition. His proposal is that we know what we think on the basis of a certain kind of imagistic experience. Inner speech is a matter of ‘auditorily imagining words’. We acquire knowledge of what we are thinking by complying with this rule:



THINK\*\* If you auditorily imagine words that are about x, believe that you are thinking about x. (Byrne 2011: 115)

Byrne puts forward a further rule for gaining knowledge of one's thinking *that p*, but for simplicity I'll focus on THINK\*\*. Byrne makes a number of claims about this rule: that it is a good rule to follow, that true beliefs acquired by its use intelligibly qualify as knowledge, and that our use of the rule explains much of the knowledge we actually have of what we are thinking about. I only want to consider the most basic question raised by these claims, whether the rule satisfies the fundamental condition of being 'truth-conducive', or generating beliefs that are for the most part true.

To be able to apply THINK\*\* one must be in a position to discriminate auditory *imagery* of words from auditory *perceptual awareness* of words. On Byrne's 'representationist' account, the two kinds of experience have much in common. They are have the same type of representational content, in the same sensory modality. But there is also, Byrne suggests (following Hume), a salient difference: imagistic content is 'seriously degraded' compared to perceptual content. It is indeterminate in ways perception is not; for example, there can be indeterminacy as regards the volume of an imagined voice. (2011: 118) Our sensitivity to the poverty vs riches of auditory content helps to avoid confusing our own thoughts with thoughts communicated to us by others.

This is not sufficient, however, to guard against possible confusion. The problem is

that auditorily imagining words is not invariably a matter of engaging in ‘inner speech’. Suppose you imagine the words ‘there is no such thing as society’ as pronounced by Margaret Thatcher. On the face of it, THINK\*\* would enjoin you to form the belief ‘I am thinking about society’. But you are not — you are imagining someone else pronouncing her thoughts about society. (Insofar as you are thinking about anything, it is most likely Margaret Thatcher you are thinking about.) Auditory imagination of words is not coextensive with inner speech.

Byrne does not seem to recognize the threat to the truth-conduciveness of THINK\*\* arising from such cases. However, it may be possible to defuse the problem, by invoking further differentiating features of the phenomenology of inner speech. In a paper cited with approval by Byrne, Peter Langland-Hassan develops just that suggestion. He distinguishes two dimensions along which the experience of inner speech differs from that of auditory speech perception. One is the characteristic poverty, or degradedness, of imagistic content mentioned by Byrne. The second dimension is labeled ‘attenuation’. It is along this second dimension that the experience of imagining *hearing* someone say ‘there is no such thing as society’ differs from the experience of imagining *saying* ‘there is no such thing as society’. The background to this is a theory according to which the perception of intended bodily effects is, as a completely general matter, characteristically ‘attenuated’, or ‘automatically overlooked’, or ‘cancelled’ — in brief, rendered non-salient —, owing to the operation of a sub-personal comparator mechanism filtering ‘reafferent’ signals through efference copies of motor commands. A possible illustration of the difference, Langland-Hassan suggests, may be the familiar difference between the phenomenology of hearing oneself while speaking and that of hearing one’s own

voice on a recorded message. The suggestion, then, is that when we imagine saying things we imagine an ‘attenuated’ auditory experience. In contrast, imagining hearing Margaret Thatcher involves imagining an experience in which the speaker’s voice is characteristically salient. Importantly, the difference in the imagined experiences may plausibly be thought to affect the character of the imagistic experience itself, which in turn should enable us to avoid mistaking imagined ‘observed’ speech with imagined acts of speech. Making this explicit, we might amend THINK\*\* on the following lines:

THINK\*\*\* If you auditorily imagine words that are about x and the imagined auditory experience exhibits attenuation, believe that you are thinking about x.

There is a further hurdle, though, that a successful neo-Rylean theory would have to overcome. THINK\*\*\* would be reliably truth-conducive only if the satisfaction of its antecedent were sufficient to ensure that the belief enjoined in the consequent was correct. But this is not so. Not every act of inner speech amounts to an episode of thinking. The ‘silent recital of a poem known by heart’ — Vygotsky’s illustration of the second traditional conception of inner speech (1986: 224) — is a case in point. Intuitively, internally reciting a war poem does not necessarily involve *thinking* about war. Again, suppose you imagine saying ‘there is no such thing as society’, simply in response to my asking you to do so. In both cases, THINK\*\*\* will produce false positives. It’s not just that you may not think (in the sense of endorsing the view) that there is no such thing as society. The problem is that your imaginative exercise is not a case of what we would ordinarily call *thinking* at all. Neo-Ryleans credit Ryle with

the insight that ‘much of our ordinary thinking is conducted in internal monologue of silent soliloquy’ (1949: 28) But THINK\*\*\* depends on the stronger and less plausible view that the occurrence of inner speech is *sufficient* for thinking.<sup>5</sup>

What has to be added to a ‘mere’ act of inner speech to make it an episode of thinking? A comparison with outer speech is instructive here. Consider Mitchell Green’s illustration of a distinction he draws between ‘acts of speech’ and ‘speech acts’:

(..) testing a microphone, as Ronald Reagan once did in preparing for a news conference, with the words, ‘The bombing of Russia begins in five minutes,’ is not a speech act because, thankfully, there was no question of Reagan’s meaning what he said. He was, for instance, not committing himself to the truth of what he says; nor was he issuing an order. (Green 2013: 14)

Reagan’s act of speech was not serious. It did not amount to asserting (or ordering) that the bombing of Russia was about to begin. It did not amount to a *speech act*, in the technical sense of an illocutionary act. ‘Mere’ acts of inner speech, such as the silent recital of a poem, similarly lack any serious import. It is tempting to think that the notion of a ‘inner’ speech act may shed light on what it means for an act of inner

---

<sup>5</sup> Ryle is very clear on this in his later writings on thinking. Compare this passage from ‘Thinking and Saying’: ‘Then is Thinking just talking to oneself? Or is it doing something Extra? Not the former, since Pythagoras might let his mind wander and just be reciting under his breath random and miscellaneous things like anecdotes, Spanish proverbs, lines from Shakespeare, jingles, and bits of the multiplication-table; and then he would not be thinking.’ (1972: 133)

speech to constitute an episode of thinking. I think this is a promising suggestion, but developing it is a delicate matter. Even if we just focus on the case of ‘inner assertions’ (or inner speech acts that are akin to assertions), at the expense of inner questions, commands etc, we should distinguish at least three importantly different kinds of cases.

First, just as one can make an assertion by saying something out loud to oneself, one may, on the face of it, also do so by saying things ‘under one’s breath’. Inner speech, on these occasions, is simply ‘private’ speech that also happens to be more or less inaudible. (This is in line with Vygotsky’s first tradition.) Whether this naïve picture can be defended depends in part on the account we should give of the intentions informing genuine acts of assertion. For current purposes, I’ll simply assume that the absence of an audience does not disqualify an act of speech from counting as a genuine speech act. Perhaps the idea of an audience is not essential to assertions at all, or perhaps in talking to oneself one can be one’s own audience. (See Green 2007 for illuminating discussion of these matters.) I’ll also assume that there can be acts of speech that are barely audible and perhaps completely inaudible. On these assumptions, talking to oneself ‘under one’s breath’ can be straightforwardly a matter of making assertions or performing other kinds of speech acts.

Second, things are more complicated in cases where ‘inner speech’ is a matter of *imagining* acts of speech. I take it Vygotsky’s second tradition is right that such cases exist and that they form an important part of the extension of what’s ordinarily called ‘inner speech’. Here the thing to bear in mind is that acts of imagination can be

intentional and subject to norms, just as acts of speech. Suppose you try to recall Austin's middle name, and after a while the name comes to mind: you find yourself 'saying in inner speech' — in the mode of *imagining* saying, let's assume — 'Langshaw'. Now while an imagined assertion is not the same as a real assertion, it can in some ways be tantamount to one, provided the imaginative exercise is informed by suitable intentions. Suppose the intention controlling your act of inner speech is to express your knowledge of Austin's middle name. Then the act of imagining saying 'it's Langshaw' incurs the same liabilities as would a real assertion that his name was Langshaw. For example, if the imagined assertion turned out to be false, one would react to one's act of imagination in much the same way as when a real assertion is exposed as erroneous.

From this second case we need to distinguish a third case, which may be described as a thought 'occurring to one', without reflecting any commitment. Much of our ordinary thinking, as the later Ryle emphasized in his work on the nature of thinking, is a matter of *trying out* assertions, or intentions, or questions, with a view to making up one's mind as to what to think, or to do, or to ask. (See Ryle 1972.) Often this takes the form of imagining not monologues, but conversations, involving assertions, pointed questions, objections, and so forth. An act of imagining asserting something, in this context, is not typically intended to express one's attitude. If one does not *have* any settled view yet, there would be nothing to express. One suggestion might be that one imagines speech acts that are less exacting members of what Green calls the assertive family, such as conjecturing or guessing.<sup>6</sup> Another suggestion is that one

---

<sup>6</sup> 'The assertive family is that class of actions in which a speaker undertakes a commitment to the truth of a proposition. Examples are conjectures, assertions,

imagines someone, perhaps an interlocutor in an imagined conversation, asserting something, leaving open (but inviting one to consider) whether one concurs with the view she is articulating. In any case, it is important that acts of imagination of this kind are intended to speak to a question, or to articulate candidates for becoming one's considered view or intention. They are after all intended to be 'exploratory' — to advance the activity of *thinking about* something, e.g. of reaching a view or intention.<sup>7</sup> It is arguably the absence of any such constraint that distinguishes 'mere' acts of inner speech (e.g. recitals of poetry) from 'inner speech acts' (i.e. thoughts).

Put in these terms, the challenge facing neo-Rylean theories is that even if there is some kind of sensory experience that may be said to be characteristic of imagined speech, this will provide no basis for distinguishing thoughts from 'mere' acts of inner speech. In other words, THINK\*\*\* is not reliably 'truth-conducive'. Nor does it seem to be a rule we are actually disposed to follow, or else we'd at least be tempted, in silently reciting poetry, to take ourselves to be thinking about its subject matter. Is it possible to amend THINK\*\*\* in a way that would make it immune to this objection? Ryle only seems to take knowledge of one's thoughts to rest on 'overhearing' internal monologues in cases where inner speech is 'unstudied' or spontaneous. This might suggest that THINK\*\*\* should only be applied on the basis of *finding oneself imagining* words that are about x.<sup>8</sup> But this would not really advance matters. Just as one may recite poetry deliberately, one may also do so spontaneously. And to find oneself imagining saying 'to be or not to be' is not to find

---

presuppositions, presumptions and guesses.' (2013: 6)

<sup>7</sup> See Vendler 1972 for illuminating discussion of the relation between the activity of thinking about something and the occurrence of what he calls mental acts, such as judging or guessing.

<sup>8</sup> This is how Cassam 2011 formulates his neo-Rylean account.

oneself entertaining suicidal thoughts. Whether an act of inner speech, even a spontaneous one, amounts to an episode of *thinking* depends on the intentions informing it. The basic problem with THINK\*\*\* is that it is not sensitive to the thinker's intentions, hence unable to track episodes of thinking.

### 3. Practical knowledge

Ryle's first account — according to which knowledge of what one is thinking can be the distinctive kind of knowledge we have of what we are doing when engaged in a 'serial operation' — has more initial plausibility for some cases than for others. Deliberating whether to buy a particular book or multiplying 79 by 45 are intentional activities. They are promising examples of the sorts of things that Anscombe and Hampshire argued can be known 'without observation' and are open to a distinctive kind of 'practical knowledge' or 'knowledge in intention'. (Anscombe 1957, Hampshire 1965) On the other hand, 'concluding that the product is 3555' is not a description under which any action is intentional, nor is 'being struck by the thought that it's going to rain'. This might lead one to conclude that even if something like Ryle's first account may satisfactorily explain our knowledge of what we are (in the process of) *thinking about*, it needs to be supplemented by a different kind of explanation for knowledge of episodes of *thinking that*, such as concluding that p or it's occurring to one that p. (Trading precision for simplicity I'll refer to such episodes as 'thoughts'.) And it might be argued that the required supplementation in turn opens up the possibility of being introspectively aware of a thought without being aware of oneself as its thinker. For example, Quassim Cassam has argued that by reflecting on thoughts that passively 'occur to us', perhaps suddenly and perhaps for no obvious reason, we can 'begin to see the full force of the Lichtenbergian suggestion that



Descartes should have said 'There is thinking' rather than 'I think'.' (Cassam 2011: 5, n.7)

I want to suggest that this line of thinking underestimates the resources of Ryle's first account, and that it overestimates the epistemological significance of the distinction (or a set of distinctions) between passivity and activity with respect to one's thoughts. Insofar as our thinking is conducted in words, as it presumably is when one is multiplying 79 by 45, reaching a conclusion, such as that the product is 3555, does involve an intentional action, viz. an 'inner speech act'. This, I want to argue, invites extending a 'practical knowledge' account to the case of thoughts (episodes of 'thinking that'). Once this point is properly appreciated, it seems to me, it is far less clear than one might initially have thought that the account cannot also deal with various kinds of cases in which a thinker is naturally said to be passive with respect to an episode of thinking occurring in her mind.

As already implied, I think Ryle's first account may be re-cast, without doing violence to it, in the terms made familiar by Anscombe's and Hampshire's work on practical knowledge. Admittedly Ryle seems more concerned with the phenomenology of deliberate, structured activities than are Anscombe or Hampshire, and less with the nature and the explanatory role of practical reasoning. Yet the serial structure Ryle is interested in may be interpreted as reflecting precisely the structure of the agent's practical reasoning, her recognition of an action as a means to an end or as a part of a whole. This would suggest that the ability at the heart of Ryle's 'ability-based' explanation of self-knowledge is that of reasoning practically and in that way

acquiring intentions, and that it's the intention informing an operation that makes the agent's knowledge of what she is doing intelligible. Connectedly, the reason the challenge 'How do you know you are multiplying 79 by 45?' would be out of place is not just that addressing it would distract the agent's attention from the practical task at hand, but that it encourages her to take up a 'theoretical' orientation on the question of what she is doing — to give an answer based on evidence or observation or some other way of finding out —, and that this would disable her from answering the question in the natural way, by expressing the intention under which she is acting. Whether this notion of practical knowledge can be sustained (whether it is possible to make knowledge of what one is doing intelligible in terms of the agent's intention without offering an account of *how she knows* what she is doing) are large issues I cannot adequately address in this paper.<sup>9</sup> I'll confine myself to the question how far the connection between intention and self-knowledge, assuming it can be made intelligible and defensible, can help to explain knowledge of one's own thoughts.

The main obstacle to extending a 'practical knowledge' account to the case of thoughts seems to be this. Very frequently the descriptions under which we are aware of our thoughts have immediate implications regarding our propositional attitudes, specifically regarding our beliefs and our knowledge. Someone who is multiplying 79 by 45 will, if all goes well, arrive at the thought that the product is 3555. It is natural to describe this in terms of her concluding, or realizing, that the product is 3555, and these are descriptions under which she may be aware of the episode of thinking. But

---

<sup>9</sup> Compare Stroud's claim that '(i)t is *because* the agent has the intention he has, or *because* he is acting under that intention, that he knows that he is intentionally doing such-and-such, but that does not explain *how* he knows.' (Stroud 2013: 8). Further recent discussions of the explanatory connection between intention and self-knowledge include Falvey 2000, Haddock 2011, Soteriou 2013, Roessler 2013a.

concluding that p (at least on one natural reading) entails *believing* that p. Realizing that p entails *knowing* that p. This strongly suggests that neither description is one under which the episode will normally be intentional.<sup>10</sup> The intention to conclude, or to realize, that something is the case would be rather blatantly irrational. And even if someone managed the feat of acting under such intentions, the intention's irrationality would arguably disable it from providing the agent with practical knowledge of what she is doing.<sup>11</sup>

The way to overcome this obstacle, I want to suggest, is to acknowledge that some of the things we know are *grounded in* practical knowledge, even though our knowledge of them is not itself a case of practical knowledge. The event of concluding or realizing that p can be an act of saying, and asserting, that p. Under these latter descriptions, the act can be intentional; and the agent's knowledge of what she is doing, under these descriptions may be explained by her acting under the intention to say, and to assert, that p. This can be so whether she is saying that p in outer or in inner speech. Now the intention to assert that p often goes together with the intention to express one's view that p.<sup>12</sup> If you have practical knowledge of expressing your

---

<sup>10</sup> I think this point is partly responsible for the resistance some philosophers show to the idea of a close connection between thinking and inner speech. Compare Charles Travis's claim that 'one chooses what to say, not what to think.' (quoted in Soteriou 2013: ) and Zeno Vendler's question 'what would it be to want to think something (...)?' (1972: 44). For illuminating discussion of these views, see Soteriou 2013.

<sup>11</sup> It is often rightly said that the 'openness of the progressive' means that knowing that you are crossing the road is compatible with never getting to the other side of the road, and perhaps never even getting started. However, if you form the intention to cross the road in the face of compelling evidence that you won't be able to do so (or perhaps even just despite not knowing that you will be able to do so), arguably you cannot be credited with 'knowledge in intention' that you are crossing the road. (See Roessler 2013a.)

<sup>12</sup> According to some philosophers, that intention needs to be invoked in explaining the nature of assertion. See, for example, Williams 2003.

view that p, then you know that you believe that p, even though believing that p is not something you intentionally do.<sup>13</sup> Suppose, further, that you appreciate that your assertion marks the termination of a knowledge-conducive process of calculation or inference. This would give you adequate grounds for re-describing your assertion in terms of concluding, and realizing or discovering, that p.

On this analysis, we often and perhaps typically know our thoughts under descriptions laden with ascriptions of attitudes. This is unsurprising, given that our interest in our own thoughts is typically an interest in the knowledge they embody, or in the practical and theoretical reasons they make, or appear to make, salient to us. That interest would not be well served just by reflection on the various sorts of inner speech acts one may be in the process of performing. However, while one doesn't intentionally conclude, realize, or have it 'occur to one' that p, knowledge of one's thoughts under these thick descriptions may be *underpinned* by practical knowledge of performing certain inner speech acts.<sup>14</sup> If you are aware of sincerely asserting that p, you know you are expressing *your view* that p, which in turn, and drawing on certain kinds of background knowledge, may enable you to know the event in question under thick doxastic and epistemic descriptions.

---

<sup>13</sup> It might be said that to be able to express your view, you need to have prior, independent knowledge of what your view is (so that you may calculate what to do in order to express it). But this is surely implausible. As Jennifer Hornsby has remarked, '(v)ocicing our thoughts is something that we are able to simply do.' (2005: 120) Intentionally stating one's view that p does not normally reflect an instrumental belief that by vocalizing in a certain way one will be able to state one's view that p, any more than intentionally raising one's arm reflects the belief that by flexing certain muscles one will be able to raise it. As a consequence, practical knowledge of expressing one's view can help *explain* one's knowledge that it's one's view that p. See Roessler (in press) for more detailed discussion and defence of this suggestion.

<sup>14</sup> As indicated in the last section, this may sometimes be, more precisely, a matter of *imagining* performing a speech act, under certain self-imposed constraints. For simplicity, I'll take that qualification as read.

#### 4. Passivity

Suppose that the practical knowledge account can successfully be extended in this way to episodes of ‘thinking that’. There is obviously more that would need to be said here. For one thing, I haven’t considered thoughts that involve questions, commands, expressions of relief or annoyance etc, rather than assertions. But the question I want to consider now is how the account fares with respect to those kinds of thoughts to which, as Harry Frankfurt put it in an influential passage, ‘we are mere passive bystanders’.<sup>15</sup> If there are such thoughts, our awareness of their occurrence is presumably not underpinned by our practical knowledge of engaging in some intentional activity. Rather we should expect there to be a sense in which we are introspectively ‘confronted’ by such thoughts — in a sort of ‘confrontation’ appeal to which would provide a good answer to the question of how we know about the occurrence of these thoughts. Doing justice to the phenomenology of ‘passivity’, it might be said, compels us to construct a ‘source-based’ account of the way we know our thoughts in these cases.

To probe this view, it will be useful to have a longer quote from Frankfurt before us (for ease of reference, I have numbered the kinds of examples he alludes to):

In our intellectual processes, we may be either active or passive. Turning one’s mind in a certain direction, or deliberating systematically about a problem, are

---

<sup>15</sup> Discussions of the epistemology of thought that are substantially influenced by Frankfurt’s claim include Stephens and Graham 2000 and Cassam 2013.

activities in which a person engages. But to some of the thoughts that occur in our minds (..) we are mere passive bystanders. Thus there are (1) obsessional thoughts, whose provenance may be obscure and of which we cannot rid ourselves; (2) thoughts that strike us unexpectedly out of the blue; and (3) thoughts that run willy-nilly through our heads. The thoughts that beset us in these ways do not occur by our own active doing. It is tempting, indeed, to suggest that they are not thoughts that *we think* at all. This would express our sense that, although these thoughts are events in the histories of our own minds, we do not participate actively in their occurrence. (Frankfurt 1988)

The passage has been seized upon by philosophers grappling with the interpretation of ‘thought insertion’, claims made by some patients with schizophrenia to the effect that someone else’s thoughts have been inserted into their minds. An immediate reaction to such statements is that it is totally obscure what patients are getting at. Thoughts, as ordinarily conceived, are not the kinds of things that can be put or placed somewhere, and in any case, for an episode of thinking to take place in my mind just seems to be a matter of me thinking something. Frankfurt’s point suggests this reaction would be precipitate. As Lynn Stephens and George Graham have argued, in the light of Frankfurt’s distinction between being the thinker of a thought and a thought occurring in one’s mind, ‘the subject’s assertion that a thought that occurs in her mind is not her own becomes intelligible. It makes sense conceptually.’ (2000: 152) Note that on this analysis, the distinction needed to make thought insertion intelligible is not particularly exotic. It is not as if psychiatry revealed some unfamiliar structure in the

concept of ownership of a thought.<sup>16</sup> Frankfurt's distinction is supposed to be exemplified by entirely familiar, pedestrian experiences, such as being struck by a thought out of the blue. We supposedly have, in these cases, no sense of ownership of the episode of thinking, as distinct from ownership of the mind in which the thought 'occurs'.

The trouble is that there is actually more than one contrast between activity and passivity in play in Frankfurt's discussion, contrary to the impression created by his unqualified disjunction 'either active or passive'. Thoughts with respect to which we are passive are said to be thoughts that 'do not occur by our own active doing'.

Examples of (2) provide perhaps the most straightforward illustration. Suppose you are suddenly struck by the thought that you left the cooker on. You weren't trying to recall whether you switched it off, nor thinking about the state in which you left the kitchen or any related matter. The thought is not prompted by, and does not form part of, some prior activity in which you were engaged. Compare and contrast the case of thinking 'so the product is 3555' when one is in the process of multiplying 79 by 45.

Thus examples of (2) may plausibly be classified as belonging on the passive side of our intellectual processes if active processes are defined as those that are generated by, and intelligible in terms of, one's activity (i.e. are 'of one's own active doing').

But a process may also naturally be labeled 'active' simply insofar as it *is*, or *involves*, an activity. Frankfurt invokes this latter sense of activity when he suggests that

---

<sup>16</sup> John Campbell takes this view: 'these reports by patients show that there is some structure in our ordinary notion of the ownership of a thought which we might not otherwise have suspected.' (1999: 610)

deliberating systematically about a problem is an activity, hence belongs on the active side of the divide. In other words, there is a risk, not clearly avoided by either Frankfurt or his followers, of equivocating on the term ‘activity’, and of drawing the unwarranted conclusion that thoughts striking us out of the blue are not active in the second sense from the plausible observation that they are not active in the first sense. One might try to amend the argument by adding the premise that only processes that are of one’s own doing are *activities*. But a moment’s reflection shows that this cannot be right. Suppose you notice a finch on the veranda or have your attention drawn to it by someone else. This is an example of engaging in an activity spontaneously — prompted by a stimulus or a co-attender, rather than as a matter of carrying out a prior intention. We may consistently say that while your watching the finch is an activity, your engaging in it is not ‘of your own active doing’.

This suggests that Frankfurt’s blanket active/passive contrast should be abandoned in favour of a more nuanced picture on which there is more than one sense in which we can be said to be active or passive with respect to our thoughts. Being struck by a thought involves passivity in a similar way in which having one’s attention drawn to a stimulus does, but this is not to say that the episode of thinking involves no activity. Arguably it does involve saying things in inner speech and indeed performing ‘inner speech acts’. Thus thoughts striking us out of the blue and thoughts that are intelligible in terms of prior activities have more in common than Frankfurt allows. In both cases we have practical knowledge of the inner speech acts we are performing, and this in turn may ground our awareness of the episode of thinking under richer description, such as concluding that something is the case or being struck by a certain thought. The difference is just that in the unexpected cases we *find ourselves* saying



things, i.e. say them spontaneously. Even so, we may have practical knowledge of what we are doing. What the latter requires is not prior deliberation but merely an intentional action, involving an intention (not necessarily a prior intention) whose content provides some sort of answer to the ‘reason-seeking’ question ‘why am I doing x?’. In the cooker example, this condition is arguably met. You are saying something in inner speech because you are affirming that you left the cooker on.

Consider next Cassam’s definition of passivity. He suggests that what he calls ‘passing thoughts’ are ‘passive in the sense that they are (i) not necessarily responsive to reason, and (ii) states from which one can distance or dissociate oneself.’ (2011: 3)

This sense of ‘passivity’ is best illustrated by examples of (1). One may be thinking about something against one’s will, as in Cassam’s case of a jilted lover finding herself obsessively thinking about her ex. Again, one may have persistent thoughts one takes to be unreasonable, as in Hampshire’s case of someone who can’t help thinking the plane will crash, despite her considered view that this is extremely unlikely (1965: 101-2). It might be said that given the lack of responsiveness to reason of such episodes of thinking, our knowledge of them cannot be grounded in practical knowledge. We may of course ask the subject ‘Why are you thinking about your ex?’ or ‘Why do you keep saying things in inner speech to the effect that the plane will crash?’, but the question would not be a *reason-seeking* question — the kind of question the agent is expected to answer in virtue of her practical knowledge. One might conclude that we need a ‘source-based’ explanation of the subject’s knowledge, invoking some kind of passive ‘confrontation’ with a thought. It is here that we might see ‘the force of the Lichtenbergian suggestion that Descartes should have said ‘There is thinking’ rather than ‘I think’.’

It seems to me that on the contrary, the case of obsessive thoughts brings out the weakness of that suggestion. Obsessive thoughts are characteristically disturbing or upsetting or at least irritating, and this is not just because they (as it were) take up mental space one would prefer to use for other purposes. What makes the experience of having such thoughts unsettling is that it involves an awareness of *oneself* doing certain things: thinking about something or making inner assertions or suggestion or asking questions, even as one's considered, reflective view implies that these activities are pointless or the views they seem to express repulsive. The Lichtenbergian view is not easy to square with the significance obsessive thoughts have for the subject. If you are not aware of such thoughts as a matter of *you* thinking things, the fact that the thoughts fail to match your reflective view would not be very surprising, and no cause for any particular concern. A more plausible alternative would be that obsessional thinking, rather than being completely disconnected from the subject's sense of reasons, reflects a lack of integration in her view of her reasons. Spontaneous thoughts may reflect a sense of one's reasons that does not match one's considered, reflective view. Despite your judgement that the likelihood of a plane crash is negligible, you may still find yourself expressing the view that a crash is imminent for the perceived reason that they happen frequently. On this analysis, obsessional thinking does involve intentional activity on the part of the subject. It's just that the activity in question — thinking about something, or performing inner speech acts — is a defective example of responsiveness to reasons, reflecting the subject's failure to resolve her ambivalent attitudes or feelings about something.

Examples of (3) — 'thoughts that run willy-nilly through our heads' — illustrate two

further senses of passivity. One is provided by the state of reverie, in which, as C.O. Evans put it, ‘the mind wanders from one thought to another in a course dictated largely by the accidental association of ideas.’ (1970: 87) The other kind of example is the case of the thoughts one has just before going to sleep. It is only the second case, I think, that puts any real pressure on the practical knowledge account. For the absence of an overarching project of thinking about something that is characteristic of the state of reverie is compatible with the subject’s activity in (spontaneously yet intentionally) performing various sorts of inner speech acts.<sup>17</sup> Pre-sleep thinking does seem to pose a challenge insofar as, in O’Shaughnessy’s words, ‘(w)hen this phase of the inner life looms up, the intentions have all but petered out (..)’ (O’Shaughnessy 2000: 218) If, along with the intentions, the subject’s ‘knowledge in intention’ has all but petered out, how should we account for one’s ability to know one’s thoughts in this phase of the inner life?

Perhaps we can make sense of the following possibility. An observation that complements O’Shaughnessy’s point about intentions is that the experience of falling asleep involves a ‘receptive attitude to auditory verbal imagery’: ‘we seem to be passive listeners to our own internal monologue’. (Brown 2009: 533) If this is right, then the experience we enjoy as ‘passive listeners’ might provide a *source of*

---

<sup>17</sup> Compare O’Shaughnessy: ‘We all know people who, as one might express it, ‘free-associate in public’, who ‘natter away’, uttering ‘the first thing that comes into their head’, and Molly Bloom’s monologue is nothing but a silent internal example of such. This phenomenon has interesting properties. To begin, it is a case of ‘talking to oneself’, and being talking cannot but be intentionally active. Now these particular intentions stand to one another in a rather special relation. (..) the connective tissue of these rapidly changing intentions is mere association and inclination. As one word-project is approaching its termination, another is already welling up into place (..)’ (O’Shaughnessy 2000: 217)

*knowledge*. We may seem to hear someone say something, and in that way come to know what is being, or at least appears to be, said. In this way, to quote Ryle once more, we may ‘eavesdrop’ on our own silent monologues. Of course, if our attitude is genuinely receptive, our experience will prompt the belief that someone is talking, or at least that it sounds as if someone were talking, rather than the belief that we are thinking or having thoughts. But it might just be possible to have retained a sufficient degree of wakefulness to realize that the ‘overheard’ monologue in fact represents one’s own thinking. In this way, one might conceivably conclude that a certain thought occurred to one, partly on the basis of one’s awareness of a quasi-auditory experience. This would provide a genuine, if somewhat *recherché*, case for adopting a ‘bifurcationist’ approach to knowledge of one’s own thoughts. Note, though, that while the subject could then be said to know what she was thinking on this sort of basis, and indeed could be said to be directly aware of what is in fact an episode of thinking, she would not be *directly* aware of an episode of thinking *as such*. For her recognition of the event as a case of thinking would reflect an inference.

## **5. Thought insertion**

None of the examples of passivity canvassed by Frankfurt, I conclude, gives us any reason to revise the view that in being non-inferentially aware of an event in one’s own mind as a thought one is aware of oneself thinking. If this is right, thought insertion remains as puzzling as ever. There would be no prospect of showing that the delusion ‘makes sense conceptually’ by reflection on familiar varieties of passivity. This result is not all negative, though. It would tell us something significant about the nature of the puzzle we face in trying to understand thought insertion. For one thing, it would deepen our understanding of why thought insertion strikes us as bizarre.

Ordinarily (including in the mildly unusual case of obsessive thoughts), awareness of an episode as one of *thinking* is inextricably bound up with an awareness of *oneself* thinking. This makes it hard to see what patients might have in mind when they insist that the thoughts they experience are not their own.

Elsewhere I have explored the suggestion that there is a way to make sense of thought insertion that simultaneously acknowledges its character as a ‘bizarre’ (i.e. in a sense ‘ununderstandable’) delusion, namely by understanding it in terms of its history. (See Roessler 2013b). Roughly put, we should think of the attitude patients express as a ‘delusional transformation’ of a belief articulated by some patients in the prodromal phase of schizophrenia, that, as one patient put it, ‘when I am thinking in this way (..) it seems to me *as if it is not me who generates*’ these thoughts. (Hesnard 1909; translation and italics by Parnas and Sass 2001: 108). Given that the patient does not, at this stage, *endorse* the claim he uses to convey the utter strangeness of his state of mind, making rational sense of this description is clearly a more manageable task than finding a rational basis for the delusion that others are inserting thoughts into one’s mind. That is not to say that the task is trivial. It is striking, for example, that none of the various forms of passivity reviewed in the previous section would ordinarily lead us to complain that it seemed to us as if we did not generate our thoughts. The basic challenge, it seems to me, is to understand patients’ *concern* with the issue of who is ‘generating’ their thoughts. A familiar move would be to claim that our ordinary awareness of our thoughts delivers a ready answer to the question of who generates them: patients’ experience is unusual insofar as it lacks the normal ‘sense of authorship’. But it is by no means obvious that we ordinarily think of ourselves as ‘generators’ of our thoughts. Our most basic way to relate to our thoughts seems to be

as their *thinkers*. And it seems to be in these terms that we ordinarily find our own episodes of thinking intelligible. Even when you are struck by a thought out of the blue (which might look like a good example of a thought you have not ‘generated’), you are aware of yourself thinking something, and we make sense of such episodes in terms of the reasons we have for performing the ‘inner speech acts’ they involve. You are not normally mystified as to why you suddenly find yourself thinking ‘I left the cooker on’. One good explanation of the thought may be that you did in fact leave it on, and suddenly remembered.

As the longer quote from the patient (reproduced and discussed in my 2013b) makes vivid, the thoughts he is complaining about are not properly intelligible to him in reason-giving terms. He refers to them as ‘associations’. Unlike obsessional thoughts, they are not even intelligible in terms of the thinker’s affective state, such as fear or resentment. The suggestion with which I want to end is that what is unusual in the patient’s description of his state of mind in the prodromal phase of schizophrenia is not just the surprising idea (not — yet — endorsed) that his thoughts might be ‘generated’ by someone other than himself, but the very question of who ‘generates’ them. There are two ways in which the suggestion might be developed. A ‘top-down’ explanation would be that a preoccupation with the possibility of ‘alien control’ is a basic feature of the state of mind characteristic of schizophrenia, and that it structures the way patients experience their own thinking. According to a ‘bottom-up explanation’, the notion that thoughts might be rendered intelligible in terms of someone’s ‘generating’ them may actually *arise from* the experience of being

engaged in episodes of thinking that resist ordinary patterns of intelligibility.<sup>18</sup> The experience may promote a transformation of our ordinary view of thoughts as intelligible inner speech acts into things for which a different style of explanation would be appropriate. A salient feature of that style of explanation is its mechanistic flavour, with thoughts taken to be ‘generated’ or (eventually) ‘transmitted’ or ‘inserted’ into one’s mind.<sup>19</sup>

---

<sup>18</sup> See Campbell 2001 for illuminating discussion of the general distinction between top-down vs bottom-up approaches to the explanation of delusions.

<sup>19</sup> Versions of this paper were presented at the ESPP meeting in Noto in 2014 and at the Society for Philosophical Analysis in Buenos Aires. For criticism and suggestions I’d like to thank Mario De Caro, Naomi Eilan, Thor Grünbaum, Diego Lawler, James Stazicker and Hong-Yu Wong. Special thanks to Sam Wilkinson for extremely helpful comments on the penultimate draft.

## REFERENCES

- Anscombe, E. (1957). *Intention*. Oxford: Blackwell.
- Brown, J. (2009). 'Inner speech: microgenetic concepts', *Aphasiology* 23, 531-43.
- Byrne, A. (2011). 'Knowing that I Am Thinking', in A. Hatzimoysis (ed.), *Self-Knowledge*. Oxford: Oxford University Press.
- Campbell, J. (1999). 'Schizophrenia, the Space of Reasons, and Thinking as a Motor Process'. *The Monist* 82, 609-625.
- (2001). 'Rationality, Meaning, and the Analysis of Delusion', *Philosophy, Psychiatry, & Psychology* 8, 89-100.
- Carruthers, P. (1996). *Language, Thought and Consciousness: An Essay in Philosophical Psychology*. Cambridge: CUP.
- Cassam, Q. (2011). 'Knowing what you believe', *Proceedings of the Aristotelian Society CXI, part 1*, 1-23.
- Evans, C.O. (1970). *The Subject of Consciousness*. London: George Allen & Unwin.
- Falvey, K. 2000. 'Knowledge in Intention', *Philosophical Studies* 99, 21-44.
- Frankfurt, H. (1988). 'Identification and Externality', in his *The Importance of What We Care About: Philosophical Essays*, pp. 58–68. Cambridge: CUP.
- Green, M. (2007). *Self-Expression*. Oxford: Oxford University Press



- 2013. ‘Assertions’, in *Pragmatics of Speech Actions*, Vol. II of the Handbook of Pragmatics, ed. M. Sbisà and K. Turner. (Mouton-de Gruyter)
- Haddock, A. 2011. “’The Knowledge That a Man Has of His Intentional Actions’”, in A. Ford *et al.* (eds.), *Essays on Anscombe’s Intention*, Cambridge, Mass.: Harvard University Press.
- Hampshire, S. (1965). *Freedom of the Individual*. London: Chatto and Windus.
- Hornsby, J. 2005. “Semantic knowledge and practical knowledge”, *Proceedings of the Aristotelian Society, Supplementary Volumes*, 107-145.
- Langland-Hassan, P. (2008). ‘Fractured Phenomenologies: Thought Insertion, Inner Speech, and the Puzzle of Extraneity’. *Mind and Language* 23, 369-401.
- O’Shaughnessy, B. (2000). *Consciousness and the World*. Oxford: Oxford University Press.
- Parnass, J. & Sass, L. (2001). ‘Self, Solipsism, and Schizophrenic Delusion’, *Philosophy, Psychiatry, & Psychology* 8, 101-120.
- Roessler, J. (2013a). ‘The epistemic role of intentions’, *Proceedings of the Aristotelian Society* 113, No. 1pt1, pp. 41-56.
- (2013b). ‘Thought insertion, Self-Awareness, and Rationality’, in K.W.M Fulford *et al.* (eds.), *The Oxford Handbook of Philosophy and Psychiatry*. Oxford: Oxford University Press.
- (in press). ‘Self-Knowledge and Communication’, *Philosophical Explorations* (special issue on Self-Knowledge in Perspective, ed. by F. Jongepier and D. Strijbos).
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson.
- (1972). ‘Thinking and Saying’, *Rice University Studies*, 1, 79-93.

Frankfurt, Harry G. 1998: 'Identification and Externality', in his *The Importance of What We Care About: Philosophical Essays*, pp. 58–68.

Soteriou, M. 2013. *The Mind's Construction*. Oxford: Oxford University Press.

Stephens, G.L. & Graham, G. (2000). *When Self-Consciousness Breaks. Alien Voices and Inserted Thoughts*. Cambridge, Mass.: MIT Press.

Stroud, B. (2013). 'Doing something intentionally and knowing that you are doing it', *Canadian Journal of Philosophy* 43, 1-12.

Vygotsky, L. (1986). *Thought and Language*. (Ed. by A. Kozulin). Cambridge, Mass.: MIT Press.

Vendler, Z (1972). *Res Cogitans*. Ithaca, NY: Cornell University Press.

Williams. B. 2002. *Truth and Truthfulness*. Princeton: Princeton University Press.