Inner speech and the nature of conscious thought

Inner speech and the nature of conscious thought.

"Much of what we call the mental, I believe, is actually things that we do."
(Korsgaard, 2009, p. 38.)

A thesis submitted for the degree of Doctor of Philosophy
at Birkbeck College, University of London

Tom Frankfort

May 2024

Declaration

The work presented in this thesis is my own.

Tom Frankfort. May 2024

Abstract

Not all of what we call thinking, or conscious thought, necessarily involves language, but a lot of it does. This thesis addresses the kind that does, and the nature of the involvement. On the face of it, some instances of this kind of thinking seem to qualify as intentional. For example, when we deliberately and purposefully engage in thinking with a view to reaching a decision or coming to a conclusion. The question is: when we engage in this kind of activity, what do we, flesh and blood creatures that we are, actually and intentionally *do*? My answer to this question is that that we speak to ourselves, either aloud or, more typically, silently; I call this the *thinking-as-speaking* thesis. If valid, it significantly undermines the sharp distinction ordinarily drawn between mental action and bodily action.

The argument for the thinking-as-speaking thesis has several strands, each the subject of a different chapter in this thesis. 1. The capacity we have for inner speech is the internalisation of, and is importantly continuous with, the capacity for overt speech. 2. Speech (both silent and overt) is a species of knowledge-how or practical knowledge. (I apply some ideas from action theory – such as the basic/non-basic distinction – to argue that the act of inner speaking is very often the basic action *by which* a non-basic action is executed.) 3. The nature of thinking-as-speaking has much in common with the nature of other kinds of skilled action: the intention which causes, guides and sustains the inner speech utterance informs the entire process of its production, from conceptualisation to articulation. 4. Crucially, the act of generating and performing the utterance makes a constitutive contribution to the content of the thought being expressed by it. 5. It follows from this that it makes no sense to think of a token thought as something which exists in the head of the thinker before she speaks. Rather, we should understand a token thought as a type of *inner speech act*, the performance of which is the means by which some cognitive goal – such as making a decision, reaching a conclusion, or solving a problem – is achieved. Metaphysically, thoughts are a type of action, not a type of object. 6. If intentions cause, sustain and guide the production of episodes of reasoning, as I claim, then it follows that not

all intentions are formed as a result of conscious deliberation or reflection, on pain of a regress. Some intentions form spontaneously.[1]

[1] Readers might also find it useful to read the Executive Summary at the end of this thesis.

Table of Contents

Chapter 1. Introduction

In a recent collection of essays on the philosophy of Inner Speech, the editors – Langland-Hassan and Vicente (2018) – begin their introduction with a thought experiment.

> In another possible world, not far from our own, inner speech occupies centre stage in contemporary philosophical psychology. Researchers there see the "little voice in the head" as the ultimate theoretical Rubik's Cube: an introspectively familiar phenomenon – more common than a house sparrow – where independent puzzles intersect, and where answering one promises to disrupt solutions to the others (p. 1).

They proceed to sketch a few of the reasons why the researchers in this possible world might find the subject so intriguing. How should we understand the apparently intimate relationship between inner speech and thinking? What should we make of the cognitive nature of inner speech on the one hand – its independence from outside stimulus – and its sensory nature on the other – the fact that, in some sense, we 'hear' an inner voice? If inner speech is where the phenomena of thought, language, and consciousness overlap, what significance does this have for debates about cognitive phenomenology and self-knowledge? If, as the science suggests, the production of inner speech involves the same processes as the production of outer speech, including, for example, those of motor planning and motor control, does that mean that inner speech utterances should be understood as a special kind of (linguistic) motor act? Langland-Hassan and Vicente conclude their thought experiment with this assessment:

> In sum, for the theoretical psychologists and empirically oriented philosophers of mind of this nearby world, there could hardly be a more tantalizing explanatory target than inner speech. It presents them with a set of seemingly intractable, intersecting "big ticket" questions about thought, language, and consciousness… (p. 2).

It is the "tantalizing" nature of inner speech, as evoked by this thought experiment, that inspires and motivates this thesis. I have identified what I consider to be some "big ticket"

questions of my own. By proposing solutions to them I hope to make a modest contribution to closing the gap between the actual world and the possible world of Langland-Hassan and Vicente's thought experiment.

The thesis has three parts. In Part I is designed to set the scene and provide some scientific background to the philosophy that follows. I make no apology for this. Philosophers of mind are not the only people interested in the phenomenon of inner speech, and a huge literature exists on the subject – from linguists, psychologists, psycholinguists, neuroscientists, and cognitive scientists. It would be not merely negligent, but perverse, to ignore that literature when examining the philosophical questions arising from the phenomenon. Chapter 2 introduces the phenomenon of inner speech, focusing on the argument that it is an internalised version of outer speech, involving most of the same language production and motor-sensory processes. In chapter 3 I deal with a question which must be addressed by anyone who theorises about the relationship between language and thought: the fact that some people experience thinking – even conceptual thinking – without experiencing words. This chapter ends by raising one of the deep questions of this thesis, one which is finally addressed directly in Chapter 6: does the action of speaking play a constitutive role in determining the content of our thoughts?

Part II contains the main philosophical meat of this thesis. It comprises four chapters, each addressing a particular "big ticket" question raised by the phenomenon of inner speech. Although each chapter was originally written as a stand-alone essay,[2] each one develops, in some way, the ideas of the previous one, so the essays lend themselves to being re-cast into the chapters of this thesis. Chapter 4 addresses an apparently simple question: is inner speech something we do, or something that happens to us? (Is it active or passive? Is it intentional or not intentional?) Introspection suggests that it can be both (although not at the same time) but this cries out for explanation. Sometimes words and phrases just pop into our heads – they form a kind of spontaneous stream of consciousness which requires no effort on our part, and which doesn't have any obvious purpose. At other times – for

[2] One of these essays has been published, one soon will be. See Frankfort (2022) and Frankfort (Forthcoming. Published online: 12 Jan 2024).

example when we are consciously thinking about what to do or what to believe – the words we silently utter to ourselves seem deliberate, motivated and purposeful. How should we understand what, in Chapter 4, I call "the two voices of inner speech"?

My answer to this question involves an appeal to a well-known idea from the philosophy of action, and one which recurs throughout the thesis. The idea is the distinction many philosophers of action make between a basic action and a non-basic action. My proposal, perhaps surprisingly, is that when we intentionally engage in an episode of (silent) deliberation (trying to work out what to think about something, or what to do) we are engaging in a non-basic (cognitive or at least quasi-cognitive) action – an action we perform by doing something else. And that something else, I argue – the basic action *by which* we perform this non-basic one – is to make one or more inner speech utterances. This proposal has the virtue of allowing us to construe episodes of deliberation as intentional, while at the same time granting that the individual inner speech utterances which comprise the episodes are not themselves motivated by independent intentions. This saves us from a possible regression. When I silently say "p" to myself during an episode of deliberation, there is no other thing I intentionally do, and do it by saying "p".

The argument in Chapter 4 is subject to an objection which comes from introspection: When we engage in deliberation it very often takes the form of a sort of inner conversation, comprising questions and answers. (The dialogical nature of inner speech has been much debated since Plato first described thinking as a silent conversation the soul has with itself.)[3] The objection is that the questions (on the one hand) and the answers (on the other hand) seem, phenomenologically, to have a different status vis-à-vis their being active, as opposed to passive. Intuitively, the questions we *ask* ourselves when we deliberate are active – our asking them seems nearly as much like an intentional action as when we ask a question aloud of someone else. By contrast, the experience of *answering* our own questions seems

[3] "[Thinking is] the talk which the soul has with itself about any subjects which it considers…. [T]he soul…when it thinks, is merely conversing with itself, asking itself questions and answering, affirming and denying…" (Plato, 1921). I speculate that Davidson (1973) might have had this famous quote in mind when he wrote, "Languages we will not think of as separable from souls; speaking a language is not a trait a man can lose while retaining the power of thought" (p. 7).

to be different – the answers just seem to come to us. (Sometimes, of course, they don't. We try to decide what to choose from the menu and we just can't. "What are *you* having?" we ask our companions, in another attempt to get the answer to come.) Frankish (2018), invoking 'dual process' theories of cognitive processing, argues that while the questions we ask ourselves are intentional ('Type 2' processing), the answers are not (he says they are 'Type 1' processing), and this leads him to conclude that an episode of reasoning is only *partially* intentional. In Chapter 5 I argue that this conclusion is wrong, and the result of not appreciating the basic/non-basic distinction I argue for in Chapter 4, as it applies to inner speech utterances. I argue that *all* the inner speech utterances involved in reasoning, both the questions and the answers, are basic, and all qualify as intentional for the same reason: in virtue of the role they play in helping bring about the non-basic action i.e., the goal of the episode of reasoning. I therefore conclude that practical reasoning is, *pace* Frankish, *wholly* intentional. But this leaves a puzzle. How should we explain the aforementioned observation from introspection – the fact that the questions *seem* more active in nature than the answers? I end this chapter by proposing a novel account of the nature of the inner speech dialogue involved in reasoning, according to which its function and its phenomenology are characteristically *exploratory*. I argue that the phenomenology of exploratory actions encourages the *impression* that some of the inner speech utterances involved in reasoning are not genuine actions, when in fact they are.

This thesis is entitled 'Inner speech and the nature of conscious thought', but in this introductory chapter, so far, I have said a lot about inner speech but not much about conscious thought. I want to say something briefly about that now. The English language is notoriously quirky, but in the case of the word 'thought' it has outdone itself. On the face of it, the word 'thought' can be used to refer to three different kinds of entity.

A) A kind of action. As in, "I thought about it, and thinking it made me sad". When used this way we might say the word refers to a token *thinking*.
B) A kind of object. As in, "Once the idea was in my head, I couldn't get rid of it". When used in this way, the word refers to an object which represents something.
C) A kind of content (such as what the object in the head represents). As in, "My thought was the same as yours".

It is even possible for the word to be used in all these different ways in the same sentence, as in, "I thought a thought, and my thought was that …" (typically followed by a proposition). (Making such an utterance would be odd, but it is not ungrammatical, as far as I can see.)

I only want to make two points about this at this stage. The first is that in this thesis I am mainly interested in the entity implied by the first kind of usage. The question I want to answer is, (A) What is the nature of 'a conscious thought' when 'a thought' is understood as 'a token thinking'? The second point is that, as I will argue later, the way we decide to answer *that* question has implications for how we should answer two others. How should we understand the nature of 'a token thought' when it is used (B) to refer to an object (something which represents something) or (C) the content of what is represented by that object. Indeed, the moral of the story is that we should properly understand the nature of thinking a thought (as an action) *before* we decide on the nature of a thought (as a noun or as content). I will return to this subject at the end of Chapter 6.

The first 5 chapters set the scene for perhaps the most important chapter in the thesis. In Chapter 6 I explore the key question: If inner speech plays the key role in episodes of thinking which I claim it plays, what precisely is the relationship between the two? More specifically: are speaking and thinking (at least sometimes) *the same thing*? I call this the *thinking-as speaking* thesis. In this chapter I make a novel argument for this thesis based on speech act theory. (Interestingly, speech act theory, as far as I can tell, has absolutely nothing to say about *inner* speech. By contrast, many inner speech theorists simply *take it for granted* that at least some inner speech utterances function as speech acts. There is, therefore, a curious disconnect in the philosophical literature which this chapter reveals.) Speech act theory has it that overt utterances are in some sense *performative* – utterances *do* something, as well as *say* something. I argue that *inner* speech utterances do things too: inner speech utterances are (very often) inner speech *acts*. If that's right, then we should understand the full *meaning* of an inner speech act as an amalgam of two things: a) what the speaker intends to get done by performing it and b) the semantic content of what is uttered. Those things, I argue, only come together with the production and performance of an utterance, which means there is no coherent sense in which *the* thought being expressed

by the utterance – what the speaker means by uttering it in the fullest sense – can have existed before the utterance was made. On my account, making the (silent) utterance (i.e. performing the speech act) is the same as making the thought. Individual thoughts are not objects (such as mental states or sentences in a language of thought) – they are a type of action. I illustrate this argument with the action of deciding something, which takes us back to a question which was raised, but not answered, at the end of Chapter 5. In that chapter I argued that an episode of reasoning is a wholly intentional activity, comprising inner speech questions and answers, but I did not address the event that occurs at the conclusion of such an episode: coming to a decision or reaching a conclusion. In Chapter 6 I show how that is done with the performance of an inner speech act, using the action of deciding something as a case study.

Chapter 7 articulates an objection to the thinking-as-speaking thesis argued for in Chapter 6, and then responds to it. The objection is based on what is sometimes called the *hybrid view* of skilled action – the view that skills have two components: a cognitive component, and a motor component. On this view the cognitive component involves propositional knowledge, and the motor component involves automatic, low-level causal processes, acquired through brute repetition. Also on this view, all the *intelligence* associated with the skill is associated with the cognitive component, and the motor component is dumb – little more than reflex. The hybrid view can be used to pose an objection to the thinking-as-speaking thesis, by arguing that *thinking*, properly understood, refers to the cognitive component of this activity, while speaking refers to the motor component; and all the intelligence involved is supplied by the cognitive component. I offer two arguments against this objection. The first is that the hybrid view is wrong in the case of skilled *bodily* actions. Skilled bodily actions show none of the characteristics one would expect them to show if the strict hybrid view was right. The second is to pose the objector a dilemma. If the objection is confined to the last stage of the language production process, then it misses its mark, since by then the language production process has produced a meaningful utterance and made its constitutive contribution to the meaning of that utterance. If, on the other hand, the objection extends to the whole of the language production process, the implication is that all the meaning of a natural language utterance is entirely determined before the

language production process has even begun. For various reasons, this is tremendously implausible.

The claim that a token thought is a kind of action is a substantial metaphysical claim, and it deserves more attention than it receives in Chapter 6. In Chapter 8 I buttress the claim by arguing for it in a novel way. Until now my method in this thesis has been to start with the phenomenon of inner speech, and then invoke action theory and speech act theory to argue that token thoughts should be understood a type of action – specifically the action of performing an inner speech act. In Chapter 8 I subject that idea to a metaphysical test. I start with a well-known metaphysical theory concerning the ontology of action and proceed to demonstrate how well it accommodates my claim that thoughts are a type of action. The metaphysical theory is Kit Fine's (especially 1982 and 2022), who argues for what he calls *qua objects*. He originally develops the idea of a qua object in the context of a puzzle about certain material objects, then extends his theory to cover certain bodily actions. I extend the theory further to certain so-called mental actions, specifically thoughts. (Given the lengths I have gone to argue that thoughts are actions and not objects, it is awkward for me to have to adopt Fine's term 'qua object' to refer to thoughts – as Evnine (2016) observes, the term *qua event* would be much more appropriate. However, I will stick with Fine's terminology to keep things simple.) Whether Fine's theory of qua objects is sound is, of course, contentious. I invite those who are sympathetic to it to read this chapter as grist to Fine's mill – another way in which Fine's theory can do useful work. I invite those unsympathetic to it to read it conditionally: were Fine's theory to be sound, it would confer metaphysical credibility to my argument that thoughts are a kind of action.

I end Chapter 8 by noting that my account of a token thought as a species of intentional action is at odds with much of the literature on intention. When the term is used in the literature, the word 'intention' typically refers to a decision or a commitment that has been arrived at *following* some form of deliberation or reflection. But clearly this conception doesn't sit comfortably with my account. Assuming an episode of practical reasoning is constituted by a series of distinct, albeit connected thoughts, and if, as I argue, thoughts are intentional actions, then an intention can't *always* be the product of practical reasoning, on pain of an infinite regress. I argue that at the onset of *any* action, mental or bodily, there is,

ultimately, a mental state which represents an intention to do something which was not itself the product or output of reasoning.

Chapter 9 is dedicated to commenting on an alternative account of the ontology of conscious thought, that of Matt Soteriou (2013). Soteriou's account is chosen for two reasons. First, it is one of the most sophisticated and comprehensive of recent years, and as a rival to my own I feel compelled to respond to it. Second, his account, like mine, takes very seriously the role speech plays in conscious thought, and yet comes to very different conclusions. One goal of this chapter is to understand why.

I want to conclude this introduction with some comments about the bigger picture. I said at the start of this introduction that the thesis addresses some "big ticket" questions arising from the phenomenon of inner speech. In their different ways, the questions explore different aspects of the same question: if thinking is an intentional activity, as we intuitively feel it is, then how do we – flesh and blood creatures that we are – actually do it? My answer to this question, in a nutshell, is that we speak, albeit silently. I emphasise that, so understood, the activity of thinking has a much more intimate connection with bodily activity than we might at first have thought.[4] This conclusion, if it's right, doesn't stand in isolation; it fits into a revisionary trend. When a person engages in an episode of practical or theoretical reasoning, most philosophers assume that whatever might be going on in the brain, it has nothing whatsoever to do with anything as biologically gross as the motor planning system, the musculature of their embouchure, or the pattern of their breathing. The human capacity for abstract thought is, in this sense, the last bastion of a certain kind of dualism. Most philosophers assume that thinking *supervenes* on the brain; this thesis argues that thinking is a lot more embodied than that.

Another philosophical distinction, almost universally observed, which is close cousin to the distinction between mind and body, is the distinction between the so-called 'personal' and 'sub-personal' levels of explanation of human behaviour. The legitimacy of this distinction,

[4] Quine (1985) takes the idea in his stride: "When we deliberately and effortfully think, presumably muscles come into play" (p. 88). My thanks to Keith Hossack for alerting me to Quine's insightful presumption.

or at least the use to which it is put, has recently come under attack: see for example Bermudez (2000) and Rupert (Manuscript)[5]. It would take me too far from the purpose of this thesis to expand on the subject here, or to explore the ways their work and mine are complimentary. Future work, perhaps. But I will touch on one implication of their work. Maintaining the distinction between personal and sub-personal explanations has played an important part in the attempts by some philosophers to find a place in their ontology for an autonomous human agent. To put it crudely (and maybe unfairly), they employ the distinction to argue that explanations of human behaviour can be causal, without being governed by causal laws.[6] Some philosophers (me included) find this claim puzzling on its face, but in any case, without the help of an ontological distinction between personal and sub-personal explanation it is hard to sustain. It seems to me that my arguments for understanding thinking as a special kind of embodied action push in the same direction as those who are sceptical about the personal/sub-personal distinction.

The nature of human agency is not, directly, the subject of this thesis (although the issue is always lurking in the background), but if what I have argued for here is on the right tracks, what I say has implications for how we should understand it. In his paper *What is 'Mental Action'?* (2019) Yair Levy argues that there is no principled way to mark the distinction between mental acts and bodily acts. One of the 'big picture' aims of this thesis is to contribute a new line of argument to support Levy's conclusion:

> It may certainly be important to learn and understand why the distinction [between mental and bodily acts] does not withstand scrutiny, however the point remains that conceiving of human agency, as manifested in essentially two different domains – the mental and the bodily – is a wrongheaded paradigm that should be discarded" (p. 989).

[5] These theorists do not to deny that explanations can and do legitimately involve both levels. They deny that the levels are ontologically or epistemically isolated from each other in the special way that supporters of so-called "agent causation" need them to be.

[6] For example, Steward (2012b). For an excellent critique of what O'Brian refers to as Steward's particular brand of "dualism" see O'Brian (2024).

Part I: The phenomenon of inner speech

In part I of the thesis I set the scene for the philosophy that follows in parts II and III. Chapter 2 is mainly about the latest scientific and psycho-linguistic research into the phenomenon of inner speech. It offers evidence in support of the motor-sensory view of the phenomenon, and it uses the predictive control account of motor planning to explain how it is that we get to 'hear' our inner voice at all. Chapter 3 address the inconvenient fact that some people experience abstract/conceptual thought without experiencing language. I explain how this phenomenon can be accommodated within the motor-sensory view of inner speech described in chapter 2.

Chapter 2. Inner speech: what are we talking about?

Section 1. Introduction

The phenomenon which is the subject of this thesis goes by a number of different names: inner language, inner speech, inner voice, covert speech, internal speech, silent speech, self-talk, internal monologue, internal dialogue, imagined speech, private speech, verbal thought, subvocalization, auditory imagery. Many of us talk to ourselves, silently; we have 'a little voice in our heads'. The phenomenon is common, although how common is still hotly debated; some people appear not to experience it at all. (This fact is clearly inconvenient to anyone who wants to investigate the relationship between inner speech and conscious thought, since those who don't experience inner speech undoubtedly engage in conscious thought. Addressing that problem is the subject of the next chapter.)

An obvious place to start when thinking about the phenomenon of inner speech is the phenomenon of outer speech (also 'covert speech' or 'speaking aloud' – I will use these terms interchangeably in what follows.) After all, it would be just remarkable if the two phenomena were not deeply connected. In practice nearly every writer on the subject agrees that they are, whether those writers are linguists, psychologists, neuroscientists, cognitive scientists, or philosophers. Many of them subscribe, in some form or another, to the ideas of Lev Vygotsky, a Russian/Soviet psychologist, who Langland-Hassan and Vicente

(2018) describe as “one of the most influential historical authors on the topic of inner speech” (p. 20). Vygotsky’s most well-known and influential thesis is that inner speech is an *internalised version* of outer speech (1987). He observed the way children acquire their linguistic skills and noted the systematic way the process developed. To begin with a child learns language by engaging with her caregivers, typically when she is engaged in some task or play, and through these linguistic exchanges she learns the social functions of verbal interaction; she learns that language can be used to instruct, to encourage, to admonish, to insist, to question, and so on. Vygotsky emphasised the *regulatory* function of these exchanges between caregiver and child. Next, she learns to internalise this behaviour in two steps, first by speaking aloud to herself (i.e., when no one else is present), and then by speaking silently to herself i.e., using inner speech. In both cases, the child uses language in the same conversational way she has learned to use it when using it socially, to speak with others. Another of Vygotsky’s well-known theses is that inner speech is an abbreviated or condensed form of speech – our inner speech frequently does not comprise syntactically complete natural language expressions, but rather incomplete phrases and expressions. More controversially, he thought that this process of condensation can reach the point where words are no longer used and the subject inner speaks in “pure meanings” (1987, p. 247). (In the next chapter, when I address the phenomenon of so-called ‘unsymbolised thought’ (UT) we shall see that some contemporary philosophers have found a way of accommodating the idea of inner speaking in “pure meanings” in a way that is continuous with the more common experience of inner speaking in words (Vicente & Jorba, 2019).) A third important Vygotskian idea is that, even when it takes a condensed form, inner speech retains its dialogical character. Since inner speech is the result of the internalisation of the social linguistic practices we learn as children, inner speech takes the form of an internal conversation with ourselves. (This idea will receive special attention in Chapter 5, when the question-and-answer nature of silent reasoning is discussed.) In fact, all three of Vygotsky’s most influential ideas will play a role in subsequent chapters of this thesis: that inner speech results from the internalisation of outer speech; that inner speech is condensed; and that inner speech is dialogic.

The rest of this chapter is not philosophy, but science, and for that I make no apology. As already noted, philosophers of mind are not the only people interested in inner speech, and

a huge scientific and psycho-linguistic literature exists on the subject; it would be perverse, and not merely negligent, to ignore that literature when examining the philosophical questions arising from the phenomenon. But I have another reason for wanting to include something on the science of inner speech which is less principled and more self-serving: the scientific consensus on one particular issue makes the arguments of this thesis much more plausible than they otherwise would be – or so I will argue in the chapters that follow. The issue in question is the rivalry between two views of the nature of inner speech: the Abstract versus the Motor-Sensory view. Roughly speaking, the Abstract view (e.g., Oppenheim & Dell, 2010) is that inner speech involves symbolic and abstract representations, divorced from bodily experience; the Motor-Sensory view is that inner speech is embodied, and involves physical processes that unfold over time. (Note: this distinction has nothing to do with the *content* of the inner speech utterances, but only the *format* of the representation – whether the representation is abstract and symbolic or concrete and physical.) As previewed in the introductory chapter, I will argue in this thesis that an inner speech utterance should be understood, at least in some cases, as a basic action in the service of a non-basic action, in much the same way as my hand going up is the basic action in the service of the non-basic action of, for example, attracting the waiter. Intuitively, this comparison is much more plausible if inner speech utterances are, like bodily movements, motor-sensory in nature, and happily for me, as we shall see in the rest of this chapter, the latest scientific consensus favours the motor-sensory view. It is also, as it happens, the view of many philosophers (but not all) who write about inner speech. Having said that, I want to make it clear that, while my arguments about the nature of conscious thought as developed in this thesis are more *plausible* if the motor-sensory view is right, they are not hostage to it. I will briefly revisit this claim in Chapter 6, Section 7, and spell out how my arguments concerning the nature of conscious thought are consistent with the abstract view of inner speech.

Section 2. Speech production

Before I offer a summary of the sensory-motor view of inner speech I want to take a moment to look at what is involved when a person moves from being in a state of intending to say something aloud (for whatever reason) and actually saying it. Exactly how this works

is still the subject of much debate and research. Since none of the steps in the process are available to us introspectively, scientists infer them based on what can be observed when subjects make mistakes, or when they perform verbal tasks designed to reveal the steps involved. Nevertheless, there is broad agreement that the process is something like that illustrated in Figure 1.

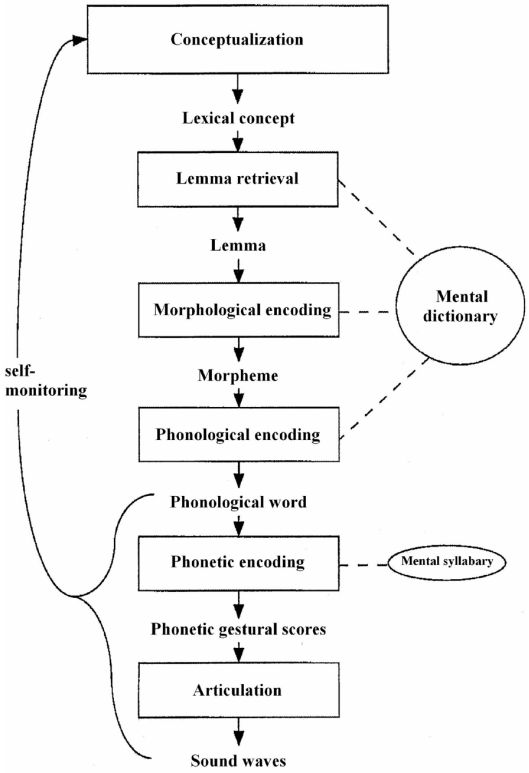

Figure 1. The theory of speech production in outline. From Levelt *et al* (1999).

Here is a summary of what each stage of the process represents. (Philosophers will immediately want to query the meaning of some of the terms used in these definitions, such as "conceptualization" or "representation"; they must be patient – these are the terms as used by the psycho-linguists who subscribe to this theory. I will address philosophical concerns about their meaning in due course):

Conceptualization. This step involves converting one type of representation (a communicative intention) into another kind of representation (a preverbal message). This

stage is so significant for what follows that I will expand on it now more than on the other steps in the process. The full name of this stage is 'conceptual preparation in terms of lexical concepts'. At this stage the concepts activated by the subject's communicative intention are converted into items for which she has words – lexical concepts. But this is far from being a matter of simple 'translation', since there are many ways the communicative intention can be realised.

> A major issue, therefore, is how the speaker gets from the [pre-linguistic] notion/information to be expressed to a message that consists of lexical concepts (here *message* is the technical term for the conceptual structure that is ultimately going to be formulated). This is called the *verbalization problem,* and there is no simple one-to-one mapping of notions-to-be-expressed onto messages (Bierwisch & Schreuder 1992). Even if a single lexical concept is formulated, as is usually the case in object naming, this indeterminacy still holds, because there are multiple ways to refer to the same object (Levelt *et al*, 1999, p. 3, original emphasis).

This passage refers to what is involved in a subject simply naming an object in a picture, say the picture of a horse. The subject might select "animal", "horse", "mare", "stallion", "pony", depending on the circumstances. "There is no simple, hard-wired connection between percepts and lexical concepts. That transition is always mediated by pragmatic, context-dependent considerations" (p. 3). A point I will be emphasising later is that, since which one of these words is selected makes a difference to the ultimate semantic meaning of the utterance of which the word forms a part, it follows that the semantic meaning cannot have been part of the pre-linguistic "notion/information to be expressed", formulated in some 'language of thought', before the process of "verbalization" began. (Looking ahead, this fact buttresses the argument I make in Chapter 6, that the content of a token thought does not exist in the head of the thinker in its final determinate form before the natural language production process transforms the pre-linguistic content into something accessible to the thinker/inner speaker. That process, I argue, makes a constitutive contribution to the content of the thought being expressed by the inner speech utterance.)

Lexical retrieval. This involves retrieving a lemma from the subject's lexicon, based on the preverbal message featured in the conceptualization phase. (A lemma is a mental representation that incorporates semantic and syntactical information.)

Morphological encoding. Morphemes are language 'atoms' – the smallest unit of language that carries meaning. This step retrieves the phonological shape of the word for the selected lemma. (E.g., *Eat* vs. *eats* vs. *eating* vs. *ate*.) (The 'tip-of-the-tongue' phenomenon is due to a momentary inability to retrieve the word form, given a selected lemma.)

Phonological encoding (also referred to as *syllabification*). For any morpheme, there will be numerous forms it could take, depending on its role in the utterance. Each form involves different syllables; this step specifies which phonemes and syllables to use, and metrical information such as stress.

Phonetic encoding. This step specifies the articulatory tasks needed to produce the utterance, but at an abstract level. It involves representing the articulatory gestures to be performed at different articulatory tiers: a glottal tier, a nasal tier, an oral tier.

Articulation. This is the execution of the abstract articulatory gestures into sounds.

I should note first that this model of the language production process is neutral on the question of the abstract versus the motor-sensory view of inner speech. On the abstract view, inner speech involves all but the last step in the process – that of articulation – and what we experience in inner speech is the words which would have been spoken had we proceeded to the next stage of articulation. (How this might work will become clearer in the next chapter.) By contrast, the motor-sensory view holds that the processes involved in the final stage – articulation – are also involved in the production and experience of inner speech. In their recent review of the cognitive and neuroscientific nature of inner speech, Loevenbruck *et al* (2018) find in favour of the motor-sensory view. I will present here a summary of their evidence, in three stages. First, I will present the evidence they offer *against* the abstract view and in favour of what they call the 'concreteness' of inner language. (They prefer the term 'inner language' to 'inner speech', since it allows them to

accommodate the phenomenon of 'inner sign' – the version of inner 'speech' experienced by deaf people.) Second, I will present their arguments *in favour* of the sensory nature of inner language. Thirdly I will describe their proposal for how to integrate the sensory-motor view of inner speech into the 'predictive-control' account of (bodily) action control.[7] This third stage gets a section in this chapter of its own (Section 3).

Evidence against the abstract view.

- Physiological measurements suggest that inner speech is physically planned, in the same way that overt speech is. For example, when subjects are engaged in inner speaking, their respiratory rate changes. The normal 'at rest' respiratory cycle is symmetrical – the duration of inspiration and expiration is the same. When speaking overtly the cycle is highly asymmetrical, with a short inspiration and a long expiration, during which speech is emitted. During inner speech, subjects show a slightly prolonged expiratory phase.
- Concerning muscular activity, electromyographic (EMG) studies show that the muscles of the lips and tongue are activated during episodes of inner speech. (When tests are performed on deaf subjects who sign, EMG records an increase in the *flexores digitorum*, a muscle in the forearm that flexes the fingers. "Behavioural studies have shown that the equivalent of inner speech in deaf signers involves internal representations of signs instead of auditory representations" (p. 139).)
- Concerning cerebral activity, studies show that both inner and outer speech production recruit essential language areas in the left hemisphere, and some are more active during overt speech than inner speech. According to Loevenbruck *et al*, the findings "support the claim that inner speech is a motor simulation of speech, including motor planning, but excluding motor execution" (p. 137).
- If the abstract view was right, then inner speech should be impoverished at the articulatory level compared with overt speech. But the evidence suggests otherwise. For example, the word "wristwatch" takes longer to pronounce than "wristband" in both the overt and the silent mode. It takes longer in the overt mode because of the

[7] As you would expect from a science paper, the original material from which this summary derived is replete with references to the original sources. These are to be found on pages 134 – 154 of Loevenbruck *et al* (2018). To make the text flow more smoothly I have excluded them from my summary.

articulatory movements required; the fact that it also takes longer in silent mode suggests that the same articulatory movements are specified, even if they are not executed.

In short, there is ample evidence to suggest that, contrary to the abstract view, the articulatory stage of speech production plays a significant role in inner speech. The key difference between overt and covert speech is that the execution of the motor plans made at the articulation stage are inhibited in the case of inner speech (more on that below). Loevenbruck *et al* summarise their findings as follows:

> Therefore, contrary to the Abstraction view, some instances of inner language seem fully physically planned, including concrete articulatory (laryngeal, orofacial, and manual) specifications that are coordinated, just like in overt language, but that are inhibited and not executed (p. 140).[8]

<u>Evidence in favour of the sensory nature of inner language.</u>

The previous section suggests that inner speech involves the specification of motor acts that are inhibited before they are executed. What about the sensory aspects of inner speech? These are just a few of the findings adduced by Loevenbruck *et al*:

- Neuroimaging studies show that during inner speech production the auditory cortex (specifically the superior temporal gyrus) is activated. "Although this activation is lesser than the one observed in overt speech, it entails that an auditory experience accompanies inner speech… The involvement of the mind's ear during silent reading has been recently confirmed by fMRI experiments" (p. 144).
- The idea of the "mind's ear" alone is insufficient to account for the sensory nature of inner speech. They argue that the sensory consequences of the "imaginary motor acts" which are involved in inner speech may be "multimodal". Inner speech, they say, involves somatosensory (i.e., bodily) sensations, such as proprioceptive

[8] An historical note: the motor-sensory view is not new, just better researched. John Hughlings Jackson (1835-1911), who studied neurological disorders, writes: "The objection that when we speak internally…there is no movement of the articulatory organs, is not of weight… We cannot surely suppose that different sets of sensori-motor processes are concerned when we "think" "gold is yellow" and when we "talk it"" (1866/1932, p. 85).

information about articulatory location, and tactile sensations, such as information about contact between tongue and palate. They coin the term "mind's *tact*" (p. 146) to capture this aspect of the inner speech experience.

They conclude: "To wrap up, the nature of inner speech is both motor and sensory" (p. 146).[9]

Section 3. The predictive control account of inner speech

You might accept, based on what I have presented so far, that the evidence appears to confirm that inner speech has a motor-sensory nature, but you might still wonder how the *experience* of inner speech is possible: How do we "hear" our inner voice if the words are not actually spoken? The answer to this question involves an appeal to the so-called "predictive control" account (also known as the "comparator model") of the motor control of *all* bodily movements. The idea is that, given the sensory-motor nature of inner speech, a version of this theory *should* in theory, and *does* in practice, apply to the phenomenon of inner speech. That version of the theory provides an answer to the question: What are we experiencing when we "hear" inner speech? I will start by providing a quick overview of the predictive control account of bodily movement generally, before showing how it applies to inner speech.

You might think that when we intentionally perform an action, any action, we have a system which figures out what movements we need to make to get it done, and then we just go ahead and make them. In a sense that's right, but it's a bit more complicated than that. First, we have a system which works out what sensory state we would *expect to be in* if the action we want to perform were to be performed successfully. That information is input to another system, the so-called inverse model, which works out, based on the sensory state we expect to be in, what movements we would need to make to be in that state. The

[9] Grandchamp *et al* (2019) reach the same conclusion: "Our data support the hypothesis that expanded inner speech recruits speech production processes down to articulatory planning, resulting in a predicted signal, the inner voice, with auditory qualities" (p. 1). By "expanded inner speech" they mean inner speech which is not "condensed" in the sense described by Vygotsky. For more on "condensed" inner speech see the next chapter on Unsymbolised Thinking (UT).

output of that system goes to another one, the motor command system, which plans the execution of the movements needed, and the plan goes to the motor system which generates the movements. But that's not all. As the movements are performed, the sensory systems perceive the actual state of the body, and this is fed back to a comparator system (one of three), which compares what the body is actually doing with what the sensory system originally determined was required for it to do for the performance of the action to be successful. Feedback from this comparator system goes to the inverse model to help it improve modelling inverse models in the future. And that's *still* not all. At the same time as the inverse model sends its input to the motor command system it also sends a copy (the so-called "efference copy")[10] to a system called the forward model. This system computes (predicts) precisely what sensory experience *would* result from executing this input, *were* it to be executed, and the output of that system is sent to another comparator system where it is compared to the sensory state the sensory system originally determined it needed to be in to meet the objective. If there is not a good match, adjustments can be made to the original desired state, and a refined version of that state can be computed. This is performed faster than it takes the other systems to produce and execute the motor plans, so this feedback can allow for early error correction *even before the motor plans are executed*. The predicted sensory feedback (from the forward model) is also sent to a third (and final) comparator system, where it is compared with inputs from the perceptual systems, to check whether there is a good match. This comparison of actual with predicted allows the performance of the forward model system to be improved over time.

So, that's an overview of the predictive control model of a common-or-garden bodily action, such as reaching for a light switch. If it is applied to the phenomenon of inner speech, it looks something like Figure 2. Figure 2 is, in effect, a more detailed view of the last box of the language production process (the box marked 'Articulation') illustrated in Figure 1 above.

[10] From the Latin, meaning, appropriately, "off on the side".

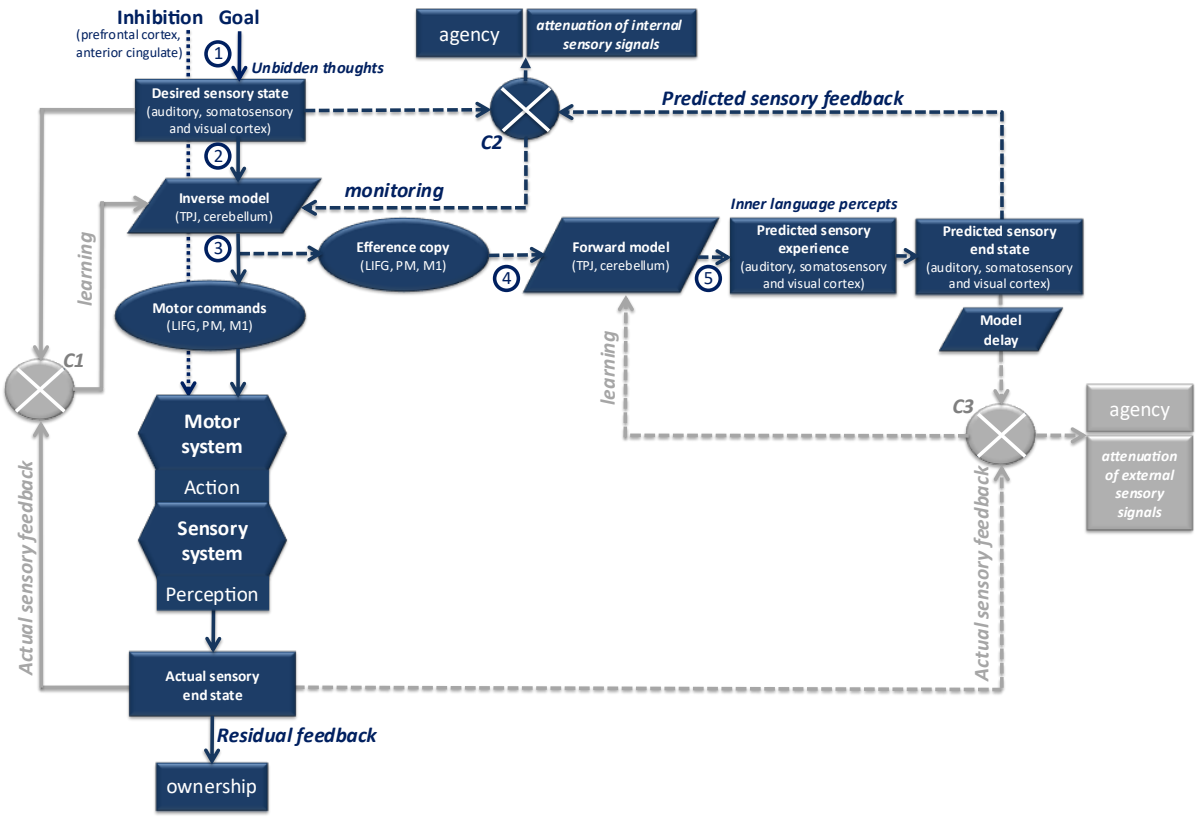

Figure 2. The predictive control account of inner speech. From Loevenbruck *et al* (2018).

Before I provide a short description of the steps illustrated in Figure 2, a key point to note is that this model of *inner* speech articulation is almost identical to the model for *overt* speech articulation. The only difference is that in the case of inner speech an instruction is sent to the motor system to inhibit the performance of the motor commands which have been specified as the ones required to voice the words. This instruction is indicated by the dotted line marked "Inhibition". During inner speech, the grey lines in this diagram are irrelevant.

So, what does this model describe? The 'goal' at the start of this model is the output of the phonetic encoding step – the penultimate step in Figure 1 above – a representation of the desired multi-sensory state, expressed in terms of articulatory and acoustic properties. (In a more recent, but unpublished paper, this has resulted in an amended version of Figure 2, in which the first step marked "Desired Sensory State" has been re-labelled "Supramodal Phonetic Goal".)[11]

[11] Helene Loevenbruck, personal correspondence. In fact, this change of nomenclature is anticipated in Loevenbruck *et al* (2018), where they describe their account as "an integrated account, [where] inner language

The supramodal phonetic goal (henceforth 'goal') is the input to an inverse model which calculates the motor commands required to achieve the goal, and these are sent to the motor system, to produce movements and sounds. In overt speech these are processed by the sensory system, which produces actual sensory experience and results in an actual sensory end state. This can then be compared with the goal (C1) where the comparison has the function of tuning the inverse model, so it produces more accurate output in the future. (C1 is irrelevant during ongoing speech, since if it was used in real time it would result in very slow speech production.) At the same time, an efference copy of the motor commands is generated and used to model what the sensory consequences would be were the commands to be executed. In overt speech, this prediction can be compared with the goal (C2) to allow for adjustments to the desired state, *even before the action is executed*. The predicted sensory feedback, to which a delay is applied, is also compared with the actual sensory feedback (C3), to improve the performance of the forward model. In covert speech inhibitory signals (the dotted line) are sent to the motor system preventing actual articulator movement from occurring. But, and this is a crucial point: *even if the motor commands are aborted, the efference copy is still produced.* The predicted sensory signal based on that copy is the inner voice that we 'hear', and the somatosensory experiences are what we 'feel', when we engage in inner speech.

In short, the predictive control account provides an explanation – a mechanism – of how it is that we 'hear' our own inner speech.[12]

Section 4. Agency

This section addresses some issues arising from the boxes in Figure 2. which are labelled "Agency", where agency is indicated as (somehow) 'flowing from' or 'being the product of'

is considered as deriving from multisensory goals, generating multimodal acts (inner phonation, articulation, sign) with multisensory percepts (in the mind's ear, tact, and eye)" (p. 132).

[12] For an alternative theory see Endicott' "Inner Speech and the Body Error Theory"(forthcoming). Endicott's hypothesis is that the experience of inner speech arises from a mix of interoception and audition. "Specifically, there is the detection of slight but well-confirmed activities in the speech musculature that occur during inner speech, which helps to transform representations of normal but quiet nonverbal sounds that inevitably occur during inner speech, from breathing to background noise, into a mistaken perception of inner speech." Which theory turns out to be correct does not make a difference to the key ideas in this thesis.

the comparison stages C2 and C3. What do the authors mean by their use of this term? This is where the predictive control account of inner speech – in fact, where the predictive control account of action generally – abruptly meets a long-standing and profound philosophical question: what is the nature of human agency?

We should start with the high-level idea that the predictive control account of motor planning plays a key role in self-monitoring for any creature, however modest, capable of even the simplest movement. (See Godfrey-Smith, 2017, Chapter 4.) Any creature capable of movement needs to be able to distinguish between (a) the changes it is experiencing perceptually which are caused by a change in its environment, and (b) the changes it is experiencing perceptually which are the result of its own movements. If the incoming sensory signals *don't* match an internally generated prediction, they might require attention and action – perhaps they are caused by an approaching predator, or prey. But if they *do* match a prediction then the creature can, in one sense, ignore them. The *experience* of actual changes matching predicted changes, it is argued, is one source of our sense of self, and of agency.

> [W]hether or not a perceived movement is experienced as one's own depends on whether or not there is match at the comparator between the sensory perception of the movement and an efference copy…. Whether you have the sense of agency [for the perceived movement] may depend on whether there is in fact match at the comparator, even though you do not have access to the content of the efferent copy. And as I said, this seems to be the simplest possible hypothesis to explain the basis of the sense of agency (Campbell, 1999, p. 613).

In other words, nature has designed a way for us to discriminate between changes in our perceived environment due to causes external to us, and changes in our perceived environment caused by us, and the latter is indicated by a particular phenomenological 'feel'; that 'feel', it is claimed, is what we are referring to when we refer to a 'sense of agency'.

According to Campbell (1999), the first person to speculate about the implications of this account of our sense of agency for *thinking*, and in particular for so-called "derangements" of thinking, was Feinberg (1978):

> Whereas the internal feedback associated with simpler motor acts is below the level of consciousness, one might postulate that the corollary discharges [i.e., efference copies] accompanying conscious thought are themselves conscious. If so, the subjective experience of these discharges should correspond to nothing less than the experience of will or intention. ... If thought is a motor process heavily dependent upon internal feedback, derangement of such feedback might account for many of the puzzling psychopathological features of the 'psychosis of thinking' (Feinberg, 1978, pp. 637-38).

While Feinberg expresses the idea in the form of a conditional ("*If* thought is a motor process") Campbell argues that it is. (His paper is entitled: "Schizophrenia, the Space of Reasons, and Thinking as a Motor Process.") His argument is brief and, in my view, incomplete, but is on the right tracks. He starts with a reasonable assumption about the source of our thoughts: they are "caused by a combination of our background beliefs, desires, and interests, together with current external stimuli" (Campbell, 1999, p. 617). He then suggests that something must "mediate" between the background beliefs and desires and the formation of the occurrent thought, and he proposes that what mediates is a motor instruction.

> What happens is that the background beliefs and desires cause the motor instructions to be issued, and that the motor instructions cause the occurrent thought (*ibid*).

His reason for proposing that motor instructions are the cause of the thought is that a motor instruction is, by hypothesis, always accompanied by an efferent copy, and the efferent copy is required "so that the ongoing stream of occurrent thoughts can be monitored and kept on track" (*ibid*). He concludes:

> [O]n this account, it is the match between the thought detected by introspection, and the content of the efferent copy picked up by the comparator, that is responsible for the sense of ownership of the thought (ibid).

Well, as I said, I think Campbell is on the right tracks, but also not quite right. How do background beliefs and desires "cause" motor instructions? What does he mean by "the thought detected by introspection"? What is the role which language production plays in the formation of a thought? Later Chapters in this thesis will attempt to provide some answers to these questions.

Using Figure 2. (above) we can fill in some of the missing detail. Recall that at comparison system C3, the predicted sensory feedback (the product of the forward model), to which a delay is applied, is compared with the actual sensory feedback (the product of perception) so that the performance of the forward model can be improved. The idea is that when the actual sensory feedback matches the predicted sensory signals, a sense of ownership and agency are experienced.

At the system C2, the predicted sensory feedback is compared with the desired sensory state; once again, the idea is that a match generates a sense of ownership and agency. What Campbell, Feinberg and others argue is that abnormalities in the functioning of the predictive control mechanism as it applies to speech explain the phenomenon of audio verbal hallucination (AVH), the experience 'hearing voices', where the voices being heard by the subject are experienced by the subject as not her own, but that of somebody else. The phenomenon is also known as 'thought insertion' (see e.g. Frith, 2012). According to these theorists,

> … if the prediction is faulty, the actual sensory consequences of inner speech are not attenuated and agency is not felt. Either because of attributional biases … or simply because self-authorship is not felt … inner speech would then be experienced as other-generated (Loevenbruck *et al*, 2018, p. 150).

The predictive control account as it applies to inner speech is thought by some to play other important roles, in addition to being responsible for a sense of agency:

> According to [Jones and Fernyhough, 2007], children start off by [engaging in] overt "private speech", simulating dialogues with interlocutors. Verbal thought would only become covert after several years, through a gradual process of internalization. During this process, it is crucial for children to be able to label the received auditory stimuli as self- or other-generated. This means that the efference copy is…ontogenetically necessary for inner speech to develop from private speech. We [Loevenbruck *et al*, 2018] further claim that distinguishing self-generated from other-generated voices remains compelling in adult inner speech… We can have imaginary dialogues, involving various voices. We claim that it is through self-monitoring that we do not mistake these internal voices for external voices, and that we are aware that we have imagined them (p. 151).

So, that explains some of the ideas which are being referred to by the boxes labelled 'Agency' in Figure 2. I want to conclude this section by touching on some of the implications for philosophical worries about agency which are raised by the arguments of this thesis.

One response to the predictive control account of motoric action has been to adduce it as evidence for a widespread illusion of conscious will. Wegner (2003) is one example of this response.[13] Wegner's claim is not that conscious thought does not cause actions, but that our impression that we cause our actions through acts of Will is misguided; the impression, he grants, is understandable given phenomena like the predictive control account of action but is nevertheless an illusion. On his account, consciousness itself has a causal explanation and the idea of a conscious Will being responsible is what he calls "the mind's best trick" (p. 65).

> The experience of conscious will is a marvelous trick of the mind, one that yields useful intuitions about our authorship – but it is not the foundation for an explanatory system that stands outside the paths of deterministic causation (2003, p. 68).

Wegner's concern here is with bodily action, and he is happy to endorse the idea that conscious thoughts cause actions. But what causes conscious thoughts? If our sense of "willing" our *bodily* actions is undermined by the predictive control account, and if language

[13] Wegner's conclusion, if not all his argumentation, is endorsed by Carruthers' paper (2007), entitled "The Illusion of Conscious Will."

production is constitutive of the thought process (as I argue in this thesis), and if the predictive control model is part of the language production process, then shouldn't the predictive control account also undermine our sense of "willing" our *mental* actions, such as our thoughts? Arguably it should. But that assumes that our thoughts are mental actions in the first place, and now we are faced with a dilemma. If we insist that thoughts *are* mental actions, (where an action is characterised as something we do intentionally) then, like our bodily actions, their status as intended or willed would seem to be undermined by the role that the predictive control account plays in their production. Their *seeming* to be intended or willed is just an illusion. On the other hand, if we deny that thoughts are actions, meaning that they are *not* intended or willed then, once again, their seeming to be intended or willed is an illusion. So, either way, thoughts are not the kind of things that are intended or willed.

A number of philosophers have expressed scepticism about the whole idea of 'mental action', as we shall see in future chapters. It has led to a recent surge of interest in what Bayne & Levy (2006) call 'will-skepticism. I have nothing to add to that debate except this: If, as I argue (particularly in Chapter 6), inner speaking and thinking are (sometimes) *the same event*, then the role of inner speech in conscious thought is at the heart of the debate about the nature of our agency.

Section 5. Summary

This chapter has been an introduction to the phenomenon of inner speech, focusing on the latest scientific and psychological understanding of the phenomenon. There is widespread support for the Vygotskian idea that inner speech should be understood as an internalised form of outer speech. There is also strong evidence for the motor-sensory view of inner speech – the idea that it is embodied, and that it involves physical processes that unfold over time. I described the different stages of Levelt *et al'*s (1999) high level theory of speech production, which is neutral on the question of whether inner speech is abstract or motor-sensory in nature, but then I offered evidence (from Loevenbruck *et al,* 2018), for the motor-sensory view. To explain how we come to experience an inner voice at all I offered a detailed description of the predictive control account of inner speech, according to which

our inner speech is, in effect, the experience of an audio image of a prediction (or forecast) of what our utterances, as generated by the language production process, *would* sound like *were* they to be voiced. The instruction to voice them is inhibited, so the words are not spoken but nevertheless 'heard' by one's 'inner ear'. I explained how the predictive control account, when applied to the language production process, is thought to be responsible for a sense of ownership of our thoughts, and therefore of our sense of agency. I ended by registering how this observation puts the phenomenon of inner speech at the heart of current debates about mental action.

Chapter 3. An inconvenient fact: the phenomenon of unsymbolised thinking

Section 1. Introduction

Anyone who wants to explore the relationship between language and thought has to address an inconvenient fact: some people claim to think without words. Of course, 'thinking' is a term used to describe a wide variety of mental activities, and it is not strange to suppose that some kinds of thinking don't involve words. The challenge posed by the phenomenon of Unsymbolised Thinking (UT) is that, on the face of it, it seems to involve the kinds of thinking that we might expect *only* to be possible with words. To illustrate what I mean, here's an example from the UT literature (Heavey & Hurlburt, 2008):

> Adam was watching two men carry a load of bricks in a construction site. He was wondering whether the men would drop the bricks. This wondering did not involve any symbols, but was an explicit cognitive process (p. 802).

You might wonder how it is *possible* for someone to have such a thought, with such determinate contents, "without the awareness of that thought's being conveyed in words, images, or any other symbols" (*ibid*). How can a person entertain a thought which contains the concepts MEN, DROP and BRICKS without using, and being aware that they are using, the words "men", "drop" and "bricks"? If it *is* possible, then the best that can be said for the relationship between inner speech and thinking (or at least this kind of *conceptual* thinking) is that inner speech is *one* way of having such thoughts, but not the only way. This possibility can't be ruled out *a priori*, but if it's true then the arguments of this thesis will be less convincing than they otherwise might be. Thankfully, I think this conclusion can be avoided.

This chapter has three parts. In Section 2 I briefly introduce the phenomenon of UT. In Section 3 I describe a version of the motor-sensory view of inner speech which attempts to accommodate the phenomenon of UT (Vicente & Jorba, 2019). In Section 4, in anticipation

of later chapters, I describe Vicente and Jorba's hypothesis that it is only when the contents of a thought are brought to consciousness – by the production of an inner speech utterance – that the content of that thought is fully determined. That idea represents a significant break with philosophical orthodoxy, which typically has it that the function of an utterance is to express a thought which already exists. It is an idea which I develop in this thesis, culminating in Chapter 6, where I argue that we should abandon the idea that a thought is an *object* of any kind at all. I argue there that, on the principle that a token thought is individuated by its content (in the broadest sense), it makes no sense to say that a token thought has one content at one time and a different content at another time. Rather, we should say that a thought is not a 'state of mind', a representation of something determinate, but something we *do*: an action we perform intentionally to get something done (with words). In short, a thought is the performance of a speech act. But I am getting ahead of myself. I turn now to the phenomenon of unsymbolised thinking.

Section 2. Unsymbolised thinking

For many years now, Russell Hurlburt has been using a technique known as Descriptive Experience Sampling (DES) to try and build a picture of our naturally occurring conscious experience.[14] The technique (like others developed for the same purpose[15]) uses a beeper that goes off at random intervals as participants go about their normal business. The participants are trained to stop what they are doing and pay close attention to their experience at the moment the beeper beeps; they make notes of their observations and are subsequently interviewed about what they noted. The training of the participants and the interviewers is aimed at preventing the influence of presuppositions about the nature of inner experience from distorting the description. The goal is a faithful description of a single, randomly chosen moment of experience at a particular time. Although all such techniques are controversial, the philosopher Schwitzgebel collaborated with Hurlburt to review DES from both a philosophical and psychological perspective (Hurlburt &

[14] See the reference section of this thesis for details. For even more details, see the reference section of Hurlbert & Heavey (2018), which lists 18 of Hurlbert's papers on this subject.
[15] Experience Sampling Method (ESM) and Ecological Momentary Assessment (EMA). Heavey & Hurlburt (2008) explain at length the virtues of DES over these other techniques.

Schwitzgebel, 2007) and "concluded at least tentatively that DES may be well suited to developing high-fidelity descriptions of moments of experience" (Heavey & Hurlburt, 2008).

One study (Heavey & Hurlburt, 2008) set out to answer a simple question: What are the base rates of the common phenomena of inner experience? (In what follows I will concern myself with the results, not with the methods used.) The authors' report of that study gives details of the five most frequently occurring phenomena as reported by the participants. These were (listed here together with Heavey & Hurlburt's description of each category):

- Inner speech. Speaking words in your own voice, usually with the same vocal characteristics as your own external speech, but with no external (real) sound or motion.
- Inner seeing. Seeing something in imagination that is not actually present.
- Unsymbolized thinking. Thinking a particular, definite thought without the awareness of that thought's being conveyed in words, images or any other symbols.
- Feeling. Affective experiences, such as sadness, happiness, humour, anxiety, joy, fear, nervousness, anger, embarrassment, etc.
- Sensory awareness. Paying attention to a particular e=sensory aspect of the environment where that sensory experience is itself a primary theme or focus apart from the object of perception.

One (or more than one – some experiences were reported as occurring simultaneously) of these inner experiences occurred in approximately one quarter of the sampled moments. Of the numerous interesting findings of this study, I will just mention two of the most pertinent. The first is the overall frequency of inner speech (i.e., the frequency with which all participants reported an experience of inner speech). At 26%, this was substantially less than other studies have reported. (Klinger and Cox (1987-1988) reported the frequency at 75%; Baars (2003, p. 106) claims "human beings talk to themselves every moment of the waking day".) Heavey & Hurlburt attribute this discrepancy to the superiority of DES over other techniques.

> For example, DES shows repeatedly that many, if not most, people who have unsymbolized thinking…will at first report such thinking to be in words. Only after repeated training as they

> iteratively confront *the apprehension of their own experience* do they come to recognize their presupposition of words as being false (p. 805-6. Original emphasis).

The second finding was the frequency of UT at 22%. In other words, participants spent nearly as much time experiencing UT as they did inner speech. Furthermore, half the participants experienced UT in at least a quarter of their samples. One might wonder whether participants who report UT are having the same kind of experience as those who report inner speech but, for whatever reason, fail to remember the words, visual images or other symbolic experiences. But that explanation can be rejected: all the participants who experienced UT in *some* of their samples also experienced words and/or images in *other* of their samples. Clearly, these participants were quite capable of distinguishing UT from inner speech.

Section 3. How can the motor-sensory view accommodate unsymbolised thinking?

This is the challenge that Vicente & Jorba (2019) take on, and they start by rehearsing the predictive control account of inner speech described above in Chapter 2. (Recall that this theory focuses on the way in which a prediction of the output of the motor commands generates a prediction of the auditory experience the subject would have were the commands to be executed; the theory claims that it is the prediction that we experience as inner speech.) Vicente & Jorba invoke evidence from psycholinguists to the effect that monitoring during speech production process might involve making predictions at many levels during that process – semantic, syntax and phonology – and not just at the articulation (motor) level. They cite Pickering and Garrod (2013) as claiming that semantic error correction, at least in overt speech, occurs faster than would be possible were it to depend on detecting the errors at the articulatory level. If that's right, then the error must be corrected by comparing the output with a predicted *meaning*, rather than a predicted sound (which is then interpreted for meaning). Vicente & Jorba set out the following hypothesis:

> [W]hen we form the intention to express a certain content and we refine that intention to the point of executing the motor commands that ultimately realize it, we issue predictions not just of how the utterance is going to sound, but also of what it is going to mean (p. 746).

Their idea, in short, is that our experience of inner speech is based on becoming conscious of two predictions, not just one: a prediction of what the utterance would sound like (if it were to be executed) and a prediction of what the utterance would mean. Now, what do they mean by "a prediction of [what] the utterance…is going to mean"? How could such a prediction be independent of a prediction of how the words would sound were the motor instructions to be executed? Here is how I understand their proposal. Suppose that, as implied by the motor-sensory account, at a certain point in the speech production process there is a representation of the thought to be expressed, in the format of natural language, with some more-or-less fully formed semantic and syntactic properties, which has nevertheless not yet been modelled for articulation by the motor planning system. It seems quite plausible that such a representation might be monitored first, to ensure that it is in line with the communicative intention which caused, sustained, and guided it, before the effort of planning for its physical articulation. For that monitoring to happen this representation is parsed by the language comprehension module, and its meaning made conscious, *before* it is processed by the motor planning system. This results, in effect, in a prediction of meaning. The benefit of this step would be that the prediction is then compared with the original communicative intention for its fit and, if it doesn't match, it can be changed (before motor planning takes place). It would make sense for this check for meaning to be performed fractionally before the motor planning system models it, since if the meaning isn't quite right then the words will need to change too. To that extent it is plausible to think that the processes are separate. The key point is that if the motor planning was aborted *just after* the meaning check was done, the conscious experience would be of meaning but not of words. And this is exactly what Vicente and Jorba propose:

> Now, what would happen if an instruction to speak never reached the motoric component? Suppose that instead of inhibiting our intention to express some thought content at the level of speech commands, we inhibited it at a previous level, say, at the level where we have given

> form to our intention in terms of semantics and syntax, but not yet in terms of phonology (ibid.)

If this were to happen, they speculate, the content of what you became conscious of would have all the structure and semantics associated with a natural language utterance, but you would not be conscious of words. (Perhaps this explains why, without training, we are so inclined to interpret UT as inner speech and need to be trained not to do it – see the Heavey and Hurlburt (2008) quote above). Vicente and Jorba call this phenomenon "aborted inner speech". They say: "We think that this view of UT accords with and can be seen as a continuation of inner speech, and is appealing once the inner voice is seen as being derived from aborted commands" (ibid.).[16]

You might still think this is all a bit quick, a bit too convenient. Isn't there an alternative explanation of UT? There is one popular alternative I will mention briefly, for the sake of completeness, before dismissing. That is the argument from phenomenal intentionality, and more specifically, from cognitive intentionality. Phenomenal Intentionality Theory (PIT) promotes the idea that the intentionality of a mental state – i.e., that which individuates it from any other mental state – derives from the phenomenology associated with that mental state (for a recent comprehensive overview see Mendelovici, 2018. Also, see the papers collected in Bayne and Montague, 2011). Cognitive phenomenology applies PIT to the contents of occurrent episodes of thinking. According to this view, what individuates the experience of, say, judging that *p*, is the phenomenology associated with judging that *p*. On this view, UT isn't problematic. There is no mystery about how we can experience a determinate thought without experiencing words, since it is the experience itself which metaphysically grounds the determinacy of the thought. The experience of words may frequently *accompany* the experience of thinking, but in those cases the words do not *determine* the meaning of the thought – that role is reserved for the experience itself.

I find this view highly implausible. I don't reject the claim that a special kind of phenomenology is associated with cognition – that seems to me highly likely. It might even

[16] They also argue that their account makes UT continuous with other phenomena that Hurlburt and others have identified, such as partially worded inner speech. I will not pursue that here.

be that each individual thought has a very slightly different phenomenology from every other thought (although I find this less plausible). Perhaps a particular kind of phenomenology *supervenes* on thinking. But the idea that the phenomenology of a particular thought metaphysically *individuates* it is utterly mysterious. How could it do that? How could something like a *feeling* ground all the rich and complex aspects of a fully specified thought? Of course, those who endorse PIT have a response to this question (see, for example Mendelovici, 2018, Chapter 7). But I find it far more plausible to think that it's *words* (and the intention which causes, sustains and guides their production) that determine the meaning of a thought, not the experience of having the thought. This is hardly an argument, but it would take me too far from my subject to mount one here. For arguments against cognitive phenomenology see Carruthers & Veillet (2011) and Prinz (2011).

Section 4. Anticipating what comes next

Before concluding this chapter, and before leaving behind Vicente & Jorba's (2019) paper on UT, I want to address, briefly, a question they raise, the importance of which will be explored in future chapters. The question is this: does an inner speech utterance merely express a thought which already exists, or does that token thought undergo some kind of change during the language production process? For everything that has been said so far, either option is still a possibility. Philosophical orthodoxy is that the content of the thought is fully determined before the language production process starts. One view, for example, is that a thought is a sentence in a language of thought (Mentalese) which simply needs to be translated into a natural language so that it can be communicated (to others) or made conscious (to the thinker). On this view, the propositional contents and its truth conditions are the same for both the pre-linguistic representation (e.g. the sentence in Mentalese) and the natural language utterance which expresses it. This orthodoxy will be challenged by me in future chapters and is challenged by Vicente & Jorba (2019).

As I read them, Vicente & Jorba offer both a weak alternative to this view and a strong one (my terminology, not theirs). The weak alternative is that a pre-linguistic thought undergoes a "representational re-description" when it takes on a linguistic form. How much re-description is not clear, but they cite some evidence that some semantic representations

"do not map neatly onto our conceptual structure". If that's right, then the contents we are conscious of will have truth conditions "given by the categories drawn by the words of our language and structures that constitute the grammar of that language" (p. 750). The implication is that these truth conditions will inevitably be different from whatever truth conditions are relevant to the pre-linguistic thought before it was re-described. I call this version 'weak' because it does not amount to a radical re-appraisal to the relationship between thought and language. It offers a slight amendment to the standard model – a thought goes through some minor re-description when it is translated into natural language, but nothing more. (This idea is revisited in more detail in Chapter 6, Section 11, which discusses the "thinking for speaking" thesis as proposed by Slobin, 1987.)

The second alternative they explore – what I call the strong alternative to the standard model – is more radical. Rather than accepting the idea that pre-linguistic thoughts have their contents fully determined before they are translated (with minor "representational re-description") into natural language, Vicente & Jorba explore the possibility that "thought content needs to be brought to consciousness for its content to be fully determined" (p. 26). This is a radical idea because it puts pressure on the idea that what is represented pre-linguistically actually qualifies as a thought at all (or so I will argue in Chapter 6). How so? Roughly: If, as is widely assumed, a particular token thought is individuated by its content, and if the content of a thought is only determined after it is made conscious in language, then the pre-linguistic content does not individuate it, so *it* (whatever entity exists before the language production process) does not qualify as *the* thought. Even if we decide it does qualify as *a* thought (of some kind), it cannot be the *same* thought as the one whose content, by hypothesis, is different, in virtue of having (linguistic-based) content which the pre-linguistic thought lacks.

In effect, Vicente & Jorba (2019) raise the question, mooted in the Introduction to this thesis, concerning the metaphysical nature of a conscious thought. Should we understand a token thought to be a token *thinking*, an object of some kind (such as a representation in the head of a thinker), or as the content or meaning of what if represented? They do not pursue these metaphysical questions themselves, being more interested in showing how

what I call their strong alternative to the standard view fits with the phenomenon of UT. But they anticipate where this thesis is going. I pick up these questions again in Chapter 6.

Section 5. Summary

Anyone who wants to explore the relationship between language and thought has to address an inconvenient fact: research suggests that quite a large percentage of people experience conceptual thought without experiencing language. How can this be? And how is such a finding compatible with the sensory motor view of inner speech, which has it that inner speech is the experience of 'hearing' an auditory image which is a prediction, in effect, of what the speech would sound like were it to be vocalised? I offer an explanation, from Vicente and Jorba (2019), which proposes that, before the utterance is submitted to the motor planning system (which is responsible for generating an efference copy of the motor plan), the utterance is available to the language comprehension system, and its meaning made available to the subject. If, immediately after that, the intention to make the utterance is abandoned, and the motor planning process aborted, the subject will have the experience of the meaning of what they were thinking without the experience of the language that determined the meaning.

Part II: Four essays on the philosophy of inner speech

As previously stated, Part II contains what I consider to be the philosophical meat of this thesis. Whereas Part I provides important background on the phenomenon on inner speech, and Part III contains further reflections on the metaphysics of conscious thought, Part II contains essays which tackle 4 related philosophical questions arising from inner speech. (1) Is it something we do, or something that happens to us? (2) When we reason (silently) in inner speech, does that activity qualify as fully intentional or not? (3) If thinking involves *inner speech acts*, as I claim, what contribution does the performance of the speech act make to the meaning of the act? (4) If thinking-as-speaking qualifies as a skilled activity involving bodily capabilities, how does it compare with other skilled activities?

Chapter 4. Action and reaction: the two voices of inner speech[17]

Abstract

Is inner speech an intentional action, something we do, or an event, something that happens to us? This chapter argues that it can be both, (although not at the same time). Some inner speech utterances are events – reactions to the circumstances we find ourselves in: they are spontaneous, they require no effort, and we are not in control of their occurring. These inner speech utterances fail to satisfy the criterion for qualifying as intentional actions, as stipulated by three popular theories of action. But some other inner speech utterances, by contrast, *are* intentional actions, performed deliberately, effortfully and with as much control as any other intentional action. When we deliberate, for example, inner speech utterances are the basic actions by which we bring about the non-basic action of trying to come to a conclusion, make a decision, solve a problem, etc. These inner speech utterances *do* meet the criterion for qualifying as intentional actions, as stipulated by three popular theories of action

Section 1. Introduction

[17] Much of the material in this chapter appears in (Frankfort, 2022). The paper's argument is mentioned briefly in the recent SEP (*Stanford Encyclopaedia of Philosophy*) entry on 'Inner Speech'.

Frankfurt (1988) anticipates some of the key themes of this chapter when he writes:

> 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 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 obsessional thoughts, whose provenance may be obscure and of which we cannot rid ourselves; thoughts that strike us unexpectedly out of the blue; and 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 (p. 59).

In this chapter I explore Frankfurt's distinction between active and passive as it applies to the phenomenon of inner speech. I start with Gregory (2020), who argues that inner speech is 'reactive', meaning, roughly, that it occurs as an automatic, spontaneous and uncontrolled response to the context a subject finds herself in. For example, on seeing new leaves on the trees in the park, the words "Spring has arrived!" just pop into a subject's head. He argues that inner speech utterances like this fail to meet the three criterion which each of three leading action theories demand of an event if it is to qualify as an action. Specifically, and unlike genuine actions, reactive inner speech utterances (1) are not performed for a reason, (2) they are not under the control of the subject, and (3) they don't involve any effort. I agree with Gregory that some inner speech utterances are reactive in this sense, but I reject the claim that all of them are (Section 2).[18]

In what follows, I use the term 'episode of deliberation' to cover a broad range of conscious mental activities that a subject can perform silently, such as deliberating, reflecting,

[18] Gregory writes, "My focus is on the inner speech utterances which form parts of the ordinary inner monologue; the ones which accompany our everyday activities; the ones we produce without seeming to think about it" (p. 57). Strictly speaking, this leaves open the possibility that he thinks there are other kinds of inner speech which might *not* be 'reactive'. But he is not explicit about this and makes no mention of any other kinds of inner speech. Arguably, by focusing on the kinds of inner speech "we produce without seeming to think about it", he implies this is the most significant kind. I disagree. The most significant kind of inner speech, I claim, is the kind that is the main subject of this thesis: the kind we make intentionally in order to get something done – work out what to think or what to do; come to a decision; reflect on our options; etc.

reasoning, considering, evaluating, and so on.[19] The characteristic of an episode of deliberation, as I am using the term, is that it has a purpose: to come to a conclusion, to reach a decision, to solve a problem, etc. Episodes of deliberation, I claim, typically involve inner speech utterances, the function of which is to help the subject bring about the purpose of the episode. For example, if I am in a restaurant, handed a menu and invited to choose a starter and a main course, I am likely to conduct a little dialogue in my head: "What is in season now?" "Is fish likely to be good here, this far from the sea?" "Does that starter go with that main course?" And so on. I will likely have to break off from conversation with my companions in order to concentrate on the task of 'trying to choose'. These inner speech utterances, I argue, are a kind of action. They are also, *pace* Gregory, examples of "inner speech utterances which form parts of the ordinary inner monologue; the ones which accompany our everyday activities".

Following Mele (2009, pp. 18-37) I emphasise the distinction between 'trying to *x*' and 'trying to bring it about that I *x*'. Some kinds of 'trying to x' are not, strictly speaking, actions; 'trying to fall asleep', for example, isn't an action, because falling asleep is something that happens to me. By contrast, 'trying to bring it about that I fall asleep' *is* an action, since there are things I can do to bring it about – counting sheep, for example, or taking a sleeping pill. So, the action of 'trying to bring it about that I *x*' is not a *basic* action (such as the action 'raising my arm'), because it requires the subject to perform other actions to execute it. I argue that, during episodes of deliberation, inner speech utterances satisfy the criterion for being actions according to two of the popular theories of action, albeit non-basic actions (Section 3).

Where *x* is the purpose of an episode of deliberation, I suggest, one of the basic actions I typically make to execute the (non-basic) action of 'trying to bring it about that I *x*', is to perform an utterance in inner speech. For example, to perform the action of 'trying to bring it about that I choose what to order from the menu' I might make the kind of inner speech utterances in the 'menu dialogue' above. In Section 4 I argue, first, that speaking *overtly* is

[19] Henceforth, all references to 'deliberation' should be taken to mean 'silent conscious deliberation', unless otherwise stated.

very often a basic action, and second, that there is no reason to think that inner speech is not also a basic action, just because it is speech which is not vocalised. In Section 5 I address the criterion for qualifying as an action as stipulated by the third popular theory of action, the criterion of control, before concluding in Section 6.

Section 2. Reactive inner speech

Gregory (2020) argues that inner speech utterances are not actions and therefore not speech acts. (I will address his own exception to this generalisation at the end of this section.) His argument involves considering three leading theories of action and then showing how our inner speech utterances fail to qualify as actions on any of them. The different action theories claim, very roughly: (1) actions are things we do which can be explained by our reasons for doing them (Davidson, 1963); or (2) actions are things we do which are under our guidance or control (Frankfurt, 1978); or (3) actions are things we can try to do (and also fail to do, despite trying) (O'Shaughnessy, 1973, and Hornsby, 1980).

2.1 Actions are things we do which can be explained by our reasons for doing them

We can't, Gregory claims, provide *reasons* for why an inner speech utterance takes place. He argues that our attempts to do so are invariably confabulations, and he does it with the help of the following illustration.

> Suppose you are walking through a park one day towards the end of winter. Noticing some green leaves, you produce the inner speech utterance, 'Spring's starting', without having consciously decided to do so. You then find yourself wondering why you produced the utterance. It seems like there are two things you might say:
> 1) 'I wanted to make the propositional content that spring is starting salient in my consciousness, and I believed that producing the inner speech utterance would achieve this.'
> 2) 'I don't know. I just did.' (p. 64).

If 1) was true, then you would be describing an action, because you would be providing a reason, in terms of your beliefs and desires, for why you uttered what you did. But 1) is very implausible, a confabulation in fact. The only plausible account is given by 2). Note that in this example you might well *rationalise* what you said, and *guess*, after the event, that the

sight of green leaves was in some way responsible for your utterance. But this is not the same thing as explaining your reasons for saying it; an explanation in terms of reasons would typically involve describing the beliefs and desires which caused you to say it. Another way of putting the same point is to say that the inner speech utterance in this example was not intentional.

2.2 Actions are things we do which are under our guidance, or control

According to Frankfurt (1978), for something to be an action it must be possible for the subject to adjust what they are doing while they are doing it. To qualify, such adjustments must be attributable to the subject, and not to some automatic mechanism possessed by the subject. For example, if I get up from my chair to get a beer from the fridge and notice the dog is in my way, I will take steps to go around the dog in order to fulfil my intention. Compare this with what happens if I get up and accidentally trip over the dog and have to make rapid adjustments to prevent myself falling on my face. The former behaviour demonstrates control by me, the subject. The latter behaviour demonstrates control by my automatic reflexes; what guides my movements is a reflex which kicks in automatically in order to prevent injury. Overt (i.e., out loud) speech is under our guidance in the required sense; it is adjusted by the speaker over the course of the speech episode, both in response to hearing our own words and to the reactions, as we perceive them, of the listeners to those words. (I will say more about overt speech in Section 4 of this chapter. I will say more about *speech acts* in Chapter 6.) Inner speech utterances, by comparison, do not involve the bodily movements involved in speech acts, according to Gregory, but only "the generation of phonological representations" (p. 68). Since we have no control of the processes involved in forming phonological representations, we have no control over inner speech utterances. So according to this theory of action, inner speech utterances are not actions. (In the next chapter I will have more to say about how we should understand the purely mechanistic automatic processes involved in speech production, those quite clearly outside our control, when I consider Dual Process theories of cognition.)

2.3 Actions are things we can try (and fail) to do

In opposition to this theory of action, Gregory's claim is that our inner speech utterances do not require any effort; we don't try to produce our inner monologue, it just happens. In

fact, it is hard to imagine how we could *prevent* it from happening. Furthermore, we never have the experience of failing to make an inner speech utterance we wanted to make, or of failing to make one the way we actually made it. The best explanation for why we never have the experience of failing, he argues, is that we never have the experience of trying.

To recap this section so far, Gregory offers arguments for why inner speech utterances fail to meet any of the criteria for being actions as stipulated by three leading theories of action: being done for a reason, being under the subject's control, and being something the subject tries to do. So, what are inner speech episodes if they are not actions? Gregory's proposal is that these kinds of utterances are neither actions nor mere reflexes but rather, "more like automatic *re*actions". What we are reacting to is the context which we find ourselves in at the time of the inner speech utterance, where context includes both the external environment (as in the "Spring's starting" example above) and our other mental states:

> [T]he automatic process that produces them is to a significant extent sensitive to context. In this way, the utterances of our ordinary internal monologues are like unbidden imaginings and unbidden memories: events which take place in the mind, which we would not consider actions, but which are closely related to our other standing and occurrent mental states (p. 71).

Following Gregory, I will call this kind of inner speech "reactive inner speech". He goes on to argue that although inner speech utterances are not actions, nevertheless we *treat* them as speech *acts*. His argument for this is that unless we *treated* them as speech acts, we would not experience them as meaningful, but merely as "auditory images", sounds without meanings. He draws an analogy with overt speech acts.

> If someone produces an audible utterance but does not take themselves to be performing a speech act – if they believe that they do not have intentions of the appropriate kind – then they must believe that they are just producing sounds… In parallel, if someone produces an inner speech utterance but does not take themselves to be acting on intentions of the relevant kind, then they must believe that they are just producing auditory imagery, not a linguistically meaningful utterance (pp. 60-61).

I have two reservations about this argument. First, while it is true that we experience inner speech utterances as meaningful, and not as meaningless sounds, it is not obvious that the only explanation for this is that we treat them as intentional. No further argument is offered by Gregory for the claim that we do. Maybe inner speech utterances are meaningful for some other reason. Isn't it possible that the cognitive processes involved in speech production, whether silent or voiced, *guarantee* that inner speech utterances are experienced as meaningful, without that experience of meaningfulness having to be underwritten by the additional psychological mechanism of treating the words as intentional? For example, as we saw in Chapter 3, Section 3, Vicente & Jorba (2019) argue that when the motor planning system is involved in speech production it not only produces a prediction of the *sounds* that executing the motor plan will produce, but it also produces a prediction of the *meaning* of those sounds. Since there is very good evidence that the motor planning system *is* also involved in the production of inner speech (see Chapter 2, Section 2) this would explain why we experience inner speech utterances as meaningful, without the additional step of our having to treat them as intentional.

Second, the claim that we treat reactive inner speech utterances as intentional is puzzling on its face. In practice, it seems to me, the defining characteristic of *reactive* inner speech utterances is not only that they are not intentional, but that we don't, *in fact*, treat them as intentional either. Rather, we treat them exactly as Gregory describes our experience of them – spontaneous, automatic, "unbidden", and as utterances "we do not consciously plan to produce" (p. 1).

A more significant worry is the following. I said earlier that I would address Gregory's claim that there are exceptions to the general rule that inner speech utterances are not actions, and that's what I will do now. He allows that "you can consciously decide to produce an inner speech utterance and then do so; the resulting utterance is an action" (p. 57). For example, a subject might consciously decide to say to herself in inner speech, 'Grass is green', and then do it. We can all agree with Gregory that in this case the inner speech utterance 'Grass is green' is an action. But I want to draw attention to what would have to occur before that action is performed. By hypothesis, this is a conscious decision to perform an inner speech utterance, but how did that decision come to be conscious? Presumably it

took the form of an inner speech utterance, such as: "I am now going to say to myself in inner speech 'Grass is green'". (For convenience, let's shorten this inner speech utterance to "I'm now going to say 'S'".) But now we can ask about *that* inner speech utterance: was it merely reactive?

There seem to be two possibilities. The first is that the utterance "I'm now going to say 'S'" arose spontaneously as an episode of reactive inner speech. I don't want to deny this possibility, and in a moment, I will illustrate how that could, in principle, happen. A second possibility is that the decision was the result of thinking about a problem and coming to a conclusion, the result of which was the decision to say, "I'm now going to say 'S'". The worry is that the second possibility is not only much more likely but is not accommodated by Gregory's analysis of inner speech as reactive.

I will expand on the second possibility shortly. For now, let me provide an example of how a conscious decision (including, by implication, the decision to say something to oneself in inner speech) could, in principle, be the result of nothing more than a series of reactive inner speech utterances. Let's take Gregory's own example from earlier, of the subject who notices signs of Spring and says to herself "Spring's starting". It seems to me possible that this reactive inner speech utterance might trigger a series of other inner speech utterances which result in the formation of an intention. For example:

a) Spring's starting
b) Spring bulbs will be coming up about now
c) I planted some Spring bulbs last Autumn
d) I wonder if they are coming up now
e) I must remember to check my garden to see if they are coming up now

In this case the inner speech monologue resulted in the formation of the intention to remember to check the garden to see if the bulbs are coming up. In the light of this example, we should allow for the possibility that a series of purely reactive inner speech utterances *could* result in the formation of an intention to say "I'm now going to say 'S'" in

inner speech. (Perhaps this would only happen to a philosopher with speech acts on his mind, but that's beside the point. In principle, it could happen.)

Now let's turn to the second possibility – the possibility that the formation of an intention was the result of setting out to solve a problem and coming to a conclusion. In this example the subject makes the following series of utterances in inner speech, one after another:

i. If it's likely to rain, I had better take an umbrella
ii. Does it look like rain?
iii. It looks like it might rain
iv. I will take my umbrella

In this case, not only did the series of inner speech utterances *result* in the formation of an intention, but the whole series was intentional *from the start*. The intention formed at the end of the episode was not an incidental feature of the episode, as in the first case we considered, but rather *its very purpose*. The utterances which comprise this episode of inner speech are all connected and guided by the same intention – the intention to decide whether to take an umbrella. The intention motivates, sustains and guides the succession of utterances in a way that is absent in the case of episodes of reactive inner speech, where the succession is merely experienced as the result of spontaneous association. While I conceded, above, that in principle the inner speech utterance "I'm now going to say 'S'" *might* have occurred purely reactively, I suggest that it is much more likely to have occurred as a result of a conscious intention. In the next section I will have more to say about episodes of inner speech that have this characteristic, and the way in which they are intentional.

In summary, Gregory's account has it that most inner speech, with the exception already discussed, is reactive. If, as I argue, some inner speech utterances qualify as intentional actions, then that account is at best incomplete. But not only does his account leave out a whole class of inner speech utterances, it leaves out what are surely the ones we care most about. Intentional inner speech is the kind we value most because the decisions we make based on our purposeful silent reflection are some of the most *important* decisions we

make, decisions which shape our lives and sustain our sense of agency. It's because so much of what I do is based on what I have decided, based on my reflection, that I feel such a strong sense of ownership and authorship for my actions – i.e., my agency. So, the criticism of Gregory's account of inner speech is twofold: first, it ignores a large class of inner speech utterances which are intentional actions; second, these are the inner speech utterances that underpin our sense of ourselves as rational agents.

Section 3. Episodes of deliberation as non-basic actions

Before proceeding with the main topic of this section I need to establish that when we engage in what I am calling 'episodes of deliberation', inner speech is often, perhaps typically, involved. As previously mentioned, I am using the term 'episode of deliberation' to cover a broad range of silent, conscious, mental activities including, for example: deliberating, reflecting, considering, working something out. The following scenarios illustrate what I have in mind when I use the term 'episode of deliberation':

a) You are seated in a restaurant and are handed a menu. Over the next 5 minutes you peruse the menu and come to a decision about what to order.
b) You have job offers from three organisations, each offering different opportunities and benefits. You have a few days to mull them over and decide which one to accept.
c) Your partner has received an exceptional job offer in another country and she is determined to accept it. You don't want to leave your own job to live in another country, but you don't want to separate from your partner either. What do you do?

If I try and imagine myself in these scenarios, I find it inconceivable that I would *not* engage in inner speech as I deliberated about what to do. In fact, I find it hard to understand what it could mean to *consciously* deliberate in these situations *without* engaging in inner speech; what would I be conscious *of* as I deliberated, if not sentences, or at least utterances of some kind, in my natural language? But maybe the claim that inner speech is *necessarily* involved in *all* episodes of conscious deliberation is too strong. After all, some people claim they don't experience inner speech at all; others claim they experience so-called

Unsymbolized Thinking – thinking without words (see Chapter 3 for details). So, I will restrict my claim to a more modest one: for at least for some people, episodes of deliberation of the kinds illustrated above just do, as a matter of empirical fact, typically involve the production of inner speech. The only argument I am offering for this claim is that I am one of them, and that the literature on inner speech is replete with others who say they are too.

In this section I argue that episodes of deliberation are a kind of action, albeit a kind of non-basic action.[20] I do that by arguing that they qualify as actions according to two of the three competing theories of action as set out in Section 2: (1) they are performed for a reason, and (2) they involve trying (and sometimes failing). (I address how episodes of deliberation qualify as actions according to the third theory of action below, in Section 5.) The first of these criteria – being performed for a reason – is the easiest to argue for. The idea that we engage in episodes of deliberation for a reason – to come to a decision or a conclusion, to make a choice, to solve a problem, to find an answer – is arguably a conceptual truth. There might be some close cousins of deliberation – such as pondering, wondering, speculating – which do not include in their essential nature the goal of coming to a conclusion. But we can exclude them from consideration for now and concern ourselves solely with episodes of deliberation which *by definition* have the goal of coming to a conclusion. In all three of the scenarios above, for example, the deliberation is purposeful – the subject intends to make a choice or come to a decision. The intention to choose, or make a decision, is what motivates and guides the entire episode, and like any other intention it can typically be explained by reference to a subject's beliefs and desires. The other two criteria for qualifying as an action – being under the control of the subject and being something the subject can try (and fail) to do - are more challenging. I will address the issue of trying here now, and address the issue of control in Section 5.

Mele (2009, pp. 18-37) makes an important distinction between 'trying to x' and 'trying to bring it about that I x'. In the Introduction I illustrated this idea with 'falling asleep', but this

[20] Or at least a kind of activity *comprised* of actions. The taxonomy of action terms, and the proper relation between an activity and an action, is explored in Chapter 6, Section 9.

might be thought irrelevant in the context of episodes of deliberation, since falling asleep is not typically thought of as a mental activity. What about the mental activity of trying to remember something? The example Mele uses is 'trying to think of seven animals beginning with 'g''. Suppose, in response to this challenge, a subject (I will follow Mele in calling her Gail) thinks 'goat'. There is nothing more to Gail's thinking 'goat', he says, than Gail becoming conscious of the word "goat". Becoming conscious of the word "goat" is something that happens to Gail, not something she does. And if that's true of goat then it's true of all the other six animals starting with 'g' that Gail thinks of. So, if thinking of seven animals starting with 'g' involves seven events which happen to Gail, one for each animal remembered, none of which is an action, then one might conclude that the entire episode of 'thinking of seven animals starting with 'g'' is an event which happens to Gail and is not an action. Indeed, that is what Mele *does* conclude – wrongly in my view, as I will later explain. He says:

> 1. Gail's thinking of 'goat' (for example) is not an action.
> 2. Gail's thinking of seven animal names starting with 'g' (her 7-ing) is not an action. (2009, p. 29)

However, there are two actions Gail does perform, according to Mele – and here I agree with him:

> 3. Gail's trying to bring it about that she 7-s is an intentional action.
> 4. Gail's bringing it about that she 7-s (her B7-ing) is an intentional action (ibid.)

Gail can try to bring it about that she 7-s by doing various things. To begin with, she might not have to do anything at all - 'goat', 'gorilla' and 'giraffe' just come to mind, effortlessly. Then she gets stuck. It might occur to her that she has only been thinking of mammals, and she has not thought about fish. She starts to focus on fish, and this gives her 'goldfish', 'guppy' and 'grouper'. Then she changes tack and focuses on insects and comes up with 'gnat'. The idea is that Gail doesn't just wait for more 'animals beginning with 'g'' to occur to her – she takes actions, mental actions, to 'try to bring it about that animals beginning

with ‘g’’ occur to her. And that fact makes it true that ‘trying to bring it about that she 7-s’ (her B7-ing) is an intentional action.

Mele makes the same point with a different example – his strategy for trying to remember what he had for dinner three nights ago:

> I have various memory-priming strategies for doing this. One is to ask myself (silently) what I had for dinner on that day and to keep my attention focused on that question. Normally, the answer does not come to mind straightaway, and I pursue the memory-priming strategy further by asking myself (silently) what other things I did on that day. An alternative strategy is to ask myself (silently) what I had for dinner last night and, if the answer comes to mind, work backward—which requires keeping my attention focused on my task (p. 19).

Notice that, in this example, Mele uses the expression “ask myself (silently)” three times. This reinforces my proposal that we should understand what Mele is doing when he tries to remember what he had for dinner three nights ago as performing a non-basic action, one he executes by performing basic actions, such as asking himself questions in inner speech. (I have more to say about the event of ‘remembering’ in Section 6 below.)

So much for asking himself questions, what about the answers? Why do I claim, *pace* Mele, that the events of remembering animals beginning with ‘g’, such as Gail’s thinking of ‘goat’ also qualify as actions? That is the subject of the next section. For now, I only want to note that, despite trying to remember what he had for dinner three nights ago by performing these basic actions, Mele might fail. (Consider the times you have tried to remember the name of someone, and despite your best efforts, you failed.) If Mele succeeds, however, we can say that his ‘bringing it about that he remembered’ was an intentional action. (This is equivalent to his claim 4. above, that ‘Gail's bringing it about that she 7-s (her B7-ing) is an intentional action’.)

To recap: this section argues that episodes of deliberation qualify as actions on two major theories of action: (1) because they are intentional, and (2) because they meet the criterion of being something a subject tries, and might fail, to do. The third possible criterion for

being an action – being under the control of the subject – will be addressed in Section 5. First, I want to say more about applying the basic/non-basic action distinction to mental action.

Section 4. Inner speech utterances as basic actions.

My claim is that episodes of deliberation, as I have defined them, should be understood in the same way as Mele's remembering what he had for dinner three nights ago. That is, they are non-basic actions which are executed by performing basic actions. In this section I argue that inner speech utterances produced as part of an episode of deliberation are basic actions. To argue for this involves two further claims. First, that speaking *overtly* is very often a basic action. Second, that there is no reason to think that inner speech is not also a basic action, just because it is speech which is not vocalised.

In arguing for my first claim – that speaking overtly is often a basic action – I am borrowing from Hornsby (2005). When a subject does something for a reason, according to Hornsby, she typically draws on knowledge of how to do it. For example, suppose what I do (for a reason) is travel to the university; I might do it *by* taking the bus; I take the bus *by* waiting at the bus stop, and I get to the bus stop *by* walking to it. Hornsby calls such knowledge "procedural"; we do one thing *by* doing another thing which we know how to do. There might be different buses I could catch to the university, and different routes by which I could walk to different bus stops; I might need to know a whole series of procedural facts to get something done. But the number of things a subject must know to get something done ultimately comes to an end; eventually there are things the subject just does 'directly'. We would not say, for example (except as some kind of joke) that the subject walked to the bus stop by *following the procedure* of putting one foot in front of the other. We would say, rather, that walking is something she is able *simply to do*.

Hornsby's central claim is that "the semantic knowledge exercised by people when they speak is practical knowledge" (p. 107).[21] As with walking, so with talking. Once we have

[21] In Chapter 7, Section 1, I address the worry some readers might have at this point that the human capacity for language is an instinct, not a skill.

learned a language, speaking is an action we are simply able to do. We have no more knowledge about the muscle movements involved in producing the words we utter when we speak them than we have about the muscle movements involved in moving our bodies when we walk. And the idea of *basicness* that's involved here is not confined to the *physical* things we do when we speak. We also have no procedural knowledge about the grammatical and syntactical principles which we put into practice when we speak – or if we do, we learn that knowledge *after* we have learned the language. Knowledge of it is certainly not a prerequisite to being a competent language user.

> Speakers can rely on the fact that producing meaningful things [i.e., overt utterances] is something that they are able to simply do. When a speaker says that p, there need be nothing such that she intentionally does it and says that p by doing it…. This elicits the force of the idea that we voice our thoughts directly (Hornsby, 2005, p. 118).

What it means to say that we "voice our thoughts directly", I suggest, is that we should not think that two actions are involved, the action of forming thoughts and the action of voicing them. Rather, we are performing one action under two descriptions: uttering the words and thinking the thought.[22]

It might be objected that, at least sometimes, we *choose* our words. I agree. But I suggest that on those occasions when we choose our words what we are really doing is making the choice of words the subject of deliberation. If this occurs during an overt episode of deliberation it should be understood as a silent deliberation within the context of a vocalised one – and as such it counts as another non-basic action in its own right. It is the non-basic action of trying to bring it about that the most appropriate word in the circumstances is selected. The basic action by which this non-basic action is performed is the action of saying the word that comes to mind. If a subject silently deliberates on which word to use in the circumstances, she can't *decide to select* a particular word to use, any more than she can *decide to remember* a particular animal beginning with 'g', for all the

[22] Another writer (than me) who makes explicit use of the basic/non-basic distinction from classic action theory in her analysis of mental action is Antonia Peacocke (2023a; 2023b). Her work is discussed in Chapter 5, Section 4.

same reasons. A better word will either present itself to her or it won't. If she is successful in 'trying to bring it about that she chooses the most appropriate word in the circumstances', then she will simply utter the word that comes to her in the circumstances, and her uttering that word is the basic action by which she performs the non-basic one of choosing which word to use. (The idea that deciding (in the sense of choosing) *just is* a kind of speech act is the subject of Chapter 6, Section 5.) In that sense, choosing the best word to use in the circumstances is like remembering an animal beginning with 'g' after prompting yourself by asking yourself a question. The answer to the question which prompted the remembering is no less intentional than the question which prompted it. (I argue for this in the next chapter.) Both earn their status of being intentional in virtue of the non-basic action they serve – trying to bring it about that one thinks of an animal beginning with 'g'.[23] And this is why Mele is wrong when he claims that his 1 and 2 are not actions. Gail's thinking of 'goat' *is* an action, albeit it a basic one, because her thinking it (and what comes to the same thing, her saying it to herself in inner speech) is the action by which she achieves her intention of trying to think of an animal beginning with 'g'. It follows that Gail's thinking of seven animal names starting with 'g' (her 7-ing) is seven actions, albeit basic ones, because her thinking of each one (and saying them to herself in inner speech) is the series of actions by which she achieves her intention of trying to think of seven animals beginning with 'g'. (The next chapter exploits this idea to argue for the idea that reasoning, with its characteristic question-and-answer format is wholly intentional, and not merely partially so, as some have suggested.)

So much for my first claim - that speaking overtly is a basic action. What about the second – that there is no reason to think that inner speech, unlike overt speech, is not a basic action, just because it is speech which is not vocalised? In the case of overt speech, the reason a speaker has for speaking is usually to communicate or express something to someone else. But as I have already argued, when a subject is engaged in an episode of silent deliberation, she too has a reason for speaking (albeit to herself) – her reason is to try to bring it about that she x-s, where x is to reach a conclusion, solve a problem, come to a decision, and so

[23] If you think this is a bit quick, and that remembering is an automatic process which can never qualify as an action, see Section 6 (Memory) below.

on. In the case of overt speech, the subject has an intention to achieve a goal (or goals) and fulfils that goal by (performing the basic action of) speaking out loud; in the case of inner speech the subject also has an intention and fulfils that goal by (performing the basic action of) speaking silently. Once again, the fact that the motor planning system is engaged in both kinds of utterance (see Chapter 2 for details) further undermines any reason for thinking that inner speech utterances are not basic just because they are not vocalised.

Section 5. Control

I still need to offer an argument for how an episode of deliberation meets the criterion of being an action, albeit a non-basic action, by being under the control of the subject; I will do that in this section. Recall how, in Section 2.2 above, I characterised Frankfurt's (1978) definition of control as follows: 'for something to be an action it must be possible for the subject to adjust what they are doing while they are doing it. To qualify, such adjustments must be attributable to the subject, and not to some automatic mechanism possessed by the subject'. If control is characterised this way, then episodes of deliberation are clearly under the subject's control. As the action unfolds – as the subject continues to try to bring it about that she x-es – she can adjust the way she performs it. For example, as we saw in the example of a subject 'trying to bring it about that she thinks of animals beginning with 'g'', she can focus first on mammals, then change to focusing on fish, then to insects, and so on.

Of course, the basic actions by which the subject performs her non-basic actions are attributable to automatic systems – whatever systems are responsible for selecting and organising into meaningful utterances the words which comprise her inner speech. But they still qualify as under the control of the subject, precisely because they are performed in the service of the non-basic action. To demand a more stringent criterion of control would be to fail to appreciate that the execution of *any* intentional action rests, ultimately, on the subject being able to simply do something, because she knows how. We should not insist that every constituent of an intended action is itself intended, on pain of a regress – as noted in Section 3 above. So, we should not insist that every basic action, when it is a

constituent of a non-basic action, is itself intended.[24] My raising my arm because I intend to attract the waiter counts as an action under my control even though I don't also form an intention to raise it, i.e., an action distinct from my intention to attract the waiter by raising my arm. (Were my arm to rise in the same way in the absence of an intention to attract the waiter – in a spasm, say – my arm rising would not count as an intentional action.) Similarly, if I produce an inner speech utterance because I intend to try to bring it about that I decide something, the utterance counts as an action under my control even though I don't first form an intention to make the utterance, i.e., an action distinct from my intention to try to bring it about that I decide something. On my account, the production of an inner speech utterance stands in the same relation to an intentional *mental* non-basic action (such as trying to decide something) as raising my arm stands to the intentional *bodily* non-basic action of, say, trying to attract the waiter by raising my arm. (The case of deciding something receives further analysis in Chapter 6, Section 5.)

Of course, there are important differences too. The production of an inner speech utterance involves cognitive resources, events, and processes which are different from those involved in the action of raising an arm. But that doesn't make any difference to the question of control, since all these cognitive resources, events and processes are managed by sub-personal systems; they are as much, and as little, under a subject's control as the bodily systems involved in raising an arm. There is another important difference, related to the first. When a subject raises her arm with the intention of attracting the attention of the waiter, whether she succeeds or fails depends on something *external* to her person – the waiter: will he notice her, or not? By contrast, when a subject utters something in inner speech with the intention of trying to bring it about that she comes to a decision, whether she succeeds or fails depends on something *internal* to her person – her sub-personal cognitive resources: will they generate an utterance that gets her closer to achieving her

[24] Jenkins (2021) makes a related argument for the claim that *reasoning* is a kind of action. If we insist that each "sub-action" involved in a chain of reasoning, such as its constituent judgments and inferences, must be under the agent's control, we generate a dilemma. Either the agent controls each sub-action by performing a distinct prior action, in which case we are off on a regress, or the agent controls each sub-action *without doing anything else*, in which case action is just mysterious. "The mistake is to think that extended actions [such as chains of reasoning] must always be made up of constituent sub-actions which can be seen as such *independently from their place in more extended action*" (p. 16, emphasis in the original). I focus on reasoning in the next chapter.

intention, or not? Despite this difference, the two cases have something in common: success or failure is dependent on a response from something which is not under the subject's personal control.

Section 6. Memory

The argument I made against Mele (2009) in the previous section might seem a bit quick. An objection might go something like this.

> Objection. Events qualify as actions only if the subject has some control over them, but remembering something is an automatic event over which the subject has no control. Certainly, a subject can stimulate a remembering, by asking themselves a question or imagining an image, and that activity qualifies as an action. But once that action of stimulation has been performed all the subject can do is wait and hope; the memory will either come or it won't. A subject can't initiate a remembering, nor intervene to stop or change it when it starts. Remembering is something that happens to you, not something you actively do. Remembering is not an action. Furthermore, all these points apply equally well to Mele's case of trying to think of 7 animals beginning 'g'. That case can be reworded in terms of trying to *remember* 7 animals beginning with 'g' (or, if you prefer, trying to *remember* which animal's names start with the letter 'g'). That being so, Mele is right when he claims:
>
> 1. Gail's thinking of 'goat' (for example) is not an action, and
> 2. Gail's thinking of seven animal names starting with 'g' (her 7-ing) is not an action. (p. 29).
>
> And if he's right about that then there are aspects of thinking which are not under a subject's control, and therefore we should conclude that thinking is only partially intentional, and not wholly so, contrary to the claim.

In the next chapter I will consider, and reject, a different argument (one based on 'dual process' theories of cognition) that thinking (specifically reasoning) is only partially intentional and not wholly so. But for now, how do I respond to this objection from the automatic nature of memory?

Happily for me, much of the heavy lifting involved in addressing this objection has already been done (persuasively, in my view) by Seth Goldwasser (2022), whose paper explicitly

takes issue with the arguments of Strawson (2003) and Mele (2009). According to Goldwasser,

> Where Strawson and Mele go wrong is in their assumption that the presence of automaticity in tokens of a type of mental event [such as remembering] is sufficient evidence for the claim that the relevant event-type is fully automatic and, thus, not an action (p. 4).

Goldwasser points out that remembering comes in two forms: He distinguishes between what he calls 'fluid' remembering and 'recurrent or intrusive' remembering.[25,26] Fluid remembering is the most common kind, he says; the content of fluid remembering is "unsurprising", and is accompanied by feelings of "expectation, knowing, familiarity, resemblance, and a sense of guidance". By contrast, as the name suggests, recurrent or intrusive remembering "feels uncontrolled" – intrusive memories occur out of the blue.

(Note how this distinction is very like the distinction I draw in this chapter between 'active' and 'reactive' inner speech. It seems to be a feature of several of our cognitive functions that they can take either an active or a passive form; examples include inner speaking, remembering, imagining, picturing, hearing a tune, seeing a face. If we allow that the term 'thinking' covers all these activities, we should allow that 'thoughts' come in two forms: those that are intentional actions ($thoughts_a$) and those that are reactive ($thoughts_r$). I argue for this idea in Chapter 8.)

Goldwasser's main argument is that fluid remembering is a skill, and that skilled activities are actions, and therefore fluid remembering is an action. To make his argument he focuses on two criteria for an activity qualifying as a skill that are common ground in the extensive philosophical literature on the subject of skill:

[25] Goldwasser uses the term 'memory' rather than 'remembering'. In what follows I prefer to use the term 'remembering', to keep the focus on memory as a type of event, rather than a phenomenon.
[26] Goldwasser explains that his distinction does not map neatly on what is sometimes called the distinction between 'voluntary' and 'involuntary' remembering. He provides a detailed account of the differences (p. 6, FN 7) but the distinction is not relevant to the points I want to make here, so I will ignore it.

> "[S]kills are, at bottom, things we can learn to do with practice and are such that their exercise admits of attributions of excellence…. Being learnable with practice and admitting of attributions of excellence define skill precisely because they illuminate the nature of control exerted in the exercise of skill—control is gained and increased over time (p. 7).

Goldwasser draws on the empirical literature on the training of working and episodic memory to make his point that remembering can be improved with training (although it would take me too far from my subject to review the evidence he adduces for that claim here). He then argues convincingly that skill is at least partly defined by its normative dimension, and that this dimension is present in remembering as in other activities. In the extreme case:

> Mnemonists are experts at remembering…. They participate in competitions in which contestants are to remember vast quantities of digits, words, poems, names, faces, playing cards, etc. in short periods of time and usually in some specific order (p. 12).

Once again, I don't want to get distracted by dwelling on the detailed empirical evidence for this claim.

But I do want to guard against giving the impression that the training involved in developing the skill of remembering is something unique to those individuals – arguably an eccentric minority – who participate in mnemonic competitions. Goldwasser distinguishes between two kinds of automaticity. First, the kind for which training and practice is irrelevant – for example the kind involved in a reflex movement – (he calls this 'untrained automaticity'), and second, the kind which occurs when a subject is engaged in an activity she knows how to do and which she performs without conscious intervention – such as making a routine commute – (he calls this 'routine automaticity'). Routine automaticity, he claims, is the more pervasive kind, involved in nearly all our everyday activities, and is the product of practice. Not only that, but its use contributes to another feature of a skilled activity – the selective use of attention to exercise control of specific aspects of the activity, while allowing other aspects to unfold without attention. The point about routine automaticity is

that it is pervasive in even quite ordinary subjects doing ordinary things, and that includes the activity of remembering something:

> [D]uring normal development, human beings start off with some basic cognitive capacities whose exercise may well result in token remembering that is outside of the control of the remembering subject [i.e., it involves untrained automaticity]…. Nonetheless, evidence from developmental psychology suggests that remembering is trained and practiced within the first years of life with guardians and peers and that it improves as a function of the quantity and quality of that training and practice…. An important part of that training and practice is getting the agent to attend to narrative aspects of her past. As ordering events in narrative structures becomes routine, the agent becomes more able to selectively attend to specific events, their relations, and their details. Eventually, she can reliably recall and elaborate entire sequences unaided. So, even if mnemonic behavior initially exhibits untrained automaticity, this alone does not show that there is no control over remembering (pp. 17-18).

In short, since fluid remembering (the most common kind) involves routine automaticity, and since routine automaticity involves practice and skill, and since demonstrating skill also demonstrates control, and since control is a characteristic of an action, it follows that remembering qualifies as a kind of action. If that's right, then the objection which is the subject of this section – that remembering can never qualify as an action because it is always beyond the subject's control – has been answered.

I want to end this section by noting that Goldwasser makes a point about remembering something which is similar to the point I made above, about silently uttering something (in response to a question we silently ask ourselves): In neither case should we demand a higher level of control over these events than the level of control we demand over perfectly ordinary bodily actions. I have much more to say about the similarity between thinking in words and (other) skilled bodily actions – that is the subject of Chapter 7.

Section 7. Summary.

Some inner speech utterances are reactive: they are spontaneous, they require no effort, and we are not in control of their occurring. These inner speech utterances fail to meet the

standard criteria for qualifying as intentional actions. But some inner speech utterances are genuine actions, performed deliberately, effortfully and with as much control as any other intentional action. For example, when we engage in an episode of deliberation, we are performing the non-basic action of trying to bring it about that we achieve some cognitive goal – coming to a decision, reaching a conclusion, solving a problem, etc. The action of trying to bring this goal about is achieved by performing basic actions – making inner speech utterances. An inner speech utterance, when made in this context, stands in the same relation to the intentional action of trying to (say) reach a decision, as raising an arm has to the intentional action of trying to (say) attract a waiter. That is, it stands in the relation of a basic action to a non-basic action.

Chapter 5. Is reasoning intentional?

Abstract

If reasoning is an intentional activity, as many suppose, what are the intentional actions a subject performs when she does it? One plausible answer is that she talks to herself - call this the *inner speech thesis*. On this account the subject engages in a silent dialogue with herself comprising a series of questions and answers. According to Frankish (2018) the questions a subject asks herself are instances of Type 2 processing and are intentional, while the answers a subject responds with are instances of Type 1 processing and are not intentional. He concludes that reasoning is only *partially* intentional. I argue against this view. My account, by contrast, has it that the questions and the answers are both intentional and conclude that reasoning is therefore wholly intentional. I also offer a novel account of the nature of the inner speech dialogue involved in reasoning, according to which its function and its phenomenology are characteristically *exploratory*. I argue that the special phenomenology of exploratory actions encourages the mistaken idea that some of the inner speech utterances involved in reasoning are not genuine actions.

Section 1. Introduction.

On the face of it, it might appear obvious that reasoning is intentional (in the ordinary sense of 'done with intention/intended'). After all: (a) we reason with the intention of working something out, (b) we do it consciously, and (c) it takes us time and effort.[27] I take it these are all hallmarks of an intentional activity. The difficulty arises when we try and specify exactly what it is we do – what actions we take or make – when we do it. If, as is commonly supposed, the "primitive sign" of intending to *A* is trying to *A,*[28] and if my intention when I engage in reasoning is to try to bring it about that I decide what to do, or what to think,[29] I should be able to describe what I actually do (what actions I in fact take) that are

[27] I acknowledge that 'reasoning' is sometimes used to refer to non-conscious cognitive activity, but this is not how I am using the term here.

[28] The reference is to Anscombe, *Intention*, p.68, where she wrote "The primitive sign of wanting is *trying to get*". Bratman (1987, p. 121), quoting Anscombe, makes the point that this is also true of *intending*.

[29] My argument refers equally to practical and theoretical reasoning. Henceforth I won't bother to make this distinction.

constitutive of my trying. If we think that reasoning is intentional, we should be able to answer the question: 'What is it we intentionally *do* when we engage in reasoning?'

One controversial response to this question is that we don't, strictly speaking, do anything at all. According to Strawson (2003) there is plenty of mental *activity* involved in reasoning, but this does not qualify as *action*:[30]

> The movement of the natural causality of reason (practical reason in this case) to its conclusion in choice or decision is lived (by some) as action when it is really just reflex; distinctively rational reflex, to be sure, but not in any case a matter of action. (p. 244)

Few would deny that there is frequently mental *activity* without action, but it is counter-intuitive to suggest that the mental activity involved in reasoning is an instance of it. Even Strawson acknowledges that some genuine action(s) might be *involved* in reasoning:

> If the issue [i.e., the problem one is reasoning about] is a difficult one, then there may be a distinct, and distinctive, phenomenon of setting one's mind at the problem, and this phenomenon, I think, may well be a matter of action. It may involve rapidly and silently imaging key words or sentences to one-self, rehearsing inferential transitions, refreshing images of a scene, and these acts of *priming*, which may be regularly repeated once things are under way, are likely to be fully fledged actions (p. 236, original emphasis).[31]

Note Strawson's suggestion that reasoning involves "silently imaging key words or sentences to one-self". This idea, I submit, is on the right tracks, and points the way to a better response to our question ('What is it we intentionally do when we engage in reasoning?') – the one that I promote in this paper. That response is the claim that the action we take when we reason is to talk to ourselves; call this the *inner speech thesis*.

[30] Strawson explicitly uses 'action' here to mean 'intentional action' (p. 234).
[31] Wu (2016, pp. 247-9) claims that Strawson's position is inconsistent. He argues, convincingly in my view, that Strawson's arguments against mental action apply just as much to these so-called "acts of priming" as they do to the "reflex" of reasoning itself. I won't rehearse his arguments for this claim here.

The *inner speech thesis* involves bringing together two ideas, each of which enjoy significant independent support. The first idea is that there are at least some types of thinking which necessarily involve language,[32] and that reasoning – intentionally trying to work out what to do or what to think – is one of those kind of thinking. (This way of putting things leaves open the possibility that there are other types of thinking which do not involve language (Camp, 2007; 2009), or even some which lack any kind of symbolic representation whatsoever (see Chapter 3).) The second idea is that at least some inner speech utterances qualify as intentional actions (Jones & Fernyhough, 2007; Martínez-Manrique & Vicente, 2015). (This way of putting things leaves open the possibility that there are some tokens of inner speech utterances which do *not* qualify as actions – something I argued for in Chapter 4.) The *inner speech thesis* combines these two ideas to make the claim that reasoning is an intentional activity, and that the inner speech utterances involved in reasoning qualify as intentional actions.

But even if the *inner speech thesis* is on the right tracks, there is a worry that it might not be enough to support an argument for the *full* intentionality of reasoning, as I will now explain. When we reflect on the nature of the inner speech utterances which comprise an episode of reasoning, they typically have a question-and-answer aspect to them. When I want to solve a problem by reasoning about it, I typically ask myself a question and then respond to what I've just asked myself; this response might prompt me to ask another question, and so on.[33] Here is a simple example of an episode of reasoning which might be conducted in inner speech:

Q: Should I take an umbrella with me today?
A: Let's have a look out of the window.
Q: Does it look like rain today?

[32] To be clear, this is not a modal claim. The claim "reasoning necessarily involves language" refers to nomological necessity.
[33] This idea has an ancient history. According to Plato: "It seems to me that the soul when it thinks is simply carrying on a discussion in which it asks itself questions and answers them itself, affirms and denies." (Theatetus 190a, Cooper & Hutchinson, 1997). For more recent theorising about the dialogical nature of thinking see Fernyhough (1996; 2008).

A: No, I will leave the umbrella behind.[34]

On the face of it, asking myself the question "Should I take an umbrella with me today?" is an intentional action – I do it consciously and for a reason, namely, to help me decide whether to take my umbrella. The worry is that *answering* my own question ("Let's have a look out of the window") has the appearance of an automatic response, a mere *reaction* to the question, and reactions, arguably, are a kind of reflex, and reflexes are not intentional actions. (Recall the distinction made in Chapter 4 between inner speech utterances which qualify as genuine (intentional) actions, and those which are merely *reactive*; the worry is that the answers to our own questions are instances of the latter, not the former.) Considerations along these lines have led Keith Frankish (2018) to claim that "intentional reasoning is not *wholly* intentional, but guided and mediated by autonomous reasoning" (p. 231, original emphasis). (Note how Frankish's phrase "autonomous reasoning" echoes Strawson's description of reasoning as "rational reflex".) Frankish's *version* of the inner speech thesis denies the claim that reasoning is wholly intentional.

Why does this matter? Whether reasoning is wholly or only partially intentional has a bearing on our sense of ourselves as agents. The actions we feel most responsible for are the ones we take deliberately after careful consideration – reasoning about what to do or what to think – and then reaching a conclusion or coming to a decision. If it turns out that the process of deliberation is itself only *partially* intentional, then it might seem to follow that we are only partially responsible for the actions and the beliefs which deliberation leads to. On the other hand, if there is a version of the *inner speech thesis* which demonstrates that reasoning is wholly intentional, then this particular challenge to our sense of agency (there are others) is neutralised. Providing such a version is the purpose of this chapter.

The structure of the chapter is as follows. I review and criticise Frankish's (2018) argument for the partial intentionality of reasoning (Section 2). I invoke (as in Chapter 4) the

[34] There is good evidence that a lot of inner speech is often highly abbreviated and does not consist of fully formed sentences, as in this idealised illustration (Alderson-Day & Fernyhough, 2015). This fact does not affect the arguments in this chapter.

traditional distinction between basic and non-basic actions and show how it can be applied to the inner speech utterances which constitute an episode of reasoning. I demonstrate why these utterances, whether they take the form of a question or of an answer, should be understood both as (a) basic and (b) intentional (Section 3). If that's right, then it follows that reasoning is wholly intentional. Finally, I offer a novel account of the nature of the inner speech dialogue involved in reasoning, according to which its function and its phenomenology are characteristically *exploratory.* I argue that the exploratory nature of reasoning gives rise to the mistaken impression that some of the utterances involved are not intentional actions (Section 4). Section 5 concludes.

Section 2. Is reasoning only *partially* intentional?

In this section I will consider an argument from Frankish (2018) which concludes that reasoning is only partially intentional. His claim is that, when we *ask* ourselves a question as part of an episode of reasoning, the action of asking it is intentional because it (the action) is under our control, but that when we *answer* our own question the answer comes to us automatically, like a reflex. Since reflexes are not under our control the answer does not qualify as an intentional action, so on this account the entire episode of reasoning is only partially intentional. I will argue that this analysis is flawed.

Frankish is a longstanding proponent of 'dual-process' theories of thinking, according to which two distinct processes, or types of processing, compete for control of our behaviour (Evans, 2003; Evans & Stanovitch, 2013; Frankish, 2009, 2010; Kahneman, 2011).[35] The first type of processing (Type 1) is characterised as fast, automatic, and non-conscious. (Also, as "low-effort, high-capacity, parallel, contextualized, associative, biased, undemanding of working memory, shaped by biology and personal experience, and independent of cognitive capacity" (Frankish, 2018, p. 229)). The second type of processing (Type 2) is characterised as slow, controlled, and conscious. (Also, as "effortful, low-capacity, serial, decontextualized, rule-governed, normative, demanding of working memory, shaped by culture and tuition, and correlated with individual cognitive capacity" (*ibid*)). In a nutshell,

[35] In this chapter I don't take any position on the Dual Process model. However, for a recent expression of scepticism about it see Borg (2022).

Frankish's version of dual-process theory – as it applies to reasoning – is that the key difference between Type 1 and Type 2 reasoning is that the latter qualifies as intentional, and the former does not.

According to Frankish, "intentional reasoning" involves a cyclical process used to break a complex problem down into more simple problems. He provides an example of how this is meant to work:

> Suppose I have been invited to a party with colleagues from work. I don't find myself strongly disposed to respond one way or the other, but I need to give an answer, so I engage in intentional reasoning…. I need some way of breaking down the problem. I begin by questioning myself to try to elicit an evaluative reaction or piece of relevant information, asking, 'Do I really want to go?', 'What will it be like?', or something similar. I hear my own utterance, my language comprehension system interprets it, and its content is globally broadcast to other mental subsystems. My mindreading faculty interprets me as requesting information about the party or an evaluation of it, and further autonomous processes throw up the prediction, based on experience, that Henry will be there. This message is selected for expression … and I utter the words: 'Henry will probably be there'. Again, this utterance is heard and interpreted. Though it does not have the form of a question, in the context it is interpreted as posing further subproblems: Do I want to meet Henry? What will happen when I meet him? Again, a response is selected and articulated: 'He'll want to talk about the budget cuts'. This in turn is heard and interpreted as posing the problem of whether I want to talk about the budget cuts. My affective response—let us suppose—is strongly negative, and I conclude by uttering, 'I can't face that; I won't go' (p. 234)

On this account of reasoning, my asking a question is an example of Type 2 processing and qualifies as an intentional action on the grounds that it has a belief-desire explanation (a desire to decide whether to accept the invitation and a belief that asking the question will help me decide). On the other hand, my responding to that question is an example of Type 1 processing and does not qualify as an action, on the grounds that I have no control over the production of the response – it just happens. Frankish concludes that "intentional reasoning is not *wholly* intentional, because it is guided and mediated by autonomous reasoning" (p. 231, emphasis in the original).

I should mention that despite my differences with Frankish on this point about the intentionality of reasoning, I agree with much of what he says in his (2018) chapter. I agree with the claim that one of the functions of inner speech is to provide a format, a representational medium, for conscious thinking (what he calls "the format view"), and also that inner speech is an activity, with many functions, continuous with those of overt speech (what he calls "the activity view). I also agree with Frankish that these views are compatible.[36] And I agree with the claim that inner speech typically does more than merely function as a channel of communication, making an existing thought conscious to the subject. Rather, as he puts it, "it must enable a distinctive kind of thought, which is conscious. It must *make a conscious thought*, rather than just *making thought conscious*" (p. 228). (This idea is put to work in Chapter 6.) So, what's wrong with his account?

Frankish assumes that if an action is intentional then it is under control, and if an action is automatic then it is not intentional. The rest of this section is devoted to showing how those assumptions are problematic. He writes: "In effect, this proposal [that it is the distinction between autonomous and intentional reasoning that lies at the core of dual-process theories] takes the property of being *controlled* as the defining feature of Type 2 reasoning, specifying the relevant form of *control* as intentional" (p. 231, emphasis added). He doesn't say what he means by the term 'intentional control', but having made the connection between an action being intentional and an action being under control he proceeds to help himself to some of the folk-psychological ideas associated with the idea of control. This, I submit, is problematic.

For example, one idea connected with the idea of control is that of being able to decide. He writes that utterances are intentional when they have "belief-desire explanations. I *decide* to perform each step [of the reasoning process] because I want to solve the overall problem and believe that the step is part of a procedure for solving it" (p. 231, emphasis added). But deciding is typically something one does *after* one has deliberated about what to do, so using the term to explain how deliberation itself works generates a regress. I don't "decide"

[36] See his (2018), sections 8.2 and 8.3, for details.

to make *any* of my silent utterances, if what is meant by that is that I first of all think about making them (by reasoning about it) and then decide (i.e. choose) to make them. If I decided *that* way, then I would have to make *other* inner speech utterances, ones connected to my reasoning about whether to decide, and so on, and this generates a regress. (The phenomenon of 'deciding something' gets a fuller treatment in Chapter 6, Section 5.)

Another example. Take this discussion of the "selectional" effect each inner speech utterance has on the direction an episode of reasoning takes:

> Different *choices* of utterance may take the process in completely different directions. By its nature as a *selective*, serial process, intentional reasoning carves out a specific route through the deliberative territory, which we would not otherwise have taken (p. 236, emphasis added).

But in what sense does the subject "choose" or "select" the utterances involved in reasoning? It can't be that she deliberates about which utterance to make before choosing one, since that would involve further inner speech utterances and, once again, a regress follows. In fact, this is as true of choosing/selecting the *questions* as of choosing/selecting the *answers*. Consider his own example of being invited to a party with colleagues from work.

> I begin by questioning myself to try to elicit an evaluative reaction or piece of relevant information, asking, 'Do I really want to go?', 'What will it be like?', or something similar.

In what sense, I want to know, is this first step – that of asking himself a question – under his control? Is he even in control of whether he responds to the invitation with an utterance in the form of a *question*, rather than with some other kind of utterance? For example, he might respond to the arrival of the invitation with a *statement*, such as, "I suppose I will have to respond to this invitation". Frankish himself acknowledges this possibility; he writes "…instead of starting off with a question, I might try to imagine the party" p. 235. But he doesn't address the question of how it comes about that he does one thing rather than the

other. He says it is under his control, a matter of choice, but that implies that he *decides* that asking himself a question would help him decide whether to accept the invitation, rather than deciding that, say, an image of the party would help him decide whether to accept the invitation. Once again, that way of understanding the matter leads to a regress of deciding.

A third example. Suppose that, when Frankish starts considering whether to accept the party invitation, an utterance in the form of a question comes to him rather than some other kind of utterance. What is the content of the question that comes to him? Frankish offers some candidates: "'Do I really want to go?', 'What will it be like?', or something similar." Once again, he clearly doesn't *choose* which question to ask himself, in the sense of making a selection between alternative candidates. If the candidate questions were somehow before him, consciously available for him to choose between, they would be already present, and therefore not of his choosing. Strawson (2003) is right about this when he writes, "…if [the content of a thought] is already there to be considered and adopted it must already have 'just come' at some previous time in order to be so available" (p. 235). I agree: If the content had 'just come' then it can't have been chosen (in the ordinary sense of that word).

In short, by making the implicit assumption that 'being intentional' entails 'being under control' Frankish licences himself to make claims which trade on a sense of 'control' which takes for granted the very thing he is attempting to explain, namely, the nature of reasoning. If, by hypothesis, we reason in order to come to a decision, he can't involve the action of deciding to explain the activity of reasoning. Specifically, Frankish is trading on the gerundive form of the verb, whereby the action of asking a question is treated as a noun - 'a choosing', 'a selecting' - ignoring the fact that the *act* of choosing or selecting is typically understood as the *conclusion* to an episode of deliberation. This allows him to make a false distinction between the nature of the questions (in an episode of reasoning) and the nature of the answers, claiming that the former are intentional and the latter are not. This allows him to conclude, falsely, that reasoning is only partially intentional.

Section 3. Why reasoning is *wholly* intentional

In this section I offer an alternative account of the nature of reasoning, one based on the well-known distinction between basic actions and non-basic actions. Whereas, in Chapter 4, I used that distinction to argue for the intentional nature of at least some of our inner speech utterances (while also acknowledging the reactive nature of others), in this chapter I apply that distinction specifically to the action of making the utterances involved in episodes of reasoning – the questions and the answers – and to argue that both kinds of utterance are basic. Since, on this view of basic action, a basic action inherits its status as intentional from the non-basic action it serves – in this case the intention of trying to bring it about that the subject reaches a decision (or comes to a conclusion or solves a problem) – it follows that both kinds of utterance (question and answer) are intentional. And from this it follows that reasoning itself is wholly intentional.

Sandis (2010) writes (rather unhelpfully, I must admit) that when it comes to the distinction between basic and non-basic actions "there are multitudes of equally legitimate conceptions of what counts as basic" (p. 13). Here is not the place to argue for the merits of one conception of basic action over the others. For the purposes of this chapter (as in the previous one) I will follow those theorists who appeal to an agent's *practical knowledge* to determine which parts of her actions are basic, and which are non-basic. According to this view, an action is basic if knowing how to do it does not depend on the agent's knowing how to do some other action (Amaya, 2017; Hornsby, 2013; Sandis, 2010). Roughly, non-basic actions are the actions a subject performs *by* doing something else, and basic actions are actions a subject does *directly*, without (intentionally) doing anything else. For example: suppose my intention is to travel to the university. This is a non-basic action because I need to perform other actions to execute it. I do it *by* taking the bus; I take the bus *by* going to the bus stop; and I get to the bus stop *by* walking to it. Hornsby (2005) calls the knowledge we possess which allows us to perform these actions "procedural"; we do one thing *by* doing another thing which we know how to do. But the number of things a subject must know to get something done ultimately comes to an end; eventually there are things the subject just does *directly*. We would not say, for example (except as some kind of joke) that the subject walked to the bus stop by *following the procedure* of putting one foot in front of

the other. Walking is something she is able *simply to do* – she doesn't need to follow a procedure to do it, and to form a separate intention to take each step. The fact that basic actions are things we can simply do does not impugn their status as intentional actions. They are intentional by virtue of the fact that they are performed in the service of non-basic intentional actions. When walking to the bus-stop, each step I take qualifies as intentional, even though I don't form a distinct intention to take each step.

This might seem a bit quick.[37] Does every basic action executed 'in the service of' a non-basic action qualify as an intentional action performed by me? Might there be some *events* or *happenings* involving me which, on the one hand, are arguably in the service of an intentional action of mine, but which, on the other hand, do not qualify as genuine actions of mine? Suppose, for example, I was hooked up to a system which, once it had identified that I had formed an intention to walk to the bus stop, took over from me the business of making the necessary bodily movements. (Perhaps I am wearing some kind of robotic suit which operates my limbs; once it 'knows' I want to walk to the bus stop it just walks me there.[38]) We might want to say in this case that, while my movements are being performed in the service of my intention (to walk to the bus stop), they are not being performed *by me*; they are events happening to me rather than actions I am performing. This possibility might seem far-fetched in the case of bodily movements, but it becomes a much more pressing worry when, as below, I apply the basic/non-basic distinction to inner speech utterances. I will address this worry at the end of the section.

With the basic/non-basic distinction in place, the question now is: How does it apply to the utterances involved in episodes of reasoning. That question can be divided into two parts. 1) *Asking* a question is plausibly intentional – a subject doesn't typically ask a question without intending to[39] – but on what grounds does it qualify as basic? 2) *Answering* a question is plausibly basic – the subject just does it directly – but on what grounds does it qualify as intentional? I will address each question in turn.

[37] Thanks to Alex Grzankowski for pressing me on this point.
[38] Something, perhaps, like the devices referred to in Aflalo *et al* (2022): "A neural prosthesis [which] translates the motor intentions of paralyzed individuals into control signals for assistive devices" (p. 2051).
[39] Spontaneous exclamations are possible exceptions: "What the …?!"

How can *asking* a question be basic? Recall that an action is basic if it is something a subject does directly, meaning that she does it without intentionally doing something else. I argued in Chapter 4 that that can be as true of the action of (overt) speaking as of any other bodily action. For example, when I (intend to) attract the waiter *by* raising my arm, I raise my arm directly, without first forming *another* intention which I fulfil by raising my arm. Similarly, when I (intend to) communicate my desire to know something *by* asking a question, I ask the question directly, without first forming another intention which I fulfil by asking the question. Of course, I need to have learned to speak a language before I can do this, but once I have that skill, I exercise it directly. In fact, if you asked me *how* I do it I wouldn't be able to tell you.

> Speakers can rely on the fact that producing meaningful things [i.e., external utterances] is something that they are able to simply do. When a speaker says that p, there need be nothing such that she intentionally does it and says that p by doing it.... This elicits the force of the idea that we voice our thoughts directly (Hornsby, 2005, p. 118).[40]

A possible objection to the claim that asking a question is basic is that, while it might be true of *overt* utterances, it is not true of inner speech utterances. But, as argued for in the previous chapter, there is ample evidence to suggest that there are no deep differences between overt and covert speech, or at least none that would justify both (a) endorsing the idea that asking a question overtly was basic, while at the same time (b) denying that asking a question in inner speech was basic, and insisting instead that it was non-basic. In short, if asking a question overtly qualifies as a basic action, there is every reason to suppose that asking a question covertly does too.

A further possible objection to the idea that asking a question is a basic action is that, at least sometimes, we choose our words before we speak. According to this objection, asking the question is not done directly, but only *by* doing something else first, namely choosing the words to use. I have addressed this objection too in Chapter 4. To recap, my response

[40] For the view that the action of phoneme selection, and not the action of making a complete utterance, is the basic action involved in speaking, see Amaya (2016).

to this objection is similar to my criticism (see section 2) of Frankish's illicit use of the concept of 'choosing'. On my account, on those occasions when we consciously choose our words during an episode of deliberation, the choosing of the words involves a distinct act of deliberation (albeit one that is nested inside the overall episode of deliberation). That act of deliberation is the non-basic action of trying to bring it about that we choose the most appropriate words in the circumstances, and the basic action by which this non-basic action is executed is the action of saying the words which come to mind *directly*.

How can responding to a question be *intentional*? Recall that a basic action is intentional if it is performed in the service of a non-basic action. I submit that when a subject responds to a question which she has asked herself, the action of uttering that response is just as much in the service of the non-basic action as the action of asking the question to which it is a response. For example, if I am trying to decide whether to take an umbrella with me and I respond to my own question ("Should I take an umbrella with me today?") with "Let's have a look out of the window", that utterance is as much in the service of the non-basic action of trying to decide whether to take an umbrella as the question it is a response to. Unlike the so-called 'reactive' inner speech utterances I considered in Chapter 4, which I characterised as 'spontaneous' and 'occurring for no apparent reason' (where that means 'being explicable by reference to beliefs and desires'), this response is neither of those things. The response is a constituent of an *activity* – the activity of trying to decide something; it is not spontaneous in the sense of occurring for no reason at all, and it is entirely explicable in the context of that activity by reference to beliefs and desires. (I say more about the difference between actions and activities in Chapter 6, Section 9.)

It is time to return to a worry I raised earlier. Does every basic action executed 'in the service of' a non-basic action qualify as an intentional action performed by me? Might there be some *events* or *happenings* involving me which, on the one hand, are arguably in the service of an intentional action of mine, but which, on the other hand, do not qualify as genuine actions of mine? This worry is particularly acute in the case of the answers we give to our own inner speech questions, for two reasons. The first reason is that the phenomenology of reasoning seems to support the intuition that the answers, in some sense, *just happen*. I will address that worry in Section 4. The second reason is that nearly

everyone has had the experience of an inner speech utterance which just "pops into my head", unbidden and without any effort on their part. As explained in Chapter 4, Gregory (2020) has argued persuasively that these inner speech utterances should be understood as *reactive* – reactive, that is, to whatever circumstances the subject happens to be in at the time they occur – and that these instances don't qualify as genuine actions; it seems we could not *prevent* this phenomenon from occurring even if we wanted to. So, the worry is that maybe the answers we give to our own questions (during an episode of reasoning) are, like instances of reactive inner speech, not genuine actions at all, *despite the fact* that they are 'in the service of' the intention which motivates the reasoning. If that's right, and the answers we give ourselves are not genuine actions, then they can't be intentional actions, and we would have to conclude, with Frankish, that reasoning is only partially intentional.

But this worry arises only if we ignore the *functional role* that the intention plays in determining the nature of a basic action. Shepherd (2015) says, "Typically, the intentions that are relevant to intentional actions *initiate, sustain, and guide* action" (p. 337, emphasis added). I will ignore for now the roles of initiation and sustenance; the idea of *guidance* is what we need to neutralise the worry that some basic actions might be both 'in the service' of an intentional non-basic action and yet not intentional. My argument is, in summary: The *answers* (to the questions we ask ourselves during an episode of reasoning) are intentional for the same reason that the *questions* are intentional – both are guided (that is, the *content* of both question and answer is guided) by the intention which motivates the episode in the first place. To make my argument, I will use the example (from earlier) of my intention to try to decide whether to take an umbrella with me.

I start with what I hope is an uncontroversial assumption: that no single intention a subject might form (under some description or other) exists in isolation of, unconnected to or uninformed by, at least some of that subject's other mental states (beliefs, desires, attitudes, etc.), standing and/or occurrent.[41] For example, I could not form the intention of trying to decide whether to take my umbrella with me when I leave the house unless (a) I

[41] I take this assumption to be in line with Bratman's (1987) influential proposal that the three constitutive norms on intention are requirements of *internal consistency*, *means-end coherence*, and *consistency with the agent's beliefs*.

knew what an umbrella was, (b) I possessed one, (c) I had an aversion to getting wet, etc. etc. These facts about me not only inform the formation of my *intention*, but they also guide the production (and the content) of the question I ask myself ("Shall I take my umbrella with me?") as motivated by that intention. To put it counterfactually: if those facts about me did not guide the formation of my intention, then the content of the question I ask myself in response would have been different. To spell it out: If I didn't own an umbrella, I would not have been able *either* to form the intention (to decide whether to take my umbrella with me) *or* ask the question ("Shall I take my umbrella with me?"). The same set of facts guide the formation of both intention and question. Similarly, I submit, the very facts about me which guide the formation of the intention and the content of my question also guide the content of my answer to the question. (Or rather, more accurately, the same facts *plus* some additional facts, such as (a) that I know that one can sometimes judge the short-term likelihood of rain by looking at the sky, (b) that I can look at the sky by looking out of the window, (c) that there is a window nearby I can look out of, etc.) To put it counterfactually again: if those facts about me (which guided both the formation of the intention and the content of the question) did not also guide the content of my answer to the question, the answer would have been different. To spell it out: if I didn't know that I can sometimes judge the short-term likelihood of rain by looking at the sky I would not have answered my own question with "Let's have a look out of the window". In summary, both the question and the answer to the question are guided by a similar set of facts about me, a set of facts which also inform the intention to which both question and answer were a response. Since the production of the answer was guided by that intention, it qualifies as intentional (as well as basic).

Let's test this proposal with an extreme case. Suppose I was linked up to a sophisticated system – the cognitive equivalent of the robotic suit we contemplated earlier in this section – which was designed to answer *on my behalf* all the questions I asked myself during an episode of reasoning. This system somehow supplies me with answers to my questions and I merely utter them (to myself). Should we say that my utterances are my actions or not? I think that depends on what information the system draws on to produce the answers. In my view, if the answers provided to me by the machine were informed *only* by all the mental states of mine, standing and occurrent, that *would have* informed the answer I

would have made had I answered my own question then, yes, the actions of uttering the answers provided by the system qualify as mine, even though I merely articulated (silently) the answers supplied to me by the device. (Similarly, if the hypothetical robotic suit that controls my limbs moved them in just the way I would have moved them given all my other mental states – my desire to be fast, but not so fast that I break into a sweat; my concern not to ruin these shoes by stepping in that puddle; my wish to appear cool to that group of people on the other side of the road; etc. etc. – and these mental states *alone* guided how the system controlled my limbs – then the actions, I claim, qualify as mine, and as intentional.)

The argument that the intention guides the answers, as well as the questions, involved in reasoning doesn't end there. Consider the *manner* in which the action of saying to myself "Let's look out of the window" might be further informed by the intention to decide whether to take an umbrella. Suppose I am fed up with the endless days of rain and I am dreading the prospect of yet another one. In this case, my intention (to bring it about that I decide whether to take an umbrella with me) might be partly informed by my disappointment concerning the recent spell of bad weather. As a result, I might make the inner speech utterance to myself ("Let's look out of the window") in tones of gloomy resignation. Suppose on the other hand I am excited by the prospect of a long drought finally coming to an end. In this case my intention (to bring it about that I decide whether to take an umbrella with me) might be partly informed by my optimism that the weather might be about to change for the better. As a result, I might make the inner speech utterance ("Let's look out of the window") in tones of eager anticipation. The point is that a psychological factor involved in determining the nature of the intention – namely, the attitude of the subject who forms the intention – is also in play in influencing (i.e., guiding) the nature of the utterance i.e., the way the question is answered. That shows that factors involved in the formation of the intention have a reach that extends beyond the content of the question, to the answer as well. This is another reason for thinking that that the action of uttering the answer is not only *in the service* of the intention but is *guided* by it too, and therefore qualifies as intentional. If that's right, then both kinds of inner speech utterance involved in reasoning – the answers as well as the questions – are intentional. They are both basic (because they both serve the non-basic intention of trying to bring it about that

the subject reaches a decision/conclusion) and they are both intentional (because they are both in the service of *and guided by* the intention which they serve). And if that's right, then reasoning is wholly intentional.

Section 4. Another argument from action theory

In the last section, I made extensive use of a distinction from action theory – the distinction (and relation) between a basic action and a non-basic one. I suggested that this distinction can and should be applied to the mental acts involved in reasoning, just as much as to the bodily acts involved in physical activities. In the previous chapter (Section 4) I mentioned that Antonia Peacocke (2023a; 2023b) does something very similar and in this section I want to consider her arguments. The first reason to do so is merely to buttress the credibility of my own use of the distinction; I take it that my application of the basic/non-basic distinction to the case of mental action is more plausible if someone else is doing the same thing (and doing so convincingly, which I think Peacocke is). The second reason for wanting to review Peacocke's work is to argue that it supports the conclusion I have been arguing for in this chapter – that reasoning is wholly intentional.[42]

In her (2023b), "How to Judge Intentionally", Peacocke's target is the standard view that it is "conceptually impossible" to intentionally *judge that p*. She provides a succinct summary of this view, abstracted from the details of the different versions of it:[43]

> Consider first what it is to judge or to believe something as a matter of intentional action: that would involve acting on an intention to judge something. Having some such intention involves exercising an adequate concept of judgment, which itself would require representation of judgment's necessary connection with truth. But to think of what you would be doing in this way—as constrained by truth in its correctness conditions and in its generation conditions—makes it incoherent for you simply to pick a content *p*, and intend to judge that *p*. What you come to judge intentionally and so self-consciously *as* a judgment must be directed by what

[42] In fairness to Peacocke I should emphasise that her papers entirely ignore any role that the act of speaking might play in the mental acts involved in reasoning. In fact, she warns that there are "good … reasons not to think of thoughts as 'internal' utterances of sentences" (2023a, p. 54). I return to this difference between us at the end of this section.

[43] See Peacocke (2023b, p. 4, FN 3) for a long list of writers who have made versions of this argument.

> you take to be the truth. What must determine the content of a genuine judgment, then, cannot be your acting on some specific intention to judge that *p*, but rather the truth of the matter, conceived of by you as controlling that content. To put the main point very roughly, this argument claims that (what you take to be) the truth, rather than you, must control any judgment of yours (pp. 3-4).

Another way of putting this might be to say that the *event* of you 'taking something to be true' is just that – an event, something that happens to you, and not an action you take. And if it's not an action you take, it can't be an action you take intentionally. So much for the standard view. Peacocke then proceeds to argue that, even if we accept the standard view (which she does, if only for the sake of argument), it is nevertheless true that there are some *other* kinds of judging which we *can* do intentionally. Specifically, she argues, we can intentionally *judge whether p*, or *judge which thing is F*, or *judge wh-* for all sorts of *wh-* questions.

To make her case, Peacocke invokes two principles from what she calls 'classic action theory' and applies them to the mental act of making a judgement. The first principle is that, "Anything that you do intentionally might also constitute an intentional action of another kind, but not every token of a new event-type that your action constitutes will also count as an intentional action of yours" (2023b p. 6). The classic illustration of this principle in the literature, which she cites, is from Davidson: "On returning home in the evening, you might intentionally *flip the light switch*. Given the electrical wiring of your house, your action of intentionally *flipping the light switch* can also constitute an instance of another kind of action: an instance of intentionally *turning on the lights*. If, unbeknownst to you, there is also a prowler in your house, this intentional action might also constitute an event of *alerting the prowler to your presence*. But that event itself is not here an *intentional* alerting of the prowler to your presence" (*ibid*. Original emphasis). She then claims that this principle can be applied to intentionally making a judgement: "Similarly, an event that executes your intention to *judge whether p*, or your intention to *judge which thing is F*, might indeed constitute a *judgment that p* without thereby constituting an *intentional judgment that p*." (*ibid*).

The second principle from classic action theory is this:

> When a given event constitutes your performing an intentional action, that event will have an indefinite number of determinate properties that are *not* specified by the intention you thereby execute. For instance, any flip of the light switch will be a flip with some determinate associated force—say, 1.23 Newtons. That light switch flip can be intentional as a light flip switch, and the event that executes it can have this determinate force of 1.23N, while it is *not* the case that you intentionally *flipped the light switch with a force of 1.23N*. You might leave the determination of specific applied force to the skillful motor routines your body runs when you execute that intention (*ibid*.)

Once again, she claims this principle can be applied to intentionally making a judgement.

> Along the same lines, a judgment might be intentional as a *judgment of which thing is F*, and the event that executes that intention might have a determinate content that *p*, even though it is *not* the case that you thereby intentionally *judge that p*, with that determinacy (pp. 6-7).

To illustrate these two principles at work in a case of intentional judgement she supposes that someone might have as their intention *judging the given name of Nietzsche*. (I am simplifying considerably here. Peacocke's argumentation is more sophisticated than I can do justice to here but capturing the gist of it is all I need to do to make my point.) Suppose, having formed this intention, the name you call to mind is "Friedrich". Peacocke's claim is that you have successfully executed your intention to judge Nietzsche's given name, without doing the conceptually impossible thing of *intentionally thinking that Nietzsche's given name was Friedrich.* You have successfully done one (non-basic) thing – judging the given name of Nietzsche – by means of doing another (basic) thing – *bringing to mind the given name you associate with Nietzsche,* namely, "Friedrich"). "[I]t is because you take *the given name you associate with "Nietzsche"*—that name you intend to call to mind—just to be Nietzsche's real given name that the one action can constitute the other action as well" (p. 13). (Incidentally, note how similar this illustration is to the case from Mele (2009, pp. 18-37), explored in detail in Chapter 4, Section 3, concerning Gail, and her *trying to recall 7 animals beginning with g*. Also note that, whereas on Mele's account, the event of *bringing [something] to mind* (e.g. *Nietzsche's given name* or *7 animals beginning with g*) is not an

intentional action, on Peacocke's account, and on mine, it is, and for the same reason: the event of the name being *coming to mind* is a basic action by which an intentional non-basic one is executed.)

In summary, according to Peacocke the action of *bringing to mind the given name you associate with Nietzsche* qualifies as intentional because it is the basic action by which you execute the non-basic action of *intentionally judging the given name of Nietzsche*. Frankish would presumably classify the event of *bringing to mind the given name you associate with Nietzsche* (i.e., recalling his given name) as a Type 1 event, and he would therefore classify the entire action of *intentionally judging the given name of Nietzsche* as only partially intentional. But by applying the basic/non-basic distinction from action theory to intentional mental actions, Peacocke shows this to be a mistake.

Peacocke's analysis, I claim, could be applied equally well to the dialogical (question-and-answer) nature of reasoning as I characterised it in the preceding sections of this chapter. To recap my position: the answers that we give to our own questions during an episode of reasoning are the basic actions by which we execute whatever intention was motivating the episode – trying to reach a conclusion, trying to come to a decision, trying to solve a problem, etc. Peacocke herself argues for something like this in her (2023a). In that paper, as in her (2023b), she appeals to the basic/non-basic distinction from action theory and applies it to cases of mental action. Specifically, she uses it to argue for what she calls 'content plurality' – the idea that a mental action with one content can be constituted by a mental action with a distinct content. One of the merits of her analysis, she argues, is that it addresses the long-standing philosophical problem of inference (which I take to be a paradigm case of reasoning). Peacocke diagnoses the longstanding problem of inference (which I will not rehearse here) as stemming from the assumption that, as she puts it, "any inference must be executed in transition from one judgment to another" (p. 53). But on her account, one mental judgement – for example, that p – can also constitute a different mental judgement – for example, that q; so, on her account of 'content plurality' no transition from p to q is involved. An example will illustrate how this works. Suppose you believe that *p iff q*; from this it follows that you can see that *figuring out whether p* is a way

of *figuring out whether q*. Peacocke spells out how a single mental action can have both these contents:

> (i) you think that *figuring out whether p* is a way to *figure out whether q* in your circumstances because *p* has the same truth value as *q*;
> (ii) you act on an intention to *figure out whether p in order to figure out whether q*;
> (iii) all it takes to figure out whether *q* in your circumstances is to think of the truth value of the content of a token judgment whether *p* as the truth value of *q*; and
> (iv) you execute both intentions (to *figure out whether p* and to *figure out whether q*) just by intentionally *figuring out whether p* in such a way that
> (v) the content of your *figuring out whether p* is qualitatively distinct from the content of your *figuring out whether q* (*ibid*.)

This way of understanding inference is precisely analogous, I claim, to what happens when you answer a question you have asked yourself as part of an episode of reasoning. Let's return to the simple example of an episode of reasoning from the introductory section of this chapter:

Q: Should I take an umbrella with me today?
A: Let's have a look out of the window.
Q: Does it look like rain today?
A: No, I will leave the umbrella behind.

This example can easily be adapted to fit with Peacocke's principle:

i. I think that *figuring out whether it will rain today* is a way to *figure out whether to take my umbrella today.*
ii. I act on an intention to *figure out whether it will rain in order to figure out whether to take my umbrella*.
iii. all it takes to *figure out whether to take my umbrella* in my circumstances is to think of the truth value of the content of a token judgment *whether it will rain* as the truth value of *whether to take my umbrella*; and

iv. I execute both my intentions (to *figure out whether to take my umbrella* and to *figure out whether it will rain*) just by *figuring out whether it will rain today*.

v. The content of my *figuring out whether it will rain* is qualitatively distinct from the content of my *figuring out whether to take my umbrella.*

In summary, Peacocke's use of the basic/non-basic distinction from classic action theory, and her application of it to an argument for 'content plurality', can also be used to argue for the central claim of this chapter, which is that reasoning is wholly intentional.

I want to end by returning briefly to something mentioned earlier in a footnote – Peacocke's claim that there are "good … reasons not to think of thoughts as 'internal' utterances of sentences" (2023a, p. 54). She does not spend much time on expanding on this; what she says is:

> One of the most important reasons not to do this has to do with the attitudinal aspects of thought. For example: if a judgment could only consist in an 'internal' utterance of a sentence, it is not clear why judging that *p* would ever constitute the doxastic commitment that judging that *p* really does constitute. Utterances can always be sincere or insincere, but there is no such thing as a sincere or insincere judgment. Recognizing the normative commitments involved in occurrent thoughts already recommends against any model on which thoughts are 'internal' utterances at all (2023a, pp. 54-55).

There is not much to go on here, so I won't do more than respond briefly to it. It seems to me that Peacocke's way of seeing things is what you would expect if you assumed that the only function of uttering a sentence was to communicate something to someone (else), either sincerely or insincerely. But if internal utterances are speech acts with illocutionary force, and if they have many more functions than merely communicating something (as I will argue in the next chapter) then that conception of an inner speech utterance is wrong. On my account, one of the functions of an internal utterance can be to *make* the doxastic commitment that Peacocke associates with a judgment. If, as I argue in the next chapter, thoughts are a kind of *inner speech act,* the "attitudinal aspects of thought" are not understood in the traditional way – as a subject's relation to a proposition – but rather as

the illocutionary force which characterises a speech act (or an inner speech act). Inner speech acts on this view are, like overt speech acts, in a broad sense *performative*, and some of them *just are* acts of entraining one's commitment to something (see Geurts, 2018). This idea will be more fully developed in the next chapter.

Section 5. The exploratory nature of reasoning

In the previous sections I have argued that, despite claims to the contrary, the inner speech utterances which comprise an episode of reasoning, whether those utterances take the form of questions or answers, have the same status as actions; I have argued that both are basic and both are intentional. But this might leave some readers feeling unsatisfied. If the questions and the answers are *the same,* why do they *feel* so different? When Strawson describes thinking as a kind of "rational reflex", and when Frankish describes the answers we give to our own questions, not as *actions* but as "automatic" and "autonomous" *events*, we intuitively know exactly what they are getting at. There *is* a sense in which thinking, even the purposeful kind of thinking involved in reasoning, seems to *just happen*. So how do we explain the tension between the account I am proposing, which argues that reasoning is a wholly intentional activity, with the phenomenology associated with reasoning, which suggests that it isn't? In this section I offer a tentative explanation based on the phenomenology of *exploratory* actions.

To begin with, I want to focus on the nature of exploratory *bodily* actions. I suggest that what makes them different from ordinary (non-exploratory) bodily actions is that they involve two particular types of mental state. First, they involve what Friedman calls an *interrogatory attitude* (2019, p. 5). Second, they require a special kind of *attention* from the subject performing them. Before I address each of these features in turn, let me pump your intuitions with some contrast cases:

a) Compare the action of reaching for a light switch you can see, with reaching for a light switch you can't see (because the room is completely dark).
b) Compare the action of sniffing a rose you know to have a beautiful scent, with the action of sniffing something you suspect of having a foul smell.

c) Compare the action of squeezing some toothpaste onto your toothbrush normally, with the action of squeezing the tube merely to make the paste appear at the end of the tube, and no further.

The actions in each pair are *similar* to each other in terms of the bodily movements involved required to realise them, but I hope that it will be intuitively obvious that they also differ in some significant ways, both functionally and phenomenally. Roughly speaking, the first action of each pair is done confidently, the second more tentatively; the first is routine, the second has a novel aspect to it; the first is done casually, the second more carefully. I think those differences can be distilled into two key properties: an interrogative attitude and a special kind of attention.

The *Interrogative Attitude.* Friedman (2019) makes a convincing case for thinking that inquiry is an activity, and that action alone is not sufficient for inquiry: "At the centre of any genuine inquiry is a certain kind of mental state or attitude" (p. 2). She calls this mental state or attitude the *Interrogative Attitude*, which she characterises as "a questioning or "asking" attitude… Inquirers have questions open in thought" (p. 5). Friedman goes so far as to consider it a matter of *necessity*:

> My thought about the *necessity* of the IAs [Interrogative Attitude] for inquiry is that it captures the sense in which a genuine inquirer *has to* be in a particular kind of goal-directed state. Perhaps this is another way to think about the force of the claim: every inquirer is trying to figure something out. If some subject is not genuinely trying to figure out *Q*, then she's not genuinely inquiring into *Q*. The IAs can be thought of as ways of trying to figure something out, they are manifestations of that sort of effortful state (p. 6, emphasis added).

I think that's right. Consider the first action in each pair of the contrast cases above - there is nothing questioning or inquiring about them. By contrast, the second action of each pair involves an implicit question of some sort: Is the light switch somewhere here? Has this food gone off? How much pressure is needed to make the toothpaste merely appear at the end of the tube? In short, the second action in each pair involves the interrogative attitude.

Attention. Exploratory actions involve a particular kind of *attention*, one we might characterise as *anticipatory* or *expectant*. Reaching for a light switch I can see is such a routine action I can perform it with the most cursory kind of attention. (As is well known, we can perform quite complex sequences of basic actions in the service of a non-basic action without even being aware of the basic actions we are performing. Driving a car along a well-known route is a notorious example.) But when I am groping for a light switch in the dark, I am concentrating on the anticipated results of my groping. I hope and expect to make contact with the switch or at least the wall, and that expectation, I suggest, is a special form of attention. Something similar can be said for the other contrast cases too. The reason that exploratory actions demand this kind of attention from the subject performing them is obvious: since the function of the action is to find something out, the action would not perform its function unless the subject attended to the response it was designed to elicit.

Note that the exploratory actions we have been looking at – the second action in each pair of the contrast cases above – are no less basic, nor less intentional, than the non-exploratory ones – the first action in each pair. Reaching for a light switch I can see is a basic action – there is no procedural knowledge needed for me to do it, I just do it directly because I know how. And it is intentional because it serves the intentional non-basic action of turning on the light. The same is true for reaching for a light switch I *can't* see. Groping for a light switch in the dark is a basic action – I do it directly – and it serves the non-basic action of finding the light switch.

To summarise, I have argued that some bodily actions can be characterised as *exploratory*, and that exploratory actions possess two distinguishing features: they involve an *interrogatory attitude,* and they require a special kind of *attention* from the subject performing them. I will now argue that at least *some* of the inner speech utterances we make when engaged in reasoning possess the same two properties. (I will return to this qualification at the end of the section.)

Reasoning involves an interrogatory attitude. As previously stated, 'reasoning' as I am using the term covers a range of activities – deliberating, reflecting, calculating – which are all

motivated and guided by an intention to work out what to do or what to think. The "what" in that sentence is the clue: it implies a question, and asking a question is a paradigm case of interrogation. So, clearly reasoning – as I've defined it – involves an interrogatory attitude.

Inner speech utterances (involved in reasoning) require a special kind of attention. It is not possible to be engaged in reasoning about one thing – trying to bring it about that I come to a decision about something, for example – while also attending to something else. The reason is obvious: for the series of utterances to succeed in performing their function of helping to bring it about that I reach a conclusion, I must attend to (concentrate on) what they mean. Only if I attend to what they mean can I expect to be in a position to respond to them. Attending to them in this expectant way is a special kind of attention.

The proposal, then, is that these two special kinds of mental state – the interrogative attitude and attention (based on expectation) – account for the special phenomenology associated with exploratory actions generally, and the exploratory actions involved in reasoning (the questions and answers performed in inner speech) in particular. The further claim is that it is this *phenomenology* which encourages the impression that the questions have a different agential status to the answers, and that this impression is mistaken. How so? The attitudes of expectation and anticipation gives the subject the impression that she is *waiting for something to happen*. After all, what normally happens when you ask a question is that you wait for your interlocutor to answer it. But in the case of reasoning, what happens next is that the subject herself answers the question. The claim, then, is that the attitudes of expectation and anticipation involved in reasoning gives the subject the *impression* that what follows her asking the question is a mere *happening*, as it would be if the subject were waiting for another person to answer her question, and that this impression misleads her into thinking of the next event – her answering her own question – as a mere happening and not an action. But it isn't, it's another action.

An important qualification. I mentioned earlier that only *some* of the inner speech utterances which constitute an episode of reasoning are exploratory. That qualification is necessary because if the episode is successful, and results in the subject coming to a decision (or reaching a conclusion, or solving the problem, etc.) then her *final* inner speech

utterance (e.g., “No, I will leave the umbrella behind”) is clearly *not* exploratory in nature. On reaching a conclusion the interrogatory attitude is replaced by a different attitude – one associated with the matter having been somehow settled. This chapter has been concerned with the actions which, during an episode of reasoning, lead up to the point at which the reasoning results in a conclusion of some kind – a decision for example. The nature of that conclusion (or, better, the nature of the phenomenon of concluding an episode of reasoning) is addressed in the next chapter.

Section 6. Summary

Reasoning is an intentional activity which is at least partly constituted by the action of making inner speech utterances. These utterances are the basic actions by which the non-basic action of reasoning-with-a-view-to-deciding-something is executed and as such they qualify as intentional. The inner speech utterances are intentional whether they take the form of questions or answers. *Pace* Frankish (2018) reasoning is wholly intentional. During an episode of reasoning the inner speech utterances which result in the subject coming to a decision have the same characteristics as other exploratory actions – they involve an interrogative attitude and a kind of expectant attention. These features give reasoning its characteristic phenomenology, and this is responsible for the misleading impression that some of the inner speech utterances involved in reasoning are actions (albeit basic ones), and some are not when, in fact, they all are.

Chapter 6. A New Argument for *Thinking-as-Speaking*[44/45]

Abstract

Sometimes, thinking a thought and saying something to oneself are the same event. Call this the "thinking-as-speaking" thesis. It stands in opposition to the idea that we think something first, and then say it. One way to argue for the thesis is to show that the content of a thought (its full and final *meaning*) cannot be fully represented by a mental state before the production of the utterance which expresses it. I make an argument for that claim based on speech act theory. Speech act theory has it that (overt) speech acts involve both propositional content and illocutionary force. My claim is that the content and the force only come together with the production of a linguistic expression, and that the action of producing a linguistic expression is intrinsic to the action of combining content with force, which is intrinsic to the performance of the speech act. So much for overt speech acts. Many inner speech utterances are inner speech *acts*, containing both content and force. As with overt speech acts, what an inner speech act *means* is more than its propositional content, it also includes the force with which it was performed, and those two aspects of an inner speech act only come together with the production of an inner speech utterance. If the determinate silent thought (in the sense of 'what the speaker fully meant by what she silently uttered') includes the illocutionary force with which the utterance was performed, then it follows that the determinate thought cannot have been represented 'in the head' of the speaker before she performed the inner speech act. This raises a metaphysical question: If a silent thought is not an *object* in the head of the thinker (such as a Fregean proposition, or a mental representation of propositional content, or a sentence in a language of thought) what is it? I suggest it is not an object at all, but a kind of action; the action of performing a speech act.

[44] Some of the material in this chapter appears in Frankfort (forthcoming).

[45] This expression should not be confused with the similar sounding "Thinking for Speaking", introduced by Slobin (1987). An appendix to this chapter (Section 11) is dedicated to explaining what this term means, and to summarising Vicente's (2022) argument for why it might be misconceived.

Section 1. Introduction

Ever since Plato first suggested that thinking should be understood as silently talking to oneself,[46] some philosophers have argued that, in some sense, thinking *just is* inner speaking, and that inner speaking *just is* thinking. Vygotsky (1987) seems to be saying something like this when he says: "Thought is not merely expressed in words; it comes into existence through them" (p. 218). Likewise Frankish (2018) when he says: "Rather than just enabling thoughts to become conscious, [inner speech] must enable a distinctive kind of thought, which is conscious. It must *make conscious thought*, rather than just *making thought conscious*" (p. 228). Since I'll be referring to this idea a lot, I need to give it a name: call it the *thinking-as-speaking* thesis. The *thinking-as-speaking* thesis is not that *all* instances of inner speaking are instances of thinking, or that *all* instances of thinking are instances of inner speaking. It is generally agreed that inner speaking is involved in a wide range of cognitive functions; the *thinking-as-speaking* thesis is only that one of those functions is thinking. It is also generally agreed that thinking can involve multiple forms, not just a linguistic one; the *thinking-as-speaking* thesis is only that some instances of thinking are also instances of inner speaking. Explaining exactly what the thesis amounts to has proved challenging, not least because, while the idea has a lot of intuitive appeal, another intuition pulls strongly in the opposite direction; we feel that thoughts are things we 'have first' (whatever that might mean), and only afterwards do we put them into words, typically in order to communicate them to someone else.[47] We wonder, reasonably enough: If my thought didn't exist first, how would I know what to say? And this second intuition is buttressed by a great deal of philosophical tradition.

Several theorists have recently argued for different versions of what I'm calling the *thinking-as-speaking* thesis. Roessler, 2015; Gauker, 2018; Vicente and Jorba, 2019; Kompa,(2023);

[46] "[Thinking is] the talk which the soul has with itself about any subjects which it considers…. [T]he soul…when it thinks, is merely conversing with itself, asking itself questions and answering, affirming and denying…" (Plato, 1921)

[47] According to Roessler (2015), Bernard Williams (2002) associates this view specifically with Rousseau, whom he quotes as saying, "we first and immediately have a transparent self-understanding, and then go on either to give other people a sincere revelation of our belief (...) or else dissimulate in a way that will mislead them" (Williams, 2002, p. 193).

Vicente, (2022).[48] The idea common to all of them, as I see it, is that, whatever content is represented 'in the head' before the onset of the language production process (see Chapter 2 for details of this process), that content goes through a process of *transformation*, such that the content of the final utterance is different from whatever was originally represented 'in the head'. We can formalise the idea like this:

> *Thinking-as-Speaking* thesis: Inner speaking makes a constitutive contribution to the production of a token thought, and to determining its content.

Other theorists who are sympathetic to the thinking-as-speaking thesis have gone some way in this direction but seem to hold back from going all the way. Vicente (2022), for example, writes this:

> While it is usually assumed [by those who subscribe to the LoT hypothesis] that the content of what we tell ourselves is exactly the content of a non-linguistic thought, I argue that there can be a lot of transformation in the process of converting a thought into words. Thus, the content of what we tell ourselves, being intrinsically linguistic, is different from the content of the thought our speech transmits (p. 1).

I agree with the spirit of this claim, but I think that this way of putting things is problematic. Note that in this passage he claims that "a thought" contains two different contents at different stages of its development: the content of a linguistic utterance ("what we tell ourselves") and the content of the original ("non-linguistic") thought. But this is incompatible with the principle – widely accepted – that a token thought is individuated by its content. *The* individuated token thought being expressed by the thinker can't be identical with either of the two objects referred to in this passage as "the thought". It can't be identical with the posited non-linguistic object because that object is different from – has different content from – the linguistic object (sentence, expression) produced. And *the* thought can't be identical with the linguistic object either, because everyone agrees that the *semantic* meaning of an utterance, in the absence of context and illocutionary force (more on this below), underdetermines the full meaning of that utterance. I might silently utter

[48] At least this is how I read them. They might disagree.

the very same sentence on two different occasions, and the thought which it expresses might be quite different, depending on the context in which, and the force with which, I utter it. To be clear, I agree with Vicente's claim that whatever is 'in the head' of the inner speaker at the onset of her speaking is transformed during the speech production process, but I reject the idea that this should be identified as the speaker's determinate *thought*. For want of a better term, let us call whatever is in the speaker's head, and which informs the production of what becomes a token thought: *information.*[49] Note that Levelt *et al* (1999), whose model of the different stages of the language production process we considered in Chapter 2, also uses the term 'information' when he describes the process of going from 'conceptualisation' to 'lemma retrieval':

> A major issue, therefore, [with converting the concepts activated by the subject's communicative intention into items for which she has words] is how the speaker gets from the [pre-linguistic] notion/*information* [my emphasis here] to be expressed to a message that consists of lexical concepts (here *message* is the technical term for the conceptual structure that is ultimately going to be formulated). This is called the *verbalization problem,* and there is no simple one-to-one mapping of notions-to-be-expressed onto messages…. (Levelt *et al*, 1999, p. 3, original emphasis).

All of this supports my suggestion that whatever "notion/information" exists in the head of a subject who has a communicative intention 'in mind', it is only transformed into a fully determinate thought when the language production process is engaged.

I make an argument for this suggestion by appealing to speech act theory. The argument, roughly, goes like this:

1) Speech act theory has it that speech acts involve both propositional content and illocutionary force.
2) The content and the force only come together with the production of a natural language expression.

[49] My thinks to Alex Grzankowski for this suggestion.

3) The action of combining content with force is intrinsic to the action of producing a natural language expression, which is intrinsic to the performance of the speech act.
4) When we attribute a thought (in the ordinary sense) to someone who performs an overt speech act we are not referring to *propositional* content alone, but to the full meaning of the speech act, including the illocutionary force with which it was uttered.
5) Therefore, what we are attributing to the speaker (as their thought) cannot have been represented 'in their head' before they performed the speech act.

As with overt speech acts, so with inner speech acts. When we think silently, we do it (at least sometimes) by making inner speech utterances. These inner speech utterances are inner speech *acts*, containing both content and force. The thought being expressed by the performance of the speech act is not fully represented in the head of the thinker before the inner speech act is performed. A token thought, I suggest, is (at least sometimes) the token action of performing a speech act.

Allow me a short digression. Famously, Hanks (2007, 2015, 2018) argues that the content-force distinction is ultimately unsustainable:

> I think that the content-force distinction should be abandoned altogether. I am skeptical of the idea that there are propositional contents that represent states of affairs independently of what speakers do in making assertions or forming judgments. An account has to be given of how these contents represent states of affairs that does not make any appeal to the intentional actions of speakers. I doubt that there is any plausible way to do this. Representations are things that we produce in speaking and thinking about the world (2007, p. 143).

Of course, Hanks' position is contentious, and I don't propose to defend it here. It is interesting to note, however, how his position complements my own.[50] Hanks has, over many years, tried to make sense of the nature of propositions, and concluded that it can't be done "independently of what speakers do in making assertions or forming judgments". It

[50] I don't mean to imply that Hanks would agree with the arguments of this thesis.

follows logically from that claim that the propositional contents of an utterance which represent states of affairs are *dependent* on what a speaker does in making assertions or forming judgments. According to Hanks, orthodox thinking about propositions has got things the wrong way round:

> Instead of speech acts and propositional attitudes inheriting their representational contents from propositions we can say that propositions inherit their representational features from the actions subjects perform in performing speech acts or forming attitudes…. This reverses the normal order of explanation and locates the source of representational content in an intentional act of the subject. In doing so it shifts the problem [of the unity of the proposition] away from understanding how propositions represent states of affairs to the problem of understanding how speakers represent states of affairs in the production of speech acts and propositional attitudes. (2007, pp. 159-160).

If Hanks is right, it seems reasonable to infer that the action of performing the speech act makes a constitutive contribution to its content (broadly construed) and that, of course, is exactly what the *speaking-for-thinking* thesis says. And if that's right then it seems reasonable to infer that the action of performing the speech act makes a constitutive contribution to the content of the *thought* (in the ordinary sense of that word) being expressed by the utterance. And from that it follows logically that the thought did not exist in its final and determinate form before the speech act was performed.

Needless to say, this is all very congenial to my way of thinking. So, one way of justifying my claim that a thought (in the ordinary sense) does not exist as a mental state before it gets expressed would simply be to appeal to Hanks. But Hanks's view is a long way from being orthodoxy; most theorists continue to endorse (at least implicitly) some form or another of the content-force distinction. So I will make the argument sketched above, and argue that, even if, *pace* Hanks, the standard picture from classic speech act theory is right, it can still be shown that a thought (in the ordinary sense) does not exist as a mental state in the head of the thinker before it gets expressed in natural language.

End of digression.

I am not the first to claim that thoughts don't exist 'in the head' before they are uttered in inner speech. Geurts (2018) and Deamer (2021), also make this claim in the course of exploring the role inner speech plays in performing inner speech *acts*.[51] But, while this is a *claim* they both make, it is not one that either of them offers an *argument* for, and that is the goal of this chapter. I should stress that my goal is not to challenge their claim, but to argue for it. The rest of the paper goes as follows. Section 2 provides a short account of the arguments of Geurts (2018) and Deamer (2021). I explain that, while their respective arguments lead them both to *claim* that the full and final content of an inner speech act is not represented beforehand by a mental state, neither provide an argument for this claim. In section 3 I make an argument for this claim based on speech act theory – the argument summarised above. Section 4 explains why, despite the complete absence of any reference to inner speech in the speech act literature, it is nevertheless true that inner speech *utterances* are frequently inner speech *acts*. Section 5 uses the action of deciding something as a case study. I show that the event of deciding something, either overtly or silently, is a speech act. Since the speech act *is* the making of the decision, and since the decision *is* the determinate content of the thought, the thought cannot have been represented in the head of the decider before the speech act was performed. Section 6 addresses an important epistemological question raised by previous sections: What kind of knowledge do we have of our own thoughts? Section 7 asks how much of what I argue for in this Chapter depends on accepting the motor-sensory view of inner speech outlined in Chapter 2. (I argue: "Not much".) Section 8 argues that my account of inner speech acts fits into a broader picture of the nature of our mental lives. I suggest that the arguments of this chapter are complementary to those of Tim Crane (2017), who argues that the content of a conscious mental state is something we actively bring into being, rather than something that is an object formed in the head and then made conscious. Section 9 clarifies some of the technical terms involved in the philosophy of bodily action (activity, actions, accomplishments and achievements) and shows how it can be applied to thinking as an activity. This section also returns to the subject not mentioned since Chapter 1: if a token

[51] Wilkinson (2020) also makes use of speech act theory when writing about inner speech. In his case he does it to argue for the agentive the role of inner speech in self-knowledge. I say more about the knowledge we have of our thoughts in Section 6 of this chapter, below.

thought can be understood as a token thinking, and specifically as the performance of a speech act, what are the implications for the nature of a token thought when the term is being used to refer to an object which represents content? Section 10 concludes.

Section 2. Inner speech utterances as inner speech acts

This section is about Geurts (2018) and Deamer (2021). Both authors make convincing use of speech act theory to explain the function of some instances of inner speech, thereby buttressing the idea, developed below in Section 4, that inner speech utterances are very often inner speech acts. Also, both authors deny, or at least doubt, that before a speaker performs an inner speech act, the thought she is expressing exists as a mental state. But they only do this in passing, and they don't offer an argument for that claim. What they have to say sets the scene for my own argument: that it is inconsistent with speech act theory to suppose that the thought being expressed by a speech act *could* exist as a mental state.

Geurts (2018) uses speech act theory to answer a puzzle about inner speech: if the main purpose of *overt* speech is to communicate something to someone else, why would we need to use *inner* speech to communicate with ourselves? Don't we know what we think already? His answer to the puzzle, roughly, is that inner speech utterances are speech acts, and the function of speech acts is to get something done; in the case of *inner* speech acts what is being done is 'entraining commitment'. When I say to myself "Time to leave the house now, or I'll be late" I am performing the speech act of committing myself to leaving the house now, so I won't be late. When I say to myself "It's sunny today" I am performing the speech act of committing myself to behave in the future in ways that are consistent with what I said. (To pack a sun hat, for example, or leave my umbrella behind.) Note something else Geurts says: "… I oppose the widely held view that, inevitably, a sincere speech act must be preceded by the mental state it expresses. It is perfectly coherent to hold that a self-addressed statement may be a way of *forming* a belief, that a self-addressed command may be a way of *forming* an intention, and so on" (p. 278, emphasis added). In other words, Geurts rejects the view that there is a mental state which represents a fully specified thought (in his words: a "belief" or an "intention") which is subsequently expressed in

words. But note, also, that he does not claim that what he says about the 'commitment entrainment' function of inner speech is an *argument* for rejecting the view that "a sincere speech act must be preceded by the mental state it expresses", only that it is *consistent* with it: "[I]n this paper I will neither defend nor presuppose that self-addressed speech acts may serve to form intentions and beliefs. However, my theory is consistent with that possibility, which I consider to be an important selling point" (ibid). I mention this because the purpose of *this* paper, unlike Geurts', is precisely to argue that no mental state *could* fully represent what an inner speech act then goes on to express.

Next, let's turn to Deamer (2021). She agrees with Geurts that at least some inner speech utterances have the function of committing the speaker in some way, but she denies that they all have that function. Geurts' idea, she says, seems to work well for promises and directives, but less well for questions and assertions. If I ask myself (silently) "When do I need to leave the house to get my train?" it is hard to see what I am committing myself too. More likely I am simply engaged in deliberating about when to leave home to catch my train. And if (once on the train) I say to myself "The state of this train is disgraceful!" it is very unclear what I am committing myself to, exactly. Deamer doubts the subject is committing herself to anything.

Deamer offers an alternative explanation for the function of at least some utterances of what she calls "self-talk":[52] We talk to ourselves because we are "to some extent, self-blind (p. 431)" when it comes to our "communicative intentions", and we need self-talk to find out what they are. (She reminds us of the well know quote from E. M. Forster: "How can I tell what I think, until I see what I say?") What she means by "communicative intention" is "(the thought/intention/belief) underlying the utterance" (p. 429). In keeping with the interpretivist position of Carruthers (2009, 2011) and Cassam (2011) – i.e., the idea that we have to interpret our own inner speech to find out what we think – she claims we don't

[52] Both Geurts and Deamer (whose paper is partly a response to Geurts') use the term 'self-talk' rather than the term inner speech. They do this to draw attention to *what* is being done, rather than *how* it is being done. Geurts (2018) writes: "Much of the previous research on self-talk is focused on inner speech. To my knowledge, there is no evidence that the dichotomy between inner and outer speech is of any great significance, and therefore I prefer the term "self-talk", which is neutral between "inner" and "outer"" (p. 273).

have access to our mental states *in general*. She builds on this idea to argue that our communicative intentions are among the mental states we don't have access to *in particular*:

> Just as we don't know our own mental states until they make themselves manifest to us via some medium on which we have a grip, so we don't know our communicative intentions, until we have spoken (p. 429)

(She also offers an alternative analysis of why we might sometimes have the *impression* we know what we are about to say before we say it (even though we don't really) which, with its emphasis on what the self-speaking is intended to *achieve*, is very much in keeping with speech act theory:

> [W]hat guides our self-speaking is typically not prior knowledge of what we're going to say, but rather a vaguer sense, prospectively, of what we're trying to *achieve* by speaking given the context, and, retrospectively, whether we have expressed ourselves aptly or not once we have spoken (p. 432, emphasis added).

I think that's right, but I also think there is more to say about the knowledge we have of our thoughts before we express them. I address this topic in Section 6 below.)

Deamer argues that inner speech sometimes performs a third function (i.e. as well as self-entrainment and discovering our communicative intentions), namely an "expressive" one, meaning that it sometimes *reveals* things to us about ourselves, rather than *communicates* them.

> Saying "Ouch!" *reveals* that I am in pain, it doesn't describe me as in pain (unlike saying "I'm in pain" does). Similarly, saying "To hell with you!" reveals that I am unhappy with what you've done, it doesn't describe my unhappiness (unlike saying "I'm unhappy with what you've done"). Similarly saying to yourself "Come on!", or "I'm such an idiot!", or whatever, reveals something about you. It seems that a great deal of self-talk is expressive in this sense [p. 431. Emphasis in the original].

In her final remarks Deamer makes a claim that echoes Geurts':

> Contrary to many views of communication, there is no thought "in mind", present and accessible, that is then articulated in language (p. 431).

In other words, there is no representation of what the thinker is about to say instantiated in the thinker's head, the content of which is fully determined, and which is then subsequently made available to the thinker by being translated into a natural language format. The speech act involved in communicating the thought (to oneself) makes it (determines) the thought that it is.

I agree with Deamer's claim. The only problem I have with it is that it doesn't follow from her argument (i.e., her argument that the function of inner speech is to reveal to us our communicative intentions, or to be expressive). It would be quite consistent to say that inner speech utterances are used to perform all three of the functions she and Geurts argue for and yet still maintain that these thoughts existed as mental states, with determinate content, before the utterance was made. Perhaps, when I said to myself "Time to leave the house, or I'll be late", I already instantiated a mental state which represented the content 'time to leave the house or I'll be late', and the inner speech utterance was my way of committing to it. Perhaps, when I asked myself, "When do I need to leave the house to get my train?", I already instantiated a mental state which represented the question 'When do I need to leave the house to get my train?' and the utterance was a way of revealing to myself that I need to catch a train. When I said to myself "I'm such an idiot!", perhaps I already instantiated a mental state which represented 'I'm such an idiot!', and the utterance was my way of expressing it. In short, inner speech utterances might have all the functions Deamer and Geurts say they have, but that does not show conclusively that "there is no thought "in mind", present and accessible, that is then articulated in language." For all either of them has said, there might be.

Section 3. An argument from speech act theory

In this section I make an argument for Guerts' and Deamer's claim that, as Deamer puts it, before we speak "there is no thought "in mind", present and accessible, that is then articulated in language." (In fact, my argument will be even stronger: that there is no thought 'in mind', present and *in*accessible, that is then articulated in language.) My argument is based on speech act theory itself; I will argue that it is inconsistent with speech act theory to suppose that a thought *could* be in mind, present and (in)accessible, before being articulated in language. To make that argument it will be helpful to have before us a view of what speech act theory says. According to Recanati (2013, p. 1), the 'standard picture', due to John Searle, has two parts:

1. Leaving aside the speech acts which are devoid of representational content ('expressives', in Searle's classification), the content of a speech act is always of the same semantic type: it is a proposition. Thus, the order 'Go to the store, John!', the question 'Will John go to the store?' and the assertion 'John will go to the store' differ by their *illocutionary force* (that of an order, a question, and an assertion, respectively), but they (allegedly) share the same content: they represent the same state of affairs (John going to the store), corresponding to the same proposition (that John will go to the store). So understood, speech act content is 'force-neutral'.
2. The so-called 'sentence moods' – that is, the devices, whatever they are, that differentiate clause-types into declarative, interrogative, imperative, exclamative, etc. – encode illocutionary force and do not affect representational content….

This 'standard picture' is, of course, contentious, and Recanati's paper is principally a summary of several of the ways it has been contended. But the key idea is that every speech act can be factored into these two components – propositional content, which is force neutral, and Illocutionary force, which indicates the speaker's primary purpose in performing the speech act – each of which can vary independently of the other. That feature of speech act theory is all I need to make my argument.

To make my argument I need to say something more about the expression 'a token thought'. For many philosophers, the term 'thought' refers to the proposition which is the force-neutral content of the speech act. In the example from Recanati's example in the above quote, this is the proposition *that John will go to the store*. My claim is that this use

of the term 'thought' is a special use, a term of art, and that, even if such an object were to be represented 'in the head' of a speaker before she spoke, it would not qualify as her full and final thought in the *ordinary* sense of that term. That's because when, in ordinary discourse, we ascribe a thought to someone who has performed a speech act, we include in what we ascribed to them the force with which they said it. A thought, in the ordinary sense of the term, i.e., when we use the term to refer to *what someone was thinking*, is not force neutral.

To illustrate what I mean, consider the three examples from the Recanati quote above:

a) 'Go to the store, John!'
b) 'Will John go to the store?'
c) 'John will go to the store'

These are all different speech acts, each with a different force: imperative, interrogatory, and declarative respectively. Because these speech acts have a different force, the speakers who perform them can be said to *mean* different things by them. When we ascribe a token thought to the speaker who has performed a speech act, we take into account both what they said (the words we heard) and what they intended by those words (the force with which they said it).[53] That is, when we ascribe a thought to someone, we don't ascribe only propositional content, and nor do we ascribe only linguistic or semantic meaning, i.e. the meaning of the natural language expression as understood by a competent speaker of the language. Rather, we include in our ascription the force (or forces[54]) with which they uttered the expression. Let me illustrate the point. Suppose you were asked, for each of the (silent) speech acts above, "What was she *thinking*, the person who thought that?" For a) you might reply something like, "She wanted John to go to the store, so she told him to go". For b) you might say something like, "She wondered whether John would go to the store, so she asked if he would go." For c) you might reply, "She judged that John would go

[53] When I refer here to what a speaker means, I am not referring to 'speaker meaning' in the Gricean sense of conversational implicature. I am merely referring to what a speaker says when speaking literally and sincerely.
[54] A speaker can do more than one thing with the same speech act. An example from Hanks (2018): "When the umpire says, 'You're out', he does two things at once: he asserts that you are out, and he makes a declaration to the effect that you are out" (p. 20).

to the store, so she predicted that he would go." These are the ways you might ascribe to the speaker the thought (in the ordinary sense) being expressed. You would not, I claim, say that they were *all thinking the same thing*, and that what they were all they were thinking was "That John goes to the store." (You might well say, however, that part of the pre-linguistic *information* each speaker had in mind was *John going to the store*, but that *information* does not qualify as the subject's determinate thought.) If that's right, then they do not all have the same thought 'in mind'; the token *thought* being expressed by the speech act is different in each case. That being so, the thought cannot have been represented by a mental state before the performance of the speech act.

Orthodox speech act theory would say these thinkers all 'grasp' or 'entertain' the same propositional content. But in my view these metaphors are in long-standing need of explanation and are part of the deep problem with the orthodox view. I take it that grasping and entertaining are things which are done consciously, but it is mysterious how a subject could be conscious of raw, pre-linguistic content, propositional or otherwise. On my account, the language production process, which synthesises pre-linguistic content (information) with illocutionary force, does two things: it both contributes to what the final utterance means, and it makes the subject conscious of that meaning.

None of this is to deny that there is plenty that *is* being represented in a speaker's minds before she performs a speech act, and that some of what is being represented is common to all the thinkers of a) through c) above. If asked *about* what these thinkers were thinking you might well say that they were all thinking *about* the same thing, namely, John going to the store. And no doubt the idea of John going to the store was represented by a mental state of some kind before the speaker spoke. But that content (whether we call it propositional content or information), since it lacks any force, is not the same as the thought the speaker expresses with the speech act they perform, since the performance includes the element of force. Also, there must be a mental state which is the intention to do *something* (whatever that is) with the idea of John going to the store, before the speaker performs the speech act: for example, an intention to issue an instruction, ask a question, or make an assertion, etc. But once again, this mental state, since it is content neutral (according to speech act theory), does not represent the determinate thought the speaker expresses with the

performance of the speech act. The determinate thought only comes into being when the speaker uses language to bring the two things together – the force-neutral content and the content-neutral force – and she only does that when she performs the speech act.

Let's consider a possible objection. Take the case of a very simple thought, the thought *that snow is white*. This is the kind of proposition philosophers often use as an example of a simple proposition which can surely be 'entertained' as a thought before it is expressed in words. If it can be entertained as a thought before it is put into words then it must be represented, somehow, in the head before it is put into words, and I am therefore wrong to claim that thoughts are not objects but speech acts. My response is to say that when philosophers invoke this example they are not, in fact, *thinking* about the colour of the stuff called snow, except in rather unusual sense; what they are really thinking about is what example they can use to illustrate the case of stating a simple proposition. They have this ready-made example to hand, and they take it off the shelf, so-to-speak, because they are thinking about thinking, and they need to insert an x in the sentence 'x is an example of a simple proposition'. If a person was really *thinking* about the colour of the stuff we call snow, what they would be representing in their minds before they perform the utterance is (a) the stuff we call snow, (b) its colour, and (c) white, but they would not have generated the (silent) *assertion* that the colour of snow is white until they formed and asserted the utterance "snow is white". Until then, they are processing (unconsciously) all the cognitive material (information) necessary for the action of producing the thought that snow is white, but they haven't actually, consciously, *thought* it.

Section 4. Some inner speech utterances are inner speech acts

I claimed in section one that at least some inner speech utterances function as inner speech *acts*, and, notwithstanding the arguments of Geurts and Deamer (Section 2), some readers will find this surprising, or even implausible. The purpose of this section is to justify the claim.

One good reason to find the claim surprising is that the idea that inner speech utterances are, at least sometimes, inner speech acts, appears to be entirely absent from the speech

act literature.[55] Inner speech is not mentioned by either Austin (1962) or Searle (1976) and, to take a more recent example, neither of the terms 'inner speech' nor 'silent speech' get a single mention in any of the 15 essays published in *New Work on Speech Acts*, (Fogal, D., Harris, D. W., Moss, M (eds.), 2018). But this absence of any reference to inner speech acts in the speech act literature is surprising, for several reasons. Firstly, as discussed in the previous section, standard speech act theory has it that what a speaker means by what they say is underdetermined by semantic content alone and depends also on illocutionary force. If that's true of overt speech, then we need a special reason for why it isn't also true of inner speech.[56] Inner speech utterances share many features with outer speech utterances. For example, as with some overt speech, some inner speech utterances *lack* illocutionary force; for example: silently reciting lines from a play or poem, reminding yourself what someone said by quoting their exact lines to yourself, reading a book and silently speaking the words. These utterances might be described as 'acts of speech' which are nevertheless not 'speech acts'.[57] Other interesting questions arise: Where inner speech *does* have illocutionary force, does the force function in just the same way as when that same utterance – an utterance with the same semantic content – is made overtly? How do Gricean theories of meaning accommodate the phenomenon of inner speech acts? (I take it that the key Gricean idea is that for a speaker to mean something they intend to produce certain effects on the addressee, and to have that intentioned recognized by the addressee. How does that work when the addressee is the same as the speaker?) Are there some kinds of illocutionary force which are *exclusive* to inner speech acts, given its private nature? If *inner* speech acts are not accommodated by the Searlean classifications of speech acts (see below), in what classification *are* they accommodated? These all seem to me to be interesting questions, but I am not aware of any attempts to address them.

Another reason I find this gap in the speech act literature surprising is that, by contrast, among many philosophers who write about inner speech, the concept of 'inner speech acts' is more or less taken for granted. This follows from the widespread agreement among

[55] At least, I haven't managed to find any references to it.
[56] For example, in a different context, Hornsby (2009) writes, "[T]he non-explicit performance of illocutionary acts is as ubiquitous as the use of language" (p. 905). So why wouldn't this apply to the *silent* use of language?
[57] This this way of putting things comes from Roessler (2015).

many philosophers of inner speech that "IS [inner speech] has almost as many functions or uses, as we can discover in OS [overt speech]" (Martinez-Manrique & Vicente, 2015, p. 7). Since one of those common functions or uses of overt speech is performative (in the broad sense) one would expect that inner speech would, at least sometimes, be performative too.

Geurts (2018), for example, thinks it follows straightforwardly from the Vygotskian idea (see Alderson-Day & Fernyhough, 2015) that inner speech is a form of internalised outer speech:

> Speech acts start their career as a form of social interaction, but almost as soon as they begin to talk, children will also talk to themselves, using speech acts to shape their own behaviour… social talk becomes private talk, which at first is mostly overt, but is increasingly internalised to become inner speech, or "verbal thought" (p. 272).

Wilkinson & Fernyhough (2018) think the "primal use" of inner speech is to make speech acts:

> [W]hat is [inner speech]? In line with a number of other theorists (Vygotsky 1987/1934, Fernyhough 1996, Martínez-Manrique & Vicente 2010) our answer is: *it is speech*. It is speech in two important senses. First, it is a productive rather than re-creative activity. Second, its primal use is in making *speech acts*: asserting, questioning, insulting etc. (p. 247, emphasis in the original).

Machery (2018) writes: "Most speech act types are found in inner speech. In inner speech, one finds assertive speech acts, directive speech acts, commissive speech acts, and expressive speech acts (Searle, 1969)" (pp. 262-263). I agree with Machery. To illustrate the point, suppose we adopt Searle's (1976) taxonomy of illocutionary acts into five mutually exclusive and jointly exhaustive classes. The following is a list of the five classes, plus a definition of the class. I have added to each one an example of a speech act which, I suggest, could easily have been uttered in inner speech:

- *Representative* or *assertive*. The speaker becomes committed to the truth of the propositional content. Example: asserting (silently to oneself), "It's raining."

- *Directive*. The speaker tries to get the hearer to act in such a way as to fulfil what is represented by the propositional content. Example: (silently) commanding (oneself to): "Hurry up!"
- *Commissive*. The speaker becomes committed to act in the way represented by the propositional content. Example: (silently) promising (oneself): "I'll finish the washing up tomorrow."
- *Expressive*. The speaker simply expresses the sincerity condition of the illocutionary act. Example: saying (silently, to oneself), "He's arrived – thank God!"
- *Declarative*. The speaker performs an action just representing herself (silently) as performing that action. Example: "I hereby promise not to touch another drop."

In short, those philosophers of inner speech who take it for granted that inner speech utterances are, at least sometimes, speech acts, seem to be fully justified.

Section 5. A case study: deciding as an inner speech act

Readers might be persuaded by the previous section that inner speech utterances can sometimes function as inner speech *acts*, but still not be convinced that the contents of a speech act, silent or otherwise, cannot be represented in a mental state. In this section I want to illustrate my argument with the case of deciding something.[58]

The term 'deciding', of the kind I am interested in here – the kind associated with an episode of deliberation – has two meanings. It can refer to the activity of trying to bring it about that one comes to a decision (call this *deciding-as-activity*), and it can refer to the event of deciding something in particular (call this *deciding-as-action*). Let's illustrate the difference with an example that includes both senses of the term. Suppose I am trying to work out the best time to leave home for a meeting; that is, I consciously engage in an episode of deliberation with the explicit purpose of trying to decide, at the end of it, when

[58] It has recently become clear to me that deciding something can be a more complex phenomenon than suggested by this section. For example, *granting consent* appears to be a type of deciding which involves both an internal (private) aspect, and an external (public) aspect. Arguably, two decisions are involved (see Goodin, 2024). As far as I can see my account of deciding something can accommodate the phenomenon of granting consent, but I do not have the space to justify that claim here.

to leave home. How do I go about this? I might ask myself a few pertinent questions: How much time ahead of the meeting do I want to arrive? Which mode of transport will get me there fastest or most reliably? How long will the journey take? etc. All these questions will receive a response of some kind. (Chapter 5 explains why the questions and answers which comprise an episode of deliberation are all *actions*. By my lights they are also inner speech *acts*, but I am not going to argue for that here. For an argument for this claim see Kompa & Mueller, 2022).[59] This inner dialogue is an example of deciding-as-activity. Suppose, at the end of this episode of deliberation, I decide that 10.30 is the time I should leave home. This is an example of deciding-as-action. We can now ask two questions. First: What form will the event of my deciding-as-action take? Second: is the content of that event – the decision itself – represented by a mental state before I become conscious of what I've decided? In answer to the first question, I claim that it will take the form of an inner speech act, the function of which in this case, and in line with Geurts' (2018) proposal (see Section 3), is to entrain my commitment – to commit myself to leaving home at 10.30. In answer to the second question, I claim that the content of this inner speech act (i.e. the decision) is not represented by a mental state before the performance of the speech act. I argue for these two claims below.

Exactly what I say to myself when I perform the speech act might vary. For example, if the meeting is very important, and being late is of great concern to me, I might make my commitment a strong one. I might say to myself "I must leave the house at 10.30 at the absolute latest." If I am more relaxed about being on time, I might say something like "Around 10.30 should be fine." If my main concern is leaving myself an acceptable margin of error, I might say, "10.30 will get me there with lots of time to spare". But whatever I say to myself, it won't function as a decision until and unless I utter the words and thereby give form to the force of the utterance. Until I say the words the decision is not made. Saying the words *is* making the decision; that's why it's a speech *act*.

[59] They write: "We will argue that *pragmatically expanded* inner speech may assist deliberative thinking and reasoning by re-purposing pragmatic principles that also guide conversations with others" (p. 2, original emphasis). It is notable that when explaining their use of the term *'pragmatically'* they refer explicitly to Austin (1962).

You might be tempted to think, in line with traditional speech act theory, that before I say the words to myself, I already have 'in mind' the force-neutral content 'leaving the house at 10.30', and that therefore the decision to leave the house is represented by a mental state before I speak. But this is wrong. It might be true that 'leaving the house at 10.30' is represented by a mental state (or states) of some kind before I speak, but it is wrong to suppose that that mental state is a *decision*. It isn't a *commitment* to leave the house at 10.30, or anything like one, because it is force neutral. The very same content might be represented by a mental state as a constituent of a different thought, such as "I've noticed that my neighbour always leaves his house at 10.30". You might also be tempted to think that since, by hypothesis, I embarked on this episode of deliberation (this deciding-as-activity) with a view to coming to a decision, the *intention* to decide when to leave the house is also a mental state before I actually make a decision. I agree. But a mental state which is merely the *intention* to decide when to leave the house is not a decision *when* to leave the house, so once again, that mental state is not the decision to leave the house at 10.30 either. Finally, you might be tempted to think that, just before I bring together the force-neutral content (leaving the house at 10.30) and the force (commitment) with which I make the utterance, I am conscious of what I am about to do, and you might think that this consciousness is represented by a mental state. But this is also wrong. To be genuinely conscious of what you are about to do is equivalent to being conscious of what you have decided, and you can't be conscious of what you've decided until you have combined the force-neutral content with the force. Another way of putting this, which reintroduces a term for Section 1, is to say that all these aspects of the speech act – the force-neutral content, the force and the intention – are different kinds of *information* represented in the head of the subject, but that neither individually nor collectively do these things amount to a *decision*. To execute the decision requires them all to come together via the production of a linguistic utterance, and the performance of a speech act.

Another possible objection. One might think that any of these speech acts could be performed just as successfully by a much more minimal inner speech utterance – something like "10.30 it is" or even just "10.30". If that's right, then the language production process might appear to be doing almost no work at all, in which case there is no reason to think that the meaning of the speech act is not represented by a mental state before the words

are uttered – the mental state of intending to choose 10.30 as the time to leave the house. If that's right then there is, after all, a mental state which represents the meaning of the speech act. And this contradicts my key claim.

But this is too quick. If I make my decision-as-action by saying nothing but "10.30" then this is simply an abbreviated version of the full speech act: "I will leave the house at 10.30". We know that to be true because we know that saying "10.30" on its own, spoken silently or aloud, could mean any number of things; we need the context to make any sense of it, and the context in this case is that of deciding when to leave the house to catch a train. If I make my decision-as-action by saying nothing but "10.30" then part of the function which this speech act is performing, as well as the function of deciding, is the function of choosing one of the options previously considered during the episode of deliberation. Since the option being chosen has already involved the language production process (because, of course, the deliberation was conducted in natural language) the option being selected already possesses more meaning that merely "10.30". The speech act "10.30" has the function, in effect, of *pointing* at a natural language expression (the option being held in short term memory) and saying, "That one". The meaning of "10.30" on its own doesn't have enough content to achieve what the speech act has the function of achieving – committing the subject to leaving the house at 10.30. The missing content is present in the decision to select that option (via a speech act) in virtue of the language production process that generated that option in the first place (in the course of deliberating about what time to leave the house to catch the train). Phenomenologically, it may be that, before the decision is made (to select the option "10.30" by the performance of the inner speech act), I have a sense that I am close to fulfilling my intention; if so, then presumably that sense is represented by a mental state. But that sense, and therefore the state which represents it, is not the same as the decision itself.

In summary, in cases like this, deciding-as-action is like any other speech act: there is no mental state which represents the full meaning of the act because the full meaning of the act is determined by the performance of the act, and the performance of the act is achieved by the uttering of the words (in a particular context, with a particular force, using some particular natural language expression).

Section 6. What kind of knowledge do we have of our own thoughts?

It's time to address an important epistemological question which arises from the thinking-as-speaking thesis: how should we understand the relation between inner speech and knowledge of our own thoughts? It might seem to follow from the thinking-as-speaking thesis that we know our own thoughts *merely* by hearing our own speech (whether overt or silent). But such a view implies that we come to know our own thoughts in just the same way as we come to know the thoughts of others, and that seems, intuitively, wrong.[60] Even if we know what E. M. Forster is getting at when he says, famously, "How can I tell what I think, until I see what I say?" (1927), nevertheless the sense in which I 'discover' what I think by speaking is, intuitively, quite different from the way I discover what you think by hearing you speak, and interpreting what I hear. But what, exactly, is the difference? In this section I argue that the idea of practical knowledge (knowledge-how) can be extended from *speaking* (see Chapter 4, Section 4) to *thinking.* In fact, this is what you might expect if the thinking-as-speaking thesis were correct. And I follow Roessler (2015), in arguing that the special kind of knowledge that we exercise when we exercise our 'knowledge how to think' is the kind of knowledge we have when we are "alive to what we are doing" (Ryle, 1949, pp. 168-9).

The idea that we come to know our own thoughts in much the same way as we come to know the thoughts of others can be traced back to Ryle (1949):

> 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 (p. 176).[61]

[60]However, see Carruthers (2009), where he defends a version of this view.

[61] Ryle (1949) adds that we "learn to make this study of our own talk from first taking part in the public discussion of anyone's talk" (p. 176), and in saying this he appears to be implicitly endorsing the Vygotskian view of inner speech – see Chapter 2, Section 1 for details.

But, as Roessler (2015) observes, earlier in his book Ryle makes some comments on the nature of "serial operations", and on the nature of our knowledge of those operations, and these comments can be used to form the basis of an alternative (but also compatible) account of the nature of the knowledge we have of our thoughts.[62] Roessler develops such an account in his paper, and it is to that account that I turn now. (Roessler's paper has some other interesting goals too, but I will ignore them in what follows.) What Ryle says about serial operations is that there are some kinds of activities the execution of which involves steps or stages. One step may relate to another as a means to an end (e.g., mixing the dough to make a loaf of bread) or as several parts of a whole (e.g., playing consecutive movements of a sonata). Carrying out these activities, he says, 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" (1949, pp. 168-9). According to Roessler,

> 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' (p. 4).

As examples of a "complex task" Roessler gives mental activities such as "deliberating whether to buy a particular book, or multiplying 79 by 45" (p. 16). He argues that "Ryle's … 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". Putting these two ideas together (i.e., 'complex mental activities' and 'practical knowledge') Roessler proposes that the notion of practical knowledge can be extended from cases of physical tasks, via the case of speaking, to the case of thinking. I take him to mean by this that thinking is one of those activities identified by Ryle the execution of which involves steps or stages, and that when

[62] Hinshelwood (2023) makes a similar claim. His paper argues that "when one is intentionally doing something, one's representation of it as a goal to be accomplished must also be *knowledge* that one is intentionally doing that thing. [And also, that] this knowledge must itself be one's intentionally doing that thing" (p. 1). I would like to explore further how Hinshelwood's ideas complement and expand on Ryle's somewhat vague idea of "being alive to what one is doing" (see main text below), but I don't have time/space here.

we engage in thinking we exercise knowledge how to do it; specifically, we exercise a knowledge of the steps and stages involved in *thinking* (i.e., judging, inferring, deducing, supposing, etc.). In other words: we exercise our knowledge of how to think *by* exercising our knowledge of how to speak. On this account, and in line with my use of action theory in Chapter 4, we should understand thinking-as-speaking as involving doing something basic (exercising knowledge how to speak) as a means *by which* to do something non-basic (exercising knowledge how to think). (I don't claim for a moment that Roessler would put it like this.)

Roessler raises a problem for this idea, of a kind we have encountered before in this thesis – in fact it is the subject of Chapter 5. There, Frankish (2018) allows that when we ask ourselves questions as part of an episode of silent deliberation, we should grant that the events of asking the questions qualify as intentional actions (the product of Type 2 thinking), but that the responses that come to us do not qualify as actions, but rather are mere events (the product of Type 1 thinking). Roessler raises a similar issue. His example is the case of multiplying 79 by 45. If all goes well, you will conclude or realize that the answer 3555. But as Roessler agrees, "The *intention* to conclude, or realize, that something is the case would be rather blatantly irrational" (p. 19. Emphasis added). It makes no sense to say that you *intended* to arrive at the answer 3555, where that implies that you knew 3555 was the answer all along. But if, by hypothesis, you didn't know the answer all along, arriving at it can't be the exercise of knowledge, practical or otherwise; you didn't *know* the answer until you arrived at it as the conclusion to your calculation. (And something similar can be said for his other example – deciding whether to buy the book you deliberated about buying; you didn't *know* you were going to decide to buy it until you decided.)

Roessler's response to this worry is the following:

> The way to overcome this obstacle [to believing that you are exercising knowledge when you arrive at the answer, or the decision]…, 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 (p. 19. Emphasis added).

On this account, we are justified in describing the event of concluding or deciding as the exercise of knowledge because the event is *grounded* in the subject's knowledge of how to perform the activity (of calculating or deciding). And the exercise of that knowledge is *intentional* because the subject purposefully engaged in the activity (of calculating or deciding) in the first place. If my arguments in Chapter 5 are sound, this line of thought is another argument for thinking that reasoning is a wholly intentional activity, and not merely partially intentional.

This way of extending the exercise of practical knowledge from speaking to thinking is significant, because without it one can go astray. Consider this from Hornsby (2005) for example:

> Now if one thought of semantic knowledge as practical knowledge but forgot that it was exercised along with everything else required for speaking, then one might get the impression that explanations of a speaker's actions had to be confined to explanations of their saying whatever they do. And one might get the impression that a person's voicing her thought was a matter of her blurting something out: speaking might seem to be something over which we lacked control. The truth of course is that *it can be a matter for choice and reflection what to say and how to say it*, and that a speaker's attention may be occupied both in deciding what to say and in actually saying it (p. 126) (Emphasis added).

The problem with this passage is the phrase I have emphasised: "…*it can be a matter for choice and reflection what to say and how to say it".* But if the thinking-as-speaking thesis is right, then the question of what to say and how to say it *cannot* be "a matter for choice and reflection", on pain of circularity: choosing is the product of reflection, and reflection is a form of thinking. However, Roessler's Rylean account provides a way to disarm Hornsby's worry – that a subject might "lack control" of her acts of thinking aloud. In thinking aloud, the subject has intentionally engaged in an activity she knows how to do. On Roessler's account, as the subject performs the activity of thinking-as-speaking she exercises her

knowledge how to do it. Her control comes from her ability to monitor her performance as she goes along, measuring it against whatever criterion counts as success for the speech act she is performing. Whether this account satisfies Hornsby's definition of "control" is, of course, a moot point. But I suggest that to require any stronger measure of control would be unreasonable. We exercise as much control over thinking-as-speaking as we do over anything else we do which we know how to do – a point I made in Chapter 4, Section 5.

In summary. In this section I have attempted to answer the question: what kind of knowledge do we have of our thoughts? The answer, I have said, is anticipated by Ryle and developed by Roessler. Thinking is one of those activities that is characterised by steps and/or stages, and when we are engaged in it, we are "alive" to this fact, and we are exercising our knowledge of how to do it. That is the kind of knowledge we have of our thoughts: we have knowledge of what we are doing when we are engaged in the activity of thinking. And as with the exercise of any other skill, we can monitor our performance as we exercise it, and adjust our performance according to how well we are meeting the goal we set out to achieve, intentionally, at the start of the activity. In this sense, we have as much control of the exercise of the activity 'speaking-as-thinking' as we do of any other skilled activity.

Section 7. The motor-sensory view of inner speech revisited.

At the beginning of the thesis, in Chapter 1, Section 1, I made the claim that, while my arguments concerning the nature of conscious thought are more *plausible* if the motor-sensory view of inner speech is right, they are not hostage to it. In this very brief section, I want to make good on that claim. Recall that the Motor-Sensory view is that inner speech is an important sense embodied, and involves physical processes that unfold over time, and the Abstract view, by contrast, is that inner speech involves symbolic and abstract representations, divorced from bodily experience. (Recall also that this distinction has nothing to do with the *content* of the inner speech utterances, but only the *format* of the representation – whether the representation is abstract and symbolic or concrete and physical.) The question I want to address now is this: How much of what I have argued for

in this thesis could you accept if, unlike me, you endorsed the abstract view of inner speech and rejected the motor-sensory view?

In Chapter 4 I accepted that, as argued by Gregory (2020), some inner speech utterances do not qualify as actions but rather should be understood as 'reactions' to the environment a subject finds herself in. But I also argued that some other inner speech utterances are intentional actions, performed deliberately, effortfully and with as much control as any other intentional action. When we engage in deliberation, for example, I claimed that inner speech utterances are the basic actions by which we bring about the non-basic action of deliberating about something. This proposal, I said, is more *plausible* if inner speech utterances involve motor-sensory processes, because it allows us to see inner speech as more fully continuous with overt speech which, uncontroversially, involves motor-sensory processes. It allows us to understand the phenomenon of inner speech within the wider context of embodied cognition generally. Nevertheless, it is not *inconsistent* with the view that inner speech utterances are sometimes intentional actions to suppose that they are intentional actions which do not involve motor sensory processes. A plausible version of such a view might go as follows.

The capacity for inner speech, as Vygotsky says, is the result of the internalisation of the capacity for overt speech. But how far this process of internalisation goes is a moot point. We don't have to follow Vygotsky in thinking that it goes so far that we are able to think in 'pure meanings'. But we can believe, if we choose, that it goes further than supporters of the motor-sensory view think. We could believe that we are able to think in our natural language without the need to embark on the final stage of language production – the articulation stage. There is nothing in a view like this which undermines the idea that the exercise of our capacity for inner speech is, as argued earlier in this chapter, a special kind of practical knowledge, something which, once we have learned, we are able simply to do. And this is all we need for inner speech utterances to qualify as the basic actions *by which* we do something non-basic, such as to deliberate about something. A view like this is entirely consistent with the thinking-as-speaking thesis, because all we need for that view is the idea that inner speech utterances are inner speech acts and, as we've just seen, the

abstract conception of inner speech can endorse that idea without endorsing the motor-sensory view.

One might wonder how the matter should be decided. In Chapter 2, Section 1, I provided empirical evidence from Loevenbruck *et al* (2018) *against* the abstract view. But of course, others adduce empirical evidence in its favour (e.g., Oppenheim & Dell (2010).) I am content to let the scientists and the psycho-linguists decide. The point of this section is to insist that the arguments of this thesis are consistent with both the motor-sensory and the abstract view of inner speech.

Section 8. The bigger picture

The main thrust of this chapter has been to buttress the speaking-as-thinking thesis by arguing that there is no determinate thought 'in mind' before the performance of the speech act which expresses it. I argued that a token thought is not an object of any kind, and I suggested that a thought should be understood as an action: the action of performing a speech act. In Chapter 8 I explore the idea of skilled actions in more detail, but I want to end this chapter by preparing the ground.

In Chapter 3, Section 4, I described how Vicente & Jorba (2019) make two suggestions which push in the direction of the idea that a thought is not represented by a mental state before it is uttered. First, they suggest that a pre-linguistic thought undergoes a "representational re-description" when it takes on a linguistic form. Second, they reject the idea that pre-linguistic thoughts have their contents fully determined before they are translated into natural language, and they suggest that "thought content needs to be brought to consciousness for its content to be fully determined" (p. 26). When you put these ideas together it is hard to see how they leave room for the idea that a token thought exists, fully determined, as a mental state, before it is uttered. Nevertheless, as I interpret them, Vicente & Jorba still assume that, metaphysically speaking, a thought is some kind of object – a representation of some kind. They seem to understand the "representational re-description" of the thought during its transition from pre-linguistic to linguistic format as merely some kind of *precisification* of what is, ultimately, the same thing. To recap: I

disagree. I have suggested that a thought is not an object of any kind, neither (a) a non-linguistic representation 'in the head' nor (b) an expression in natural language, nor (c) an entity that goes through a transformational process from one to the other. A token thought, I have suggested, is really a kind of action: the action of performing a speech act. The purpose of this section is to buttress my position by showing that, while it goes against philosophical orthodoxy, it finds an echo in some of the ideas of Tim Crane (2017), specifically his ideas on the nature of belief. I will start by sketching the standard view of belief which Crane rejects; then I will introduce his concept of a 'world view'; then I will describe his alternative conception of the nature of belief.

The standard view of beliefs, according to Crane, is the view that,

> …beliefs are mental states that represent the world and the way they represent the world is by having representational or intentional content; their intentional content is a proposition, something which is true or false. This view is sometimes summed up by saying that belief states are relations to propositions, and implies that for each belief you have there is a distinct proposition to which you are related… Beliefs, it is often said, are 'individuated by their contents'. I take this to mean that individual belief states are distinguished from one another by the propositions which give the ways they represent the world… The distinct [mental] states are then characterised in various further ways — as dispositions to behave, or to utter sentences, as aiming at truth, as being fine-grained or coarse-grained in their contents, as being externalistically individuated etc. (p. 3).

Crane begins his criticism of this view with some gentle mockery: if beliefs are numerically distinct mental states, 'individuated by their contents', we should in principle be able to count them. How many do we have? The absurdity of the question, he suggests, should make us suspicious of the standard picture.

To develop his alternative view, Crane introduces the idea of a subject's 'worldview' – which he characterises as a subject's "unconscious representation of the world, a representation that aims to 'depict the world more or less as it is'" (p. 9).[63] The unconscious nature of the

[63] Crane is quoting Wollheim (1999).

worldview is very important – we are not (we cannot be) conscious of everything we believe (our "entire doxastic orientation towards the world") all at the same time. The propositions we *consciously* entertain, and which we ascribe to ourselves as our beliefs, are "models" of some aspect of our worldview – some aspect of our worldview that is salient at the time we consciously entertain the proposition. Typically, we become conscious of that proposition by saying something, either aloud or silently to oneself. What does Crane mean by "model"? "A model is a (concrete or abstract) object used to draw attention to some features of a system under investigation, and to make the study of the system more tractable… Models idealise and simplify in order to highlight some structural feature of the thing being modelled" (p. 8-9). So, the content which is represented unconsciously is different from the content which is represented consciously, in the form of an utterance: "[U]nconscious states of mind do not have content in the way that conscious states do" (p. 5). When the thing being modelled is ourselves, as when we attribute to ourselves a belief, the proposition we utter models some aspect our worldview. Models, unlike propositions, are not true or false; they are more or less effective at drawing attention to whatever feature of the system being modelled is relevant in the context.[64]

On the standard view, all the facts about what we believe are there (in our head) waiting to be discovered, and if we are not sure what we believe the problem is a purely epistemological one. But on Crane's view, there isn't necessarily a fact of the matter about what we believe; there is no sharp distinction between 'discovering what you think' and 'making up your mind'.

> Discovering what you believe can resolve indeterminacy and unclarity in your worldview, and producing a conscious judgement that settles things as far as you are concerned. This view treats the determinate content of a belief as 'coming into being as we probe'[65]… *The clearest*

[64] Korsgaard (2009) arrives at a similar view from a consideration of the nature of our rationality. Regarding propositions she says "Propositions are true when the concepts that appear in them are applied correctly; but I do not suppose that [our] ways of conceptualizing the world are themselves simply true or false. I think of them on the analogy of maps, since they are devices that enable us to find our way around" (p. 36). And as regards belief, she says "In other words, what I am proposing is that, for a rational animal, believing itself is an active state, it is *doing* something, it is an activity: it is *representing* the world to yourself in a certain way" (p. 37, original emphasis).

[65] Crane is quoting Dummett (1959) writing about mathematical reality.

> *way in which the content of a worldview can become the content of a conscious judgement is when one says something, either aloud or to oneself*… The central phenomenon here is interpretation. To model an aspect of a worldview with a proposition is to interpret it. To bring a belief to consciousness is to interpret something about your worldview, where interpretation is a matter of taking a system — you — and making sense of how you represent the world in general (pp. 13-14, emphasis added).

I hope it will be obvious that many of the things Crane says about belief, and how we bring a token belief to consciousness, resonate with some of the things I have said about thought, and how we bring a token thought to consciousness.[66] On Crane's view it makes no sense to say that the belief exists, like an object in a box, before it is brought to consciousness. I say the same thing about a token thought. Crane says the content of the belief brought to consciousness is a 'model' of some aspect of the subject's world view – whatever aspect (or aspects) are most relevant to the context in which it was brought to consciousness. As a model it is not subject to the test of truth or falsity, but to whether it was successful at achieving what was intended in the circumstances. Similarly, I say that a token thought is the action of performing a speech act, and as an action it is not subject to the test of truth or falsity, but to whether it was successful at achieving the purpose for which the speech act was performed. Crane speaks about "interpretation"; on his view the event of bringing a token belief to consciousness, typically by saying something, either aloud or to oneself, involves interpreting some aspect of one's world view. I talk about "production"; on my view the event of bringing a token thought to consciousness involves the act of producing a linguistic utterance – a speech act – which will achieve the purpose for which it was intended. What both ideas have in common is the idea of *creation* rather than *recreation* or *reproduction*.[67]

I want to end this section by putting both my account of thought, and Crane's account of belief, into some scientific context. In a recent book, *The mind is flat: The Illusion of mental*

[66] I don't want to imply that Crane would endorse the arguments of this thesis.

[67] Allow me a fanciful analogy: Thinking is like jazz. The jazz musician has some musical 'information' (a phrase or tune) 'in mind'; she has practical knowledge of the rules of improvisation; and she just starts playing. The (determinate) music doesn't exist before the performance. (Contrast this with reading from a score, which is more like performing someone *else's* thoughts; you might say it involves 'acts of music-making' not 'music-making acts'.)

depth and the impoverished mind (2018), Nick Chater draws on new research in neuroscience, behavioural psychology and perception, to argue that the intuitive view – that our minds have determinate contents – a hidden world of beliefs, feelings, motives, fears and desires, all waiting to be discovered, if only we had the tools to do it – is wrong. The 'contents' of our minds, he says, are not like documents in a filing cabinet, waiting to be taken out when needed. They come into existence the moment we become conscious of them.

> We generate our beliefs, values and actions in-the-moment; they are not pre-calculated and 'written' in some unimaginably vast memory store just in case they might be needed. And this implies that there is no pre-existing 'inner world of thought' from which our thoughts issue. Thoughts, like fiction, come into existence in the instant that they are invented, and not a moment before (p. 5).

We intuitively assume, he says, that we can introspect our minds – scrutinize the inner world to find out what is there – in something like the way we can use our powers of perception to understand the external world. But this is to misunderstand the nature of introspection:

> [I]ntrospection is a process not of *perception* but of *invention*: the real-time generation of interpretations and explanations to make sense of our own words and actions. The inner world is a mirage (ibid.).

(Later, he writes: "I have now, somewhat reluctantly, come to the conclusion that almost everything we think we know about our own minds is a hoax, played on us by our own brains" (p. 15). This might remind you of something Wegner says about the will, quoted in Chapter 2, Section 4: "The experience of conscious will is a marvelous trick of the mind…" (2003, p. 68).)

It might seem that Chater goes much further than Crane (or me), since it might seem that Chater is denying that there is any *content* in the mind at all, any worldview which informs what we do and what we say. But it seems to me that this might more a matter of

rhetorical style than a deep difference of view. For example, Chater talks about each of us having a "mental tradition" based on our prior thoughts and experiences.

> New actions, skills and thoughts require building a rich and deep mental tradition; and there is no shortcut to the thousands of hours needed to lay down the traces on which expertise is based. And for each of us, our tradition is unique: with those thousands of hours each of us will lay down different traces of thoughts and actions, from which new thoughts and actions are created (p. 11).

It seems to me that what Chater means by "rich and deep mental tradition" is not so very different from what Crane means by a "worldview", and that both develop in much the same way as Chater describes. I am not suggesting that both writers are saying the same thing – they are not – but I do want to emphasize Chater's use of the term "invention" in the passage above. The word, I suggest, is a member of the same family as "interpretation" and "production" in as far as it involves the idea of "creativity". The shared idea (across Chater, Crane and me) is that the stream of our conscious mental lives is not a matter of *accessing* determinate content, but a matter of *creating* something relevant and useful through the activity of thinking (broadly construed).

Section 9. A taxonomy of thinking: activity, actions, accomplishments and achievements.

As previously stated, the main purpose of this chapter has been to argue that a token thought – one individuated by its content (in the sense of its meaning) – does not exist 'in the head' of the thinker before she performs a speech act. And I have argued that in virtue of that fact we should conclude that a thought is not an object of any kind, but an action. In this section I want to make that claim a little more precise by borrowing (again) from those who write about theories of action, and specifically (again) from Jen Hornsby – on this occasion her (2012). In that paper she argues that human action can only properly be understood if we recognise (a) the existence of an ontological category of "process or activity" (henceforth I will just use the term activity), and (b) that this is not a category of

particulars.[68] This latter point is particularly important to Hornsby because she wants to create the logical space for a causal account of human behaviour that is not one based purely on events which, unlike activities, *are* particulars. (Having made the logical space for it, she does not say much more about agency in this paper, and in any case that aspect of her paper is irrelevant to my interest in it. So I will ignore it in what follows.) Her argument (which I find very persuasive) for the ontological claim that activities are not particulars, involves making some careful distinctions between the different ways we conceive of, and talk about, bodily actions. In this section, I offer a short summary of her position, and then show how her account maps onto my claims about thinking-as-speaking. I argue her position can be exploited to buttress my suggestion that a token thought is a kind of action.

The so-called 'standard story' of action, according to Hornsby, has it that "Human agency is … supposed to be fully accommodated in the natural world when the causality it involves is treated as *event-causality*" (p, 233, original emphasis). According to the standard story, processes and activities are treated as events which are extended over time. But this, says Hornsby, generates a problem:

> Suppose that an agent raises her arm and her arm goes up. Those who treat causation as always linking events (broadly understood) say that mental states of the agent cause an event of her raising her arm and an event of her arm's going up. They then face a dilemma. Is there only one event here caused by the mental states? Or does the event of her raising her arm cause the event of her arm's going up? Evidently one cannot say Yes to both questions: if the raising event just is the arm's going up, then it is not a cause of the arm's going up. Yet there is pressure to say Yes to each of the questions (p. 234).

I will not rehearse here the ways in which the standard story of action fails, according to Hornsby, to resolve the dilemma. I will move straight to her solution, which involves the idea that an agent's actions are the product of her being engaged in *activity*. The important thing about the activity the agent is engaged in is that, unlike an event comprised of it,

[68] Steward (1997) makes a similar argument. But in her (2012a) Steward credits Hornsby with making an improved version of her argument (see p. 388, Note 19) so I will mainly refer to Hornsby's (2012) account.

which might be described equally as 'her raising her arm' or as 'her arm's going up', the activity is not a particular.

> [O]nce the agent's role in her arm's coming to be up is acknowledged, she will be seen as having been engaged in activity; and an event of her raising her arm and an event of her arm's rising can then be seen as alike comprised from a bit of activity – a bit of the type of activity that someone engages in for just so long as she is raising her arm and her arm is going up. The duration of any such bit of activity is the duration of an event which might be described equally as 'her raising her arm' or as 'her arm's going up'. There are not two events here, inasmuch as her arm's going up is what she is causing at any moment at which she is raising it (p. 235).

The claim that the agent was engaged in the activity of raising her arm is a claim that something was going on, but it is not a claim that a particular event (nor a particular of any other sort) was going on. The activity itself is not, as Hornsby says, "anything of a countable sort", and as such it can't participate in casual explanations which involve only events, which *are* things of a countable sort.

Activity, then, is a noun, but not a count noun. To help make her point Hornsby makes an analogy with "another brand of non-count nouns – those which name types of stuff. Names of stuff don't pick out particulars" (p. 237). Take the stuff 'milk' for example; milk comes in bottles, cartons, puddles and spills. But to say that 'There is some milk around here somewhere' is not to say that there is a particular bottle, carton, puddle or spill here. Of course, it can only be *true* that there is some milk around here somewhere if there is found to be an instance of a bottle, carton, puddle, spill or some other token sample of milk around here, at the relevant time, but to say that 'there is some milk around here' is not to *say* that there is such an instance. The analogy with 'stuff' helps explain what Hornsby means by saying that an event is "comprised of a bit of activity". Just as a particular puddle of milk on the kitchen floor is not identical with milk (the stuff), but *constituted from* the stuff of milk, a particular walk that I walked this morning is not identical with walking (the activity) but constituted from the activity of walking. (This example comes from Crowther, 2011, called *The Matter of Events*. The paper contains a detailed analysis of the analogy

between spatial and temporal notions, and how, "Both substance-stuffs and time-occupying stuffs, respectively, fill out space and time in the same way" (p. 17).)

Returning to the nature of activity, Hornsby grants that there can be *stretches* of activity and that, when the term is used in this way, *an* activity can be individuated as a particular.[69] For example, if I start reading at 10.00 and continue until 12.00, there is a 2 hour stretch of my reading, and this can be treated as *an* activity, in one sense of that word. But this sense of the term, according to Hornsby, is just another way of referring to an action – an action which will have taken place if someone started doing something at some point in time and then stopped doing it. An action which is characterised as a stretch of activity is one which is "comprised" of activity.

Two other ways to characterise an action are as an accomplishment and as an achievement. Accomplishments are actions which come to be over – completed – with the coming of the relevant "product, upshot or outcome" (p. 240).[70] Some examples: Sarah's writing an essay; Neil's making an omelette; Jane's painting the porch. So, what's the difference between this sort of action and the sort characterised as a stretch of activity?

> What it takes for there to be an accomplishment of any sort is determined by the sort of endpoint required for an accomplishment of that sort, whereas for there to have been a stretch of activity of any kind, it is only required that there no longer be activity of that kind: a stretch of activity is over simply when it has stopped (p. 240).

Like stretches of activity, accomplishments are comprised of activity. But that's not true of achievements, which are characterised by Hornsby as "actions that may be conceived of as punctate", meaning that, unlike accomplishments, they are conceived of as lacking temporal duration. Some examples: Jerry's winning the race; Mary's finding the book; Anna's arriving home. It makes no sense to ask about each of these actions: How long did the winning/finding/arriving go on for? Since these actions are not conceived of as having

[69] Hornsby writes, "(I use 'stretches' as a catch-all term: depending on the nature of an activity, we speak more readily of bits, or of pieces, or of bouts, or of spells, or of stints.)" (p. 239)
[70] Hornsby is quoting Mourelatos (1978) when she writes this, and credits him with anticipating some of the ideas in her paper.

duration they cannot be conceived of as comprising activity, but nevertheless, they cannot happen in the *absence* of some activity; what achievement verbs do is to *report* the product, upshot or outcome of some piece of activity. "Achievements are not comprised of activity, as stretches of activity and accomplishments surely are. But activity is a *sine qua non* of achievements nonetheless" (p. 242).

The above, then, is a very brief sketch of Hornsby's taxonomy of action terms: activity, accomplishment and achievement. What I want to do now is reflect on how such a taxonomy maps on to the thinking-as-speaking thesis. My proposal is that it maps on to it very well. Hornsby makes an important point about the term 'activity' which has a parallel with the term 'thinking'. Something is considered *an* activity, she writes, by virtue of its being a *type* of activity. The examples she gives are strolling, walking, reading, etc. I suggest a parallel is to be found in respect of the term 'thinking'. Something is considered *thinking* by virtue of its being a *type* of thinking. Examples include reasoning, reflecting, considering, deliberating, imagining, wondering, calculating, picturing, etc. Put simply, thinking is an activity. (One might be tempted to say: Thoughts are comprised of *mental* activity in much the same way that bodily actions are comprised of *bodily* activity. But I don't like this way of putting it because it reinforces a distinction my thesis is designed to undermine.)

Just as particular actions can be conceived of in different ways – as stretches of activity, as accomplishments, or as achievements – so can particular thoughts. Suppose I start thinking about what to cook for dinner at 17.00 and stop thinking about it at 18.30. We might characterise this thought ("He thought about what to cook for dinner for an hour and a half") as a stretch of thinking. A thought can also be characterised as an accomplishment, and as with bodily actions, the type of accomplishment is in turn characterised by the type of activity of which it is comprised. As Hornsby puts it regarding bodily actions, "terms for accomplishments always introduce activity of a particular type" (p. 240). She goes on:

> One might give the name of *directed activity* to those activities, which, while they are under way, have as their anticipated development the product, upshot or outcome of some accomplishment (p. 241).

Thinking, I have argued elsewhere in this thesis, is also, very often, a "directed activity": when we deliberate, for example, we have as our goal the 'product, upshot or outcome' of coming to a decision, reaching a conclusion, solving a problem, etc. (See Chapters 4 and 5). When thoughts are conceived of as accomplishments then, as with other actions, "terms for accomplishments always introduce activity of a particular type". If John is deciding what to cook for dinner, the product of his deciding will be a decision as to what to cook. If Mary is choosing between two courses on the menu, the product of her choosing will be a choice. If Bob is trying to solve a problem or a puzzle, the product of his trying will be a solution. And so on.

What about achievements? These too have their parallel in the domain of thoughts. Some mental actions are conceived of as punctate. Consider the following examples. Tim calculated the answer to the maths problem: 42. Liz understood the question immediately. Fred judged the distance to be around 10 meters. These events, at least on one reading, have no duration. (It would be odd to ask Tim, "When you arrived at the answer 42, how long did the arriving at the answer 42 go on for?" It would be odd to ask Liz, " When you heard the question, and you understood it immediately, how long did understanding the question immediately go on for?" It would be odd to ask Fred, "When you judged the distance to be 10 meters, how long did the judging the distance to be 10 meters go on for?"). Nevertheless, they qualify as actions because they are the product, upshot or outcome of an activity that was engaged in intentionally – calculating, understanding, judging. Note that Hornsby describes these actions as "conceived of" as achievements, and she writes that, "[W]hat achievement verbs report (at least when used in saying what has been done intentionally) is the product, upshot or outcome of some piece of activity" (p. 241).[71] I want to emphasise that, just because this is how these verbs are sometimes *used* does not mean that the actions they refer do not, if fact, take time to perform. I will return to this issue, and in particular the case of judgment, in Chapter 9, where the alternative view plays an important role in the ontology of conscious thinking proposed by Matt Soteriou (2013).

[71] They are not always used this way. See Hornsby (2012, FN 16, p. 244) for the details.

In summary, thinking is an activity, something engaged in by a subject, and thoughts are actions comprised of the activity of thinking. Thoughts, like bodily actions, can be characterised in different ways: as accomplishments or as achievements. The fact that Hornsby's taxonomy of bodily actions provides a parallel taxonomy of mental actions supports the proposal that thoughts are a species of action. Or so I claim.

(You might have a worry that while it's fine to say that thoughts have content, it sounds odd to say that actions have content; maybe to claim that thoughts are a kind of action is to make a category error. Peacocke (2023a) anticipates this worry and rejects it:

> To dispel this confusion, it is important to see that the category of action is a determinable category, of which there are many determinates. An action can be a kick, a heist, an election, or something else. A mental action can be a judgment, a decision, a recollection, an imagining, or something else. These more determinate categories of actions can clearly have contents. Take the category of judgment: a judgment must have a content to be a judgment at all. Similarly, each decision has a content, since each decision is a decision to do something (p.34).

This seems to me entirely right, and to dispel any worry about actions having content. I won't consider the matter further.)

The subject of 'content' takes us back to an issue first raised in the Introduction to this thesis, the issue of the different *kinds* of entity the word 'thought' can be used to refer to: an action, an object which represents something, or the content itself. If, as I argue in this chapter, a token thought-as-an-action (a token *thinking*) refers to the performance of a speech act, what does this tell us about the nature of a token thought when the term is being used (as a noun) to refer to something which has *content*, where 'content' refers to what the thought *means*? An obvious move is to suggest that, if a thought (in the sense of a token thinking) is the performance of a speech act, then a thought (in the sense of what the thought *means*) is what the performance of that speech act means. But this is problematic, because the meaning of a speech act is notoriously hard to identify uniquely.

Let's starting with *overt* speech acts. There are numerous candidates for the meaning of the performance of an overt speech act:

i. What the speaker intended (to achieve) by her performance (before she started speaking).
ii. What the speaker understood her own performance to mean (after she heard and interpreted her own inner speech act).
iii. What a witness to the performance understood the speaker to mean by it.
iv. What a witness to the performance understood the speaker *intended* to mean by it.
v. What all the other witnesses to the performance understood the speaker to mean by her performance, or what she intended to mean by it.

In the case of the performance of an *inner* speech act the issue of interpretation is reduced, but not eliminated. The performance might mean:

i. What the speaker intended (to achieve) by her performance (before she started speaking).
ii. What the speaker understood her own performance to mean (after she heard and interpreted her own inner speech act).

It seems that the meaning of the performance of a speech act is always a matter of interpretation. If that's right, and if we understand the meaning of a speech act to be its content, then the conventional wisdom – that a thought is individuated by its content – is wrong.

So how should we characterise the meaning of (the performance of) a speech act, if not by reference to its content? To do justice to this question would take me beyond the scope of this thesis, so my comments here are only intended to be suggestive. Future work, perhaps. My suggestion is that the meaning of a speech act is kind of event in the life of a subject who is a witness to that speech act, and who understands it (or takes it) to mean something, including cases where the witness to the speech act is the thinker/speaker herself. I suggest that the nature of this event has two aspects to it (1) the experience itself, which one might characterise as a kind of cognitive phenomenology, and (2) a change in the state of the

subject, which one might characterise as a change in her dispositions. Specifically, she will now be disposed to have further thoughts which are rationally consistent with what she understood the speech act to mean (in the full sense of what the speech act was meant to *achieve*). The suggestion can be summed up by saying that thinking involves two events: the action of performing a speech act, and the event of understanding that action to mean something. In the case of inner speech, both events feature in the life of the same person.

A final comment, before concluding this chapter. Some philosophers argue there is no metaphysically significant distinction between objects and events, and treat both as entities of the same kind. I take no stand on this one way or another, and it has no consequence for my argument. Recall, my argument is only that a token thought – one individuated by its content (in the fullest sense) – does not exist 'in the head' of the thinker before she performs a speech act – and that in virtue of that fact we should conclude that a thought is not an object of any kind, but an action – which I take to be a type of event. For philosophers who hold that events (including actions) are, metaphysically, objects, my claim can be harmlessly re-worded as follows: the object (which is the event of my performing a speech act) does not exist 'in the head' of the thinker before she performs the act. Helen Steward (2012a) argues that actions are not events, but processes. Once again, my claim can be harmlessly reworded to the effect that the process (which is the action of my performing a speech act) does not exist 'in the head' of the thinker before she performs the act.)[72]

Section 10. Summary

Some philosophers are attracted to what I have called the *thinking-as-speaking* thesis – the idea that some thoughts *just are* inner speech utterances, and vice versa. Standing in the way of this thesis is the popular idea that the content of a thought is fully determined before it is translated from a non-linguistic format into the thinker's natural language. I argued against this by appealing to speech act theory. I argued that many inner speech

[72] Kent Bach's (1980) deserves special mention here as having anticipated (and perhaps partly influenced) Hornsby (2012) and Steward (2012a). He argues that actions are not events but instances of a relation, the relation "of bringing about between agents and events" (p. 119) Since such instances are not individuals (i.e., not particulars), Bach claims they are not subject to quantification.

utterances are inner speech *acts*. If that's right, then inner speech utterances, like other speech acts, are individuated by their meaning in the fullest sense, an amalgam of propositional content and illocutionary force. I argued that the determinate content of a token thought only comes into being with the production of a natural language expression, and therefore could not be represented beforehand by a mental state. That being so, the thought being expressed by the inner speech utterance has a different content from the content of any mental state which might have preceded the performance of the utterance. Since a thought is said to be individuated by its content it follows that the thought expressed by a speech act is not the same as whatever content/information existed (in a non-linguistic format) before the performance of the speech act. I suggested that, based on this analysis, we should reject the assumption that a thought is any kind of object, such as a mental state or an expression in a natural language, and we should recognise that a thought is a kind of action – the action of performing a speech act. I buttressed this proposal by arguing that thinking is a type of activity, with the same ontological status as other (bodily) activities, such that individual actions should be conceived of as the product, upshot or outcome of a type of activity. My proposal casts doubt on the assumption which informs much thinking about thinking – that a token thought is individuated by its content. That's because if a token thought is understood as a token thinking, and if a token thinking is understood as the performance of an inner speech act, and if the *content* of a token thought is understood as what the performance of the relevant speech act means, then, because the meaning of speech acts is indeterminate, so too is the content of the thought.

Section 11. Appendix to chapter 6: "Thinking for Speaking"

In this chapter I have coined the term "thinking-as-speaking" to refer to the proposal that the language production process makes a constitutive contribution to the meaning of the utterance being produced. But the term is easily confused with another which has some currency in the inner speech literature – *thinking for speaking* – and I owe an explanation of the difference, and how the latter idea fits into my proposal, if at all. (For a detailed but also sceptical analysis of *thinking for speaking* see Vicente (2022). I will rely on that paper extensively in what follows.)

Some authors working in linguistics assume, either implicitly or explicitly, that before the process of language production starts, not one but *two* systems of conceptual representation are involved. The first system is thought to use conceptual representations which are universal, regardless of what natural language the subject speaks. The second system uses conceptual representations which are aligned to the semantic representations of the subject's particular natural language, but which are nevertheless distinct from those representations. Levinson (2003) called the first type of representations "CR" (for "conceptual representations"), and the representations of the second type, "SR" (for "semantic representations"). I will do the same. Slobin (1987) called using the second system of conceptual representation "thinking for speaking". The reason for supposing that this second system exists is based on the observation that each natural language carves up the different "realms of reality"[73] in slightly different ways. These "realms" include, among others: the domain of spatial relationships (e.g., of inclusion and support), the domain of colour, the domain of motion events. The details are somewhat technical, and I don't have space to do more than sketch the idea here, but an example from Vicente (2022) will help convey the idea.

> In general, an English speaker will describe a scene in which an individual runs from point *a* to point *b* by making reference to the manner in which they move (*run*), while for Spanish speakers, referring to the manner of motion is expensive and anomalous. Their description will skip manner of motion but be more accurate about the trajectory (*S* went from *a* to *b*). The result is that English speakers might apply the same description to two motion events that Spanish speakers might describe in different ways, and vice versa (p. 4).[74]

But we don't want to say that the English speaker and the Spanish speaker in this example didn't have the *same thought* – the thought, for example, that Jim ran from his home to the bus stop. So, the idea is that the thought they both have might be represented *first* by the same (universal) CRs, and *then* go through a "translation" process into SRs, a process which generates a "message", which is then input to the language production process. (Note that

[73] I have borrowed the phrase from Vicente, (2022).
[74] Page numbers refer to the final version of this paper available on *PhilPapers*.

in Levelt *et al's* (1999) hierarchy of speech production, reproduced in Chapter 2, Section 2, his assumption is that the first stage of the process, the one he labels "conceptualisation", is the creation of a message in SR.) In short, "thinking for speaking" means preparing a message that meets the demands imposed by the subject's native language so that it can more easily be put into words.

Suppose that Slobin's "thinking for speaking" thesis is right. In that case, I suggest, it is very friendly to my view, since the more stages involved in producing a thought – the more kinds of transformation that must be performed on the pre-linguistic information in order for it to become a determinate thought – the more credible the claim that the thought is not determinate before it is articulated via the medium of a speech act. But as it happens, Vicente (2022) argues, very persuasively in my view, that thinking for speaking is a redundant process, and that much of the work being ascribed to the thinking for speaking thesis by its supporters is done by the (natural) language production process, which "plays a major role in shaping thought". Vicente is explicit that "our inner speech utterances have contents that we have not previously thought" (p. 15), so his rejection of the thinking for speaking thesis is also friendly to my view. In short, whether the thinking for speaking thesis is right or not, both possibilities are friendly to my view.

Chapter 7. An objection to *thinking-as-speaking*, and a response

Section 1. Introduction and Objection

In the previous chapter (Section 6) I introduced the Rylean (1949) idea that, in addition to propositional knowledge, there exists a kind of non-propositional knowledge – practical knowledge or knowledge-how. I invoked that idea in order to explore the epistemic question of what sort of knowledge we have of our own thoughts. In this chapter I want to invoke Ryle's idea again, but for a different purpose – to set the scene for a possible objection to the thinking-as-speaking thesis as I have characterised it. Ryle's idea was radical at the time, an attack on the prevailing orthodoxy, which maintained that *all* knowledge was propositional knowledge, or knowledge-that: a subject knows how to do something if and when they are standing in an appropriate relation to a proposition. Ryle referred to this view disparagingly as 'Intellectualism', and his own view came to be known as 'Anti-intellectualism.' Ryle's view was the dominant one for the next 50 years, but the consensus was broken with the publication of a paper by Stanley and Williamson (2001). They argued that, *pace* anti-intellectualism, knowing-how is a species of knowing-that. An example illustrates the basic form of their argument:

> [Y]ou know how to ride a bicycle if and only if you know in what way you could ride a bicycle. But you know in what way you could ride a bicycle if and only if you possess some propositional knowledge, viz. knowing, of a certain way w which is a way in which you could ride a bicycle, that w is a way in which you could ride a bicycle (Stanley, 2011, p. 209).

Stanley and Williamson's paper spawned a substantial literature.

As Pavese (2022) observes in her *Stanford Encyclopaedia of Philosophy* entry on knowledge-how, "The most recent debate on knowledge-how has intertwined with a debate on the nature of skills". In the course of that "recent debate" a position emerged (one might think of it as a sort of compromise) which Fridland (2017) and Mylopoulos (2020) have characterised as the *hybrid view*. The hybrid view, as the name suggests, allows that skilled action might involve both kinds of knowledge. In her (2020) Fridland offers a careful

analysis of what each party gets right, and what each gets wrong, about the phenomenon of skilled action. Here is her summary:

> Overall then, my take on the philosophical debate between possessing skills and the ability to perform them is as follows: if anti-intellectualists think knowing how is simply the instantiation of know how in action, they are wrong. Skills critically involve internal, learned, control states that account for the ability to implement skilled actions at different times. But anti-intellectualists are right to think that skills and instantiations of skilled actions are necessarily and intimately connected. That is, anti-intellectualists are right that action is at the heart of skill. Intellectualists, on the other hand, are right that skill is not identical to the ability to instantiate skill on any particular occasion or group of occasions. But they are wrong if they conclude from this that skill and instantiations of skill in action are only contingently related. They are also wrong to conceive of the standing states of a skilled agent in terms of propositional knowledge, but they are right to insist that the internal standing states are central to the possession of skill (p. 245).

In this introductory section I will show how the hybrid view, or at least one version of it, can be used to mount a challenge to the thinking-as-speaking thesis. On this version of the hybrid view, skills (like playing the piano or riding a bicycle) have two components: a cognitive component, and a motor component:

> "The cognitive component is usually cashed out in terms of propositional knowledge or intentional states and the motor component is construed in terms of automatic, low-level, causal processes, which are acquired through brute repetition" (Fridland, 2017, p. 1540).

To give it a name, and to distinguish it from other versions of the hybrid view, call this the *strict hybrid view*, to reflect the fact that it makes a strict demarcation between the parts of the skill which are *intelligent* and the parts which are not. On the strict hybrid view, the propositional knowledge associated with the cognitive component of a motor skill is solely responsible for the *intelligence* involved in executing the skill. To put it another way: if a subject possesses a skill, all the intelligence associated with the skill is due to her propositional knowledge, (the cognitive component) and none of the intelligence is due to her skills of execution (the motor component). For now, 'intelligence' should be understood

as meaning "the ability to systematically respond to various conceptual, semantic, intentional and otherwise, higher-order cognitive, personal-level contents of goal-directed representations" (Fridland, 2017, p. 1547) during the performance of a skilled action. (This idea will become clearer in what follows.) The motor component is essential for the *execution* of a skilled activity, but this component merely consists in 'motor acuity': the ability a subject possesses to make the detailed fine-grained movements involved in executing a skilled action. Motor acuity is construed as "a low-level, non-knowledge-involving, bottom-up, brute causal process…responsible for the detailed kinematics of motor skill instantiations" (Fridland, 2017, p. 1547).

So, finally, we can articulate a possible objection the thinking-as-speaking thesis. If the strict hybrid view is right, an objection might go as follows:

> Objection. Suppose I grant your point about the motor planning process playing an essential role in the performance of an inner speech act. I can grant that and still deny that the motor skill involved in *speaking* makes the kind of contribution to *thinking* that you claim it makes. You claim that the action of speaking (i.e., the process of forming a natural language expression) contributes not only to the *performance* of a speech act, but to its very *meaning* (in the fullest sense). But if we apply the strict hybrid view of motor skills to the case of thinking-as-speaking we might reasonably suppose that all the *intelligent* work involved in thinking-as-speaking – all the work that involves making the speech act *meaningful* – is done by the cognitive component, and the speaking is purely motoric. That is, the motoric aspect of speaking is an automatic, low-level, causal process which merely executes instructions which have been generated by the (intelligent) cognitive component. If that's right, then your claim that speaking makes a constitutive contribution to the *meaning* of a speech act, and consequently to the thought, is wrong. In a slogan: thinking is smart, speaking is (so to speak) dumb.

My response to this objection has two parts. The first part is to attack the strict hybrid view (Section 2). If the strict hybrid view is wrong as far as skilled bodily actions are concerned, then its application to the thinking-as-speaking thesis is otiose, and the objection neutralised. But I can do better than that. The second part of my response to the objection is that even if the *articulation* stage of the language production process is conceived of as

reflex-like (as the strict hybrid view maintains), that fact does not undermine the thinking-as-speaking thesis – at least, not if all the *other* stages of the language production process are granted to be intelligent. That's because by the time the articulation stage of the language production process is reached, the full and final content of the speech act (its *meaning* in the fullest sense) has been determined, and only requires vocalisation to be executed. That being so, the language production process has already made its constitutive contribution to the full meaning of the speech act, just as the thinking-as-speaking thesis maintains. As a final attempt to make the objection stick one might try to argue that the *entire* process of language production is reflex-like, but this option has some extremely implausible implications and should be rejected (Section 3). Section 4 concludes.

Before attacking the hybrid view in the next section I want to briefly raise, and reject, another possible objection. Some readers will we aware that many linguists, including such luminaries as Chomsky (2000) and Pinker (1994), have long held that the human capacity for language is not a skill, but an instinct. If that's right, then it might be thought that both the objection above, and my response to it, are beside the point; skill has nothing to do with speaking, so it can have nothing to do with *thinking*-as-speaking. But this is too quick. There are good reasons to believe that the traditional sharp division between the so-called Exceptionalists (who emphasise the innate nature of the language capability) and the so-called Anti-Exceptionalists (who reject it entirely) is too simplistic. Armstrong & Pavese (Manuscript), for example, argue that a middle way is much more plausible than either of these two extreme positions. They argue that human linguistic competence is a skill, but one that is guided by a "biological endowment", one that grounds a number of other social and environmentally dependent skills.[75] Since I am entirely convinced by their argument (and the evidence they adduce in its favour) I will not consider this possible objection any further.

Section 2. The 'intelligent view' of skilled action

[75] Balcarras' (2023) position is similarly conciliatory between the two positions.

In this section I respond to the objection by attacking the strict hybrid view as characterised in Section 1. In the realm of skilled *bodily* action, the strict hybrid view has been challenged by a number of writers and generated a substantial literature. A common theme is to argue for the intelligence of 'motor representations'.[76] Here are some examples of writers who take this approach. Ellen Fridland (2017) and Neil Levy (2017) argue, roughly, that the skill that's involved in performing a skilled action (and therefore involved in demonstrating knowledge-how) does not reside purely in the cognitive component of the skill, but also resides in the motor component – specifically the motor representations which execute the action. Shepherd (2019) says something similar: "Agents have the capacity to specify motoric parameters for action execution at the personal level, and when they do so their intentions can lead a double life, taking both propositionally and motorically formatted contents" (pp. 301-2). Brozzo (2017) argues that some motor representations are matched with corresponding motor *intentions*, and as such are open to rational appraisal: "Motor intentions…provide an extension of the reach of an agent's *reasons* into very detailed descriptions of her bodily action – descriptions that encompass specifications of her bodily movements" (p. 253, emphasis added). Sinigaglia & Butterfill (2015) present empirical evidence that "…some motor representations carry information about outcomes of the kind sometimes identified in thought as the goals of actions…" (p. 1927). Christensen, Sutton, and McIlwain (2016) argue that "controlled and automatic processes are closely integrated in skilled action, and that cognitive control directly influences motor execution in many cases" (p. 43). Ferreti & Caiani (2021) argue that "MRs [motor representations] are intelligent representations, as they are sensitive to propositional information and conceptual categorization" (p. 22). Mylopoulos (2020) even argues that " the motor system is intelligent in its own right" (p. 262). I could go on; for a list of additional papers in the same vein see Fridland (2019, p. 780, FN 21.)

The idea that is common to these writers is that knowledge-how (skilled action) is constituted by the possession of both propositional knowledge *and* motor representations, and that the motor representations do more than merely execute an action, the skill of

[76] The idea of a motor representation comes from Pacherie (2007, 2008). What are they? "Motor representations encode action goals together with the motoric means for achieving them and do so in a motoric format directly suitable to action execution" (Mylopoulos & Pacherie, 2018, p. 3).

which is explained purely by reference to propositional knowledge. Rather, the intelligence of the motor representations makes a constitutive contribution to the skill (or intelligence) of the action. I will follow Fridland (2017) and refer to this view as *the intelligent view*. According to Fridland, we should reject the strict hybrid view of motor skills and understand motor skills as "intelligent all the way down" (p.1539).

The rest of this section is devoted to an argument from Fridland (2017) that the strict hybrid view does not fit the empirical data associated with paradigm skilled bodily actions, and that the intelligent view does. Fridland's (2017) paper is targeted at two recent accounts of skilled action, from Stanley and Krakauer (2013) and Papineau (2013), both of which posit that motor control (i.e., the motor component of paradigm bodily actions) is essentially reflex-like. She argues that this position implies predictions which are not borne out by the data and that, to the extent this is true, their position is disconfirmed.[77] Specifically, she claims that if their accounts are correct then we should find that motor control has the following characteristics: it is (1) ballistic, (2) invariant, (3) independent of general action trajectories, (4) insensitive to semantic content, and (5) independent of personal-level intentions. And in fact, she claims, we find that it is not.[78]

In what follows, I offer a summary of Fridland's arguments under the heading of each of the 5 characteristics listed above. If you accept her arguments, you will conclude that the strict hybrid view is flawed. Under the same headings I propose that thinking-as-speaking also fails to display the 5 characteristics which you would expect it to display if the motor component of thinking-as-speaking was reflex-like. Of course, my proposals, unlike Fridland's arguments, are not backed up by empirical evidence, so they lack the same force. Even so, if you find Fridland's arguments persuasive, then you might wonder why what she says about the motoric component of skilled action shouldn't also apply to the motoric component of thinking-as-speaking. And if you find *my* suggestions persuasive, you might

[77] Additional arguments to the same end, with further empirical evidence, can be found in Fridland (2019).
[78] Empirical arguments are not the only way to go. An a priori argument is provided by Dickie (2012), the power of which was recognized even by Stanley (a leading intellectualist): "On certain conceptions of propositional knowledge, Dickie's objection is powerful, and indeed devastating" (Stanley 2012: p. 763). Furthermore, Ferretti (2020) shows how empirical evidence from cognitive science on motor representation can be marshalled to support Dickie's (a priori) argument.

conclude that the motor control involved in thinking-as-speaking, like any other skilled action, should also be understood as intelligent "all the way down". If that's right, then the objection to the thinking-as-speaking thesis, which was that the speaking component does not make an 'intelligent' contribution to the action but is a brute causal process that simply executes the motor planning involved in speaking, is refuted.

1. Is motor control ballistic?

According to Fridland, if motor controls were ballistic, they would automatically run in their entirety once initiated, with the subject unable to intervene or interfere as the action unfolds. But intuitively it seems obvious that the more skilled we are at something, the more ability we have to control the way an action unfolds. We would not *expect* the motor routines that are involved in complex skills to be ballistic. And this is what we find empirically. Fridland cites studies which conclude that the motor processes that underlie skilled action are both automatic (in the sense of not requiring conscious attention) *and* controlled (in the sense that subjects can nevertheless make adjustments to what they are doing, as they are doing it, when prompted).

I suggest that when we consider the way that language is produced when we engage in thinking-as-speaking we observe something similar. On the one hand there is clearly a sense in which our speech just 'unfolds' automatically. (Indeed, as I have already observed (see Chapter 5, Section 2, where I cite Strawson, 2003) there is something incoherent in the idea that we deliberately *choose* the content of our thoughts.) And yet at the same time, we frequently make adjustments to it as we go along. It is not as though the process of speaking, once started, cannot be interrupted, as would be the case if the process were "ballistic". On the contrary, as explained in Chapter 2, Section 3, we have sophisticated monitoring systems for allowing us to make corrections even before the vocalisation of a thought commences.

2. Is motor control invariant?

The idea here is that "if automatic motor routines comprising skilled actions were reflex-like, then we should expect them to be more or less invariant or fixed" (p. 1551). But the evidence from studying the performance of athletes (for example) is that movements are

inherently variable. Athletes exhibit two apparently conflicting properties of the motor system: the ability to perform high level goals reliably and repeatedly, and yet to do so by varying the fine motor movements needed to accomplish the goal. (See Todorov and Jordan, 2002, p. 1226).

Analogously, I suggest, when a subject makes a particular utterance, there are numerous ways that utterance type can be instantiated i.e., voiced. I can utter a token of the type "Pass the salt, please" in a great variety of ways, depending on the precise circumstances. The differences in each case are all due to differences in the way I articulate the request, i.e., the fine motor movements I make when I speak the words. It seems the motor controls involved in language production are not invariant.

3. Are trajectory and motor control independent?

One feature implied by the strict hybrid view that Fridland rejects is that the detailed kinematic strategies executed in motor skills blindly implement a general, pre-planned trajectory. But studies show that this is not what happens. In studies of subjects reaching for an object, for example, where perturbations are introduced after the action has begun, certain perturbations are corrected for, others not. Those that are not corrected for are those that are irrelevant for task success, suggesting that "fine-grained sensorimotor control is flexible insofar as corrections are made in an intelligent way – not simply to conform to a pre-determined trajectory, but in order to achieve one's goal" (p. 1554). Fridland concludes that action trajectory planning and execution are not independent.

Consider what happens when a subject is thinking aloud, and someone interrupts her. (This is analogous, I suggest, to the perturbation which interferes with a skilled bodily action once it has started.) Will the subject carry on regardless, saying what she was going to say, or will she make an adjustment? I suggest that, as with bodily actions, this will depend on the relevance of the interjection. If it is irrelevant to the subject's goal, i.e., the reason for her thinking aloud in the first place, then she will ignore it and say what she was going to say anyway. If, on the other hand, the interruption is relevant she will make adjustments to what she says and take the interjection into consideration. (The nature of thinking out loud is considered in more detail in Chapter 9.)

4. Is motor control sensitive to semantic content?

The previous three points combine to generate a question. If motor controls are not ballistic, not invariant, and not independent of trajectory, how is it that they are nevertheless relevant to the high-level goal of the action? The only explanation, according to Fridland, is that the motor control system "is directly sensitive to the semantic content of personal-level goals" (p. 1556).

This point seems highly intuitive in the case of language production. It would be hard to imagine that what we say, as we say it, was *not* sensitive to the "semantic content of our personal-level goals". In the case of speaking, the "semantic content of personal level goals" is best understood, I suggest, as content of the information which is transformed by the language production process into a speech act. In particular, the content of the communicative intention which causes, sustains and guides the production and performance of the utterance.

5. Is motor control independent of intentional states?

According to the strict hybrid view criticised by Fridland, our intentional states are only causally related to the motor control systems which execute them. She concludes that such a view is unsustainable. "[T]he best evidence we have indicates that fine-grained, automatic motor processes instantiated in motor skills are not simply causally connected to intentional states but, rather, continue to be semantically sensitive and responsive to personal-level goals throughout execution" (p. 1557). I see no reason to believe that this is any less true of the motor skills involved in language production than it is for the motor skills involved in skilled bodily actions. On this new view, there is a kind of "intentional cascade",[79] from the formation of a communicative intention right through the language production process, at each stage of which the relevant representations possess enough information about the original goal of the utterance to perform their function intelligently.

[79] I have taken this term from Pacherie (2008)

(Ironically, some of the literature on motor planning, and on so-called 'motor knowledge', has adopted the metaphor of *language* to describe some features of motor control. For example:

> [The] learning mechanism within the motor system, known as *motor reinforcement*, structures F5 as a sort of *motor vocabulary*, in which words are composed by groups of neurons encoding very specific parameters of action and a specific motor act is represented as an ensemble of different *motor words* (rather than a simple movement) (Rizzolatti and Sinigaglia 2008). These *words* can refer to different motor aspects: the goal, the execution, etc. (Rizzolatti and Sinigaglia 2008; Gallese and Metzinger 2003: p. 367; Ferretti 2016: 4.2; 2017: p. 7.1). As Jeannerod puts it, in the *motor vocabulary* "actions are encoded element by element" (2006: p. 12) (Ferretti, 2020, p.11, original emphasis).)

In summary, if the intelligent view is as true of thinking-as-speaking as it is of other skilled actions, and if, therefore, the strict hybrid view is wrong, then the objection – that the motor planning component of thinking-as-speaking makes no intelligent contribution to the performance of the speech act – is refuted.

Section 3. The objection revisited

Maybe you are still not convinced. You might feel that there is something so special about the activity of thinking that it is just absurd to suppose that the physical process of producing the sounds required to vocalise the words can make any intelligent contribution to the meaning of the thought which the speech act performs. But if you want to insist on the strict hybrid view when it comes to thinking-as-speaking, you need to have an answer to the following question: At what point in the thinking-as-speaking event does the process change from being intelligent to being dumb? To illustrate the problem, here is Levelt *et al's* model of the speech production process again, originally reproduced in Chapter 2.

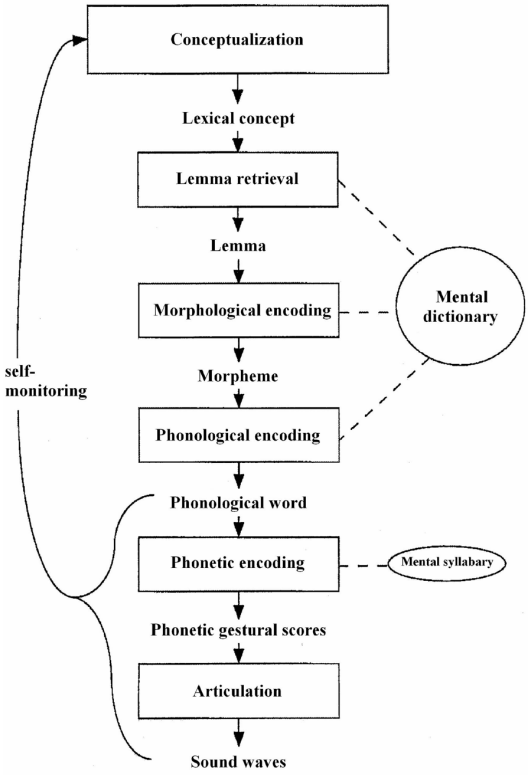

Figure 1. The theory of speech production in outline. From Levelt *et al* (1999).

Suppose, to restate the objection, you reject the idea that the last stage of the process – Articulation – makes any intelligent contribution to the meaning of the thought which the speech act expresses. In that case, the question can be restated: Which, if any, of the previous stages of speech production make a contribution to the meaning of the thought which the speech act expresses?

Let's explore two possible answers to this question, the two extreme cases: all and none. In the first extreme case the answer is that *all* the previous stages of the language production process make a constitutive contribution to the meaning of the thought, on the grounds (a) that these stages are the ones that generate semantic or linguistic meaning, and (b) because it is patently absurd to suggest that semantic meaning doesn't contribute to the meaning of the speech act. But in that case, it doesn't matter whether the articulation stage makes any intelligent contribution or not, because by the time that stage is reached, the semantic meaning is in place, and the thinking-as-speaking thesis goes through. And if that thesis goes through, then so does that idea that the thought being expressed by the speech act does not exist (in its full, final and determinate form) before the event which is the

performance of the speech act. In short, this option doesn't undermine my claim for thinking-as-speaking in any way, since it accepts that the language production process makes a constitutive contribution to the meaning of the speech act. (This point is the same as that made in Chapter 6, Section 7, where I explained how the thinking-as-speaking thesis is compatible with the abstract (rather than the motor-sensory) view of inner speech.)

The second extreme answer to the question posed (Which stages, if any, of the language production process, contribute to the meaning of the thought which the speech act expresses?) is to argue that *none* of them do. On this view, all the stages of the language production process, from the conceptualisation stage through to articulation, share the same characteristics as the strict hybrid view attributes to motor acuity – that is, they are automatic, low-level, causal processes, acquired through brute repetition. On this account, all the intelligence involved in thinking has been done before the language production process gets under way. All the *meaning-making*, as it were, is done before that process starts. But this is very implausible. How could the semantic meaning of a natural language utterance have existed before the language production process even begins? To believe this you would have to believe that all the semantic meaning generated by the language production process is nothing but a *literal translation* from some pre-linguistic language such as a language of thought, into natural language. You would have to believe that all the elements of meaning that we associate with a natural language expression are already present, albeit differently represented, in the pre-linguistic language. It is not at all clear how one could even make sense of this idea. In any case, as I argued in the previous chapter, it is overwhelmingly likely that all the pre-linguistic content (or *information* as I prefer to call it) which exists in the head of the subject before she speaks (and indeed causes, sustains and guides the language production process), goes through a process of *transformation* during the language production process. Indeed, if it didn't go through a process of transformation, one might wonder why we have two 'languages' (in the sense of two *formats*) at all, a pre-linguistic format and a linguistic one. In short, this answer to the question should be rejected as hopelessly implausible.

There might be other less extreme answers one could give to the question, other than the 'all' or 'none' answers. One might want to pick a point somewhere else in the language

production process – somewhere between conceptualisation and articulation – and argue that before that stage the process is intelligent, and after that it is all brute reflex. I doubt such a case can be made plausible. It seems likely that there is so much monitoring, feedback and subsequent adjustment at all stages of the process that such a strategy is doomed.

In summary, if the strict hybrid theorist digs in her heals and insists that the articulation stage of speech production is reflex-like, but agrees that all the proceeding stages are intelligent, that does not remotely undermine the thinking-as-speaking thesis, because on that view the meaning of the speech act has been specified before the articulation stage begins. But if they insist that the entire process of language production is reflex-like, they seem to be claiming that the entire language production process contributes nothing to the meaning of what is said. And that is very implausible. A more plausible alternative to either of these positions, in my view, is to accept that there is a "intentional cascade" (Pacherie, 2008) which runs through the whole process, top to bottom, and that the process is, as Fridland puts it, "intelligent all the way down".

Fridland (2020) argues that we should understand skills as functions, "Namely, as functions from intentions to controlled, successful actions, where the functional transformations from intention to action are implemented by control structures that have been developed through practice" (p. 247).[80] My suggestion is that we should understand thinking as a skill which can be characterised in the same way, "as functions from intentions to controlled, successful actions", where the relevant type of "successful action" is a (usually silent) speech act.

Section 4. Summary

According to what I have called the *strict hybrid* view, skills have two components: a cognitive component, and a motor component. The cognitive component involves

[80] Fridland identifies three kinds of "control structures" involved in skills: strategic control, attentional control and motor control. Her characterisation of these three kinds of control, and how they might apply to the skill of thinking-as-speaking, is not something I have space for here.

propositional knowledge, and the motor component involves automatic, low-level causal processes, acquired through brute repetition. All the intelligence associated with the skill is associated with the cognitive component; the motor component is dumb – little more than reflex. The hybrid view can be used to pose an objection to the thinking-as-speaking thesis, by arguing that thinking is the cognitive component of this activity, speaking is the motor component, and all the intelligence involved is supplied by the cognitive component. On this account, contrary to the claims of Chapter 6, the speaking component would make no contribution to the content of what is said by an utterance; the speaking component would merely execute the production of a meaningful expression already fully specified in a non-linguistic format.

I offered two arguments against this objection. The first is that the hybrid view is wrong in the case of skilled bodily actions. Skilled bodily actions show none of the characteristics one would expect them to show if the strict hybrid view was right. The second is to pose the objector a dilemma. If the objection is confined to the last stage of the language production process, then it misses its mark, since by then the language production process has produced a meaningful utterance and made its constitutive contribution to the meaning of that utterance. If, on the other hand, the objection extends to the whole of the language production process, the implication is that all the meaning of a natural language utterance is entirely determined before the language production process has even begun. For various reasons, this is tremendously implausible.

Part lll: The metaphysics of conscious thought

This part of the thesis consists of two chapters which explore the metaphysics of conscious thought. The first chapter takes the proposal made in Chapter 6, that thoughts are a type of action, and subjects that idea to Kit Fine's Theory of Embodiment, which concerns bodily actions. I argue that his theory can also accommodate the thinking-as-speaking thesis. The second chapter is different to all the others, since it is dedicated entirely to an alternative ontology of conscious thought, one proposed by Matt Soteriou (2013). Soteriou's account is chosen for two reasons. First, it is one of the most sophisticated and comprehensive of recent years, and as a rival to my own I feel compelled to respond to it. Second, his account takes very seriously the role speech plays in conscious thought, and yet comes to very different conclusions. One goal of this chapter is to understand why.

Chapter 8. Thoughts as embodied acts

Abstract

Thoughts, like inner speech utterances, can be divided into those which are active and those which are reactive. Call instances of the first a 'thought$_a$' and instances of the second a 'thought$_r$'. The former kind qualify as actions in virtue of its being caused, sustained, and guided by an intention. The latter kind are merely events; they are not the product of an intention, but a reaction to circumstances. The medium for both kinds of thought is inner (and sometimes outer) speech. Metaphysically, then, thoughts$_a$ are a kind of action, not a kind of object. They are things we *do* (with words), not things we *have* or *instantiate* (and then express with words). Their content is something we make, as we think$_a$, not something we access and then translate into natural language. How well does this proposal cohere with the metaphysics of human action? I take one theory – Kit Fine's Theory of Embodiment – and I argue that it fits very well. I consider some implications of my application of Fine's theory to my own for the nature of intention.

Section 1. Introduction

I will begin this chapter with a recap. In Chapter 4 I argued that inner speech utterances, when they are intentional ('active' and not 'reactive') are the basic actions by which a non-basic action is performed, where the non-basic action is something cognitive, such as trying to decide something, or reach a conclusion, or solve a puzzle, etc. In Chapter 6 I suggested that, from a metaphysical point of view, thoughts, when they are caused by an intention to get something done, are not objects, such as mental representations or sentences in a language of thought, but the action of performing a speech act. In this chapter these two ideas come together. Thoughts, like inner speech utterances, can be divided into those which are active and those which are reactive. Call the first a '$thought_a$' and the second a '$thought_r$'. The former qualifies as an action in virtue of its being caused, sustained, and guided by an intention. The latter is a mere event; it is not the product of an intention, but a reaction to circumstances. The medium for both kind of thought is inner (and sometimes outer) speech.

A token $thought_a$ is individuated by its content, and the content of a $thought_a$ is an amalgam of propositional content and the illocutionary force with which the words are uttered. (A token $thought_r$ is also individuated by its content, but in the case of a $thought_r$ there is only linguistic meaning in a context; there is no illocutionary force.) In the case of both $thoughts_r$ and $thoughts_a$, the production of the utterance – albeit in response to different causes – is automatic. This is analogous to the way in which, in the case of a bodily movement, the execution of the movement is automatic, whether the movement is intentional (I raise my hand to attract the waiter) or not (my hand rises – as a reflex response to an object approaching my face, say). A $thought_a$ is the basic action *by which* a non-basic (cognitive) action is performed. We make $thoughts_a$ to get something (cognitive) done – make a decision, come to a conclusion, work something out, reflect on our options, etc. A $thought_a$ is the performance of a type of speech act: an *inner* speech act.

Metaphysically then, $thoughts_a$ are a kind of action, not a kind of object. They are things we *do* (with words), not things we *have* or *instantiate* (and then express with words). The content which individuates them is something we make, as we $think_a$, not something we access and then translate into natural language.

On this account, the nature of a thought$_a$ is metaphysically comparable to Crane's (2017) account of the nature of a belief – see Section 7 of Chapter 6. To recap: a belief is not the reproduction of an object but an interpretation, made at a point in time, of the subject's world view. When a subject expresses a belief, she models some aspect of her world view, one that is relevant at the time it is expressed. The model takes the form of a proposition in the medium of (inner or outer) speech. The words used to express the proposition are 'chosen' automatically. The choice of words is guided by the purpose for which the subject's world view is being modelled. 'Expressing' a belief (interpreting some relevant aspect of your world view) is comparable to 'having' a thought (using words to get something done); both are, in their own ways, acts of *creation*, rather than acts of *reproduction* or *translation*.

As a metaphysical thesis about the nature of thinking this will strike many as surprising, to say the least. So far in this thesis my approach has been to build up to it, starting with some important scientific evidence about the (inner and outer) speech production processes (Chapters 2 and 3), and working through some of the conceptual issues arising from the phenomenon: the active and reactive nature of inner speech utterances (Chapter 4), the role of inner speech in reasoning (Chapter 5), the speech-act function of some inner speech utterances (Chapter 6). In this chapter I want to buttress the metaphysical claim from another direction. I want to start with a respectable metaphysical theory concerning the nature of action and evaluate how well my thesis fits it. I take it that my account will gain or lose credibility according to how well or badly it fits. The theory I have in mind is Kit Fine's Theory of Embodiment (1982, 2003, 2008, 2022). Of course, some readers will be sceptical about Fine's theory, in which case they will not be impressed by how well or badly my thesis fits it. But this cuts both ways. If I am right about thinking-as-speaking, and if my thesis fits well with Fine's theory, then perhaps his theory gains credibility from having broader applicability than Fine himself imagined.

The rest of this chapter goes as follows. In Section 2 I provide a summary of Fine's theory of embodiment, as it applies to bodily action. In Section 3 I show how these ideas can be extended to cover mental action – specifically, to the idea that thoughts, or more precisely thoughts$_a$, can be understood as a species of action. In my thesis, as in Fine's, *intention*

plays a key role; in Section 4 I explore some implications of my thesis for the nature of intention itself. Section 5 concludes.

Section 2. Fine's Theory of Embodiment

Fine's theory of embodiment as it applies to human action makes its first, and brief, appearance in his (1982), a paper which is mainly focused on addressing a longstanding metaphysical puzzle concerning the identity of material objects. The puzzle is that, on the one hand, a statue (say) and the bronze of which it is made, would seem to be identical, a single object. On the other hand, the bronze and the statue don't share the same properties – the bronze would survive the destruction of the statue, for example – and, as nearly everyone agrees, objects with different properties are not identical. Fine rehearses and rejects as inadequate what he calls *extensional* attempts to explain the problematic relationship-of-identity between a material object and the matter from which it is made. He argues that these problems (which I won't rehearse here) can best be solved by conceiving of the statue as what he calls a *qua object*. In the case of a statue made of bronze, the bronze matter is a constituent of the statue (its 'base'), and the statue has the additional property – not possessed by the base – of having been intentionally formed into a statue (its 'gloss'). So, on his account, there are two distinct entities before us – a statue and a lump of bronze. However, this does not generate the traditional puzzles of identity because the causal properties of the qua object, the statue – its shape, weight, hardness, etc. – overlap entirely with the causal properties of the object which is its base – the lump of bronze. An important idea here is that there is an act of creation by the sculptor, who takes something of a particular kind (the bronze) and intentionally makes it into something which qualifies as a statue (according to whatever our best theory of what a statue is).

Almost as an aside, Fine draws a parallel between the identity problems which have traditionally dogged the ontology of material objects, and similar problems in the domain of action theory: What is the relation between a bodily action (such as my raising my arm) and the bodily movements which constitute that action (my arm rising)? Fine is adamant about the ontological parity between the two cases: "It is not just that the two relationships are

analogous: they are exactly the same" (p. 102). Once again, I will not rehearse the puzzles of identity which the two events are said to generate, nor the different positions taken in the debate (which Fine characterises as a conflict between *conceptualists* and *causalists*). Fine's proposal for solving these puzzles is that we should conceive of the action of my raising my arm as a qua object. On Fine's account, the movements involved in my arm rising are the 'base' of the action, and its 'gloss' – the property which the action of raising my arm has – is that the movements are made intentionally. So, once again, there are two distinct entities involved, but in this case the entities are events, rather than material objects. The existence of two events is not metaphysically problematic because the causal properties of the action overlap entirely with the causal properties of its base – the movements of my arm.

In his (1982) Fine spends scarcely more than two pages of his paper sketching his proposal that actions are qua objects, before returning to his main topic. There is then a gap of forty years before he addresses the idea again in any depth, with the publication of his *Acts and Embodiments* (2022). (This is not to say he does not refer to the idea again, in passing, in other work. One example, which is particularly tantalizing, occurs when he is explaining why something like a lump of bronze, unlike a statue, can never be described as 'misshapen', even if it happens to be in the shape of an imperfect sphere. He writes, "Trifling as this example may appear to be, I suspect that there are many significant cases – *such as how it is that a person is able to think while his body is not* – that can be explained in a similar way (2008, p. 115, emphasis added). Sadly, he doesn't elaborate.)

Fine begins his (2022) paper by drawing a distinction between two questions which come up in the theory of action: 'what is it to act?' and 'what is an act'. (I say more on his use of the term 'act' below.)

> The first of our questions is conceptual; it is concerned with the concept of acting, of what it is to act. The second question is ontological; it is concerned, not with the concept of acting but with the objects that are the acts (p. 15).

Fine is almost exclusively concerned with the second, ontological question: "What, for example, is the (token) act of raising my arm or, more generally, what is the (token) act of A's ϕ-ing, for any doer A and doing ϕ?" (*ibid*). My concern in this chapter is also ontological; to paraphrase Fine: What is the (token) act of thinking *p* or, more generally, what is the token act of A's thinking ϕ, for any thinker A and thinking ϕ? I will be arguing that I can answer my question in broadly the same terms as Fine answers his, and that what we should conclude from this is that thoughts (or at least $thoughts_a$) qualify as acts on broadly the same terms as bodily actions do.

Fine summarises his original theory of embodiment (i.e., as it applies to objects, not actions) as follows:

> The basic idea behind rigid embodiment[81] is that an object and a property or description of the object can come together to form a new object – an object which results, so to speak, from imposing the property upon the given object.[82] We might call an object obtained in this way a *qua object* and, where *b* is the given object and ϕ the property, we might designate it as 'b qua ϕ' or as 'b under the description ϕ', and notate it as 'b/ϕ'. Thus if *b* is Socrates and ϕ is *being a philosopher*, the new object will be Socrates under the description of being a philosopher…. Given a qua object *b* qua ϕ, we call *b* its *basis* and ϕ its *gloss*. It is essential to our understanding of a qua object that it is not identical – or, at least, not generally identical – to its basis. Socrates qua philosopher is not simply Socrates. It is rather some sort of amalgam of the basis and the gloss in which the gloss preserves its predicative role and somehow serves to modify or qualify the basis[83] (p. 18).

Before showing how his theory of embodiment applies to acts, he makes some important remarks about pluralism,[84] which also explain his shift in vocabulary since his (1982) paper, from talking about an 'action' to talking about an 'act'. Let me explain, using his own

[81] 'Rigid' embodiment deals with objects whose constitution or matter is rigid; 'variable' embodiment deals with objects whose constitution or matter can vary. In his (2022) Fine is only concerned with rigid embodiment. The distinction need not concern us here.
[82] Fine has an important footnote here: "I do not want much to turn on my use of the term 'property' here. I might equally well have talked of features or forms." When I describe $thoughts_a$ as utterances which possess the property of being intentional, I mean 'property' in a similarly loose sense.
[83] Apart from the change in terminology from 'base' to 'basis', this account is broadly the same as his (1982).
[84] Fine describes himself as a "committed pluralist", and refers us to Fine (1991, 1999, 2000, 2003, 2007).

example. Suppose Smith shoots a gun and thereby kills Jones. There is then Smith's act of shooting the gun and Smith's act of killing Jones. The question is, are these acts identical or not? A monist (e.g., Davidson, 1971) will maintain that there is a single act here, intentional under the description *shooting* (since, by hypothesis, Smith intended to shoot the gun) but not intentional under the description *killing* (since, by hypothesis, Smith did not intend to kill Jones) A pluralist might be tempted, according to Fine, to agree with the monist that, at least in the case of the predicate 'intentional', the term is intensional (with an 's'). But this, says Fine is a mistake. Suppose, for example, we are talking about 'Smith's act of killing Jones' and consider the sentence 'The act we are talking about was intentional'. According to Fine, "this [sentence] can only have a de re meaning, to the effect that the act of killing was intentional; it cannot mean that the act was intentional under the description of my talking about it". Next, consider the sentence 'Smith's act of killing was intentional'. According to Fine this sentence can only have one reading – "in this case, the de dicto reading to the effect that the act was intentional under the description of being a killing" (p.17).

> It therefore looks as if there are two ways of specifying an act, one merely descriptive of the act, telling us *how* the act is (as in 'the act I was talking about') and the other definitive of the act, telling us *what* the act is (as in 'the act of killing'). When the first is used, the resulting sentence is only capable of a de re reading and, when the second is used, the resulting sentence is only capable of a de dicto reading. In neither case, do we have – or normally have – an ambiguous reading, as one would expect under the usual forms of intensionality (*ibid*. Emphasis added).

In short, there are two distinct acts involved when Smith shoots a gun and kills Jones, not one act under two possible descriptions. There is Smith's act of shooting the gun, and there is another act (Smith's act of killing Jones) which has the same basis as Smith's act of shooting the gun but possesses a property that the shooting alone lacks – the property of being intentional. It is because this second act is intentional it qualifies as an *action*.

Something more needs to be said about the 'basis' of an act when that act is a qua object. One of the key principles of Fine's theory (a principle he calls 'Foundation') says, in effect,

that where a qua object is constituted from a series of other qua objects, there is some object at the base or start of that sequence which is not itself a qua object. He describes this as "the *ultimate basis* or *core*" of the object; he calls this core basis '*c*'. The critical question in the case of acts is what this ultimate basis is. For example, "[I]n the case of Smith's act of killing, its immediate basis is perhaps Smith's act of shooting, the immediate basis of this act is then perhaps his pulling the trigger, and so on. But what of *c* itself?" (p. 19). His answer to this question in his (2022) marks a change from his answer from forty years earlier (although he doesn't seem to think which answer is the right one makes a deep difference):

> In Fine (1982), I had proposed taking the core of any bodily act to be a bodily movement. I am now more inclined to think of it as something mental, like an act of trying or willing; and this is what I shall assume in what follows. However, much, though not all, of what I say can be modified so as to accommodate alternative views as to what the core might be (p. 20).

The thrust of my argument throughout this thesis has been that to make a sharp distinction between a bodily movement and "something mental" is a mistake, and the cause of considerable confusion – even in the case of thinking, which is usually taken to be a paradigm case of "something mental". So, I am not going to take sides as to whether the younger Fine or the older Fine is correct, since I reject the dichotomy. On my account, in the case of a thought which is a qua object (i.e. a thought$_a$) the "ultimate basis or core" is an event (which is not a qua object), and that event is the a-intentional[85] formation of an intention.

Interestingly, before coming down on the side of "trying or willing" Fine briefly considers a middle way which is closer to my way of thinking:

> Whether *c* is itself an act is not so clear…. [I]t is perhaps most plausibly taken in the present context not to be an act but some kind of 'act-neutral' event. Acts would then issue from bodily movements by placing them under a suitable gloss (p.20).

[85] This seems a more agentially neutral way of putting it than 'unintentional'. My thanks to Alex Grzankowski for suggesting this term.

On my account, the act-neutral event is the formation of an intention. I have more to say about intention in Section 4.

Section 3. Thoughts as embodied acts

The previous section was intended to provide an overview of Fine's theory of embodiment as it applies to bodily action. In this section I show (a) how these ideas can be successfully applied to thoughts, and (b) that thoughts (or more precisely thoughts_a) should also be understood as qua actions. In his (1982) Fine argues for what he called the "ontological parity" between an intentional human act (i.e., an action) and a material object, such as a statue, which possesses the property of being intentional (a qua object). In this section I want to argue for the ontological parity between an intentional human thought (thought_a) and an intentional human action. If that's right, then it follows that a thought_a is a kind of qua object, one which possesses the property of being intentional. (I should note, here, how unfortunate the term 'qua object' is from my point of view, since I have gone to considerable trouble in this thesis to argue that a thought_a is not *any* kind of object – such a mental state or representation – but a kind of action. As Evnine (2016) points out, it would be much better if Fine had used the term *qua events* when talking about actions. However, I will stick with Fine's terminology to keep things simple.)

To begin with, recall how, on my account, the act of making an (inner speech) utterance is constitutive of the action of making a thought (i.e., either a thought_a or a thought_r). Utterances, I claim, are the 'matter' (or 'basis') of a thought in much the same way that bodily movements are the 'matter' (or 'basis') of an action. Note that the distinction Fine makes between 'act' and 'action' in the case of bodily events has a parallel in the case of thinking, on my view. In chapter 4 I agreed with Gregory (2020) that some of our inner speech utterances are reactive, not intentional, and therefore do not qualify as actions. When an utterance is caused (sustained and guided) by an intention it qualifies as an action – call this a thought_a. But it would do violence to our ordinary use of the term 'thought' to insist that when we produce reactive inner speech utterances – when words just 'pop into our heads' – we are not having thoughts. On my account these are thoughts_r, not thoughts_a

and, as with bodily actions, the 'matter' can be the same in both cases; what is different is the presence or absence of an intention. To illustrate what I mean, consider the case of a token bodily action, where the very same set of movements might be intentional (my raising my arm) or a-intentional (my arm going up). In the same way, the very same natural language expression might be silently uttered intentionally ("I'm going to be late" – spoken in the context of an episode of deliberation concerning when to leave the house to catch a train) or silently uttered reactively (as when the words just pop into my head as I sprint for the bus, caused by the circumstances of my being late).

My account of thinking leans heavily on the distinction between basic and non-basic actions – see Chapter 4. I argue that inner speech utterances are the basic actions *by which* a non-basic action – such as deciding something, or reaching a conclusion, or working something out – is performed. The thinking-as-speaking thesis, as I characterised it in Chapter 6, claims that the (inner) speaking is *constitutive* of the thinking, and this too is in line with Fine's conception of a qua action.

> We can get at the relevant sense [in which we can talk of one act being performed or done by way of another] by insisting that the one act must be constitutively, rather than causally or in some other way, responsible for the other, so that it is by *virtue of* performing the one act that one performs the other (p. 25, original emphasis).

On my account it is by virtue of (inner) speaking that one thinks (i.e., both $thinks_a$ and $thinks_r$).

My account of thinking – or at least, thinking which is deliberative – also fits well with Fine's conception of an activity. Fine makes the following distinction (citing Stout, 1997): "Activities are to acts as processes are to events. Intuitively, an event or act is something that happens or occurs while a process or activity is something that is going on or occurring"

(p. 27).[86] He gives as an example a stroll (an activity) which is constituted by a series of steps (acts or events) taken over time. He goes on:

> I should like to suggest that an activity be identified with a variable embodiment whose manifestations are particular acts. Thus in the case of a stroll, there is some principle ϕ which picks out the different acts which might constitute the stroll; and the stroll itself will be the variable embodiment /ϕ/. We can in this way explain why an activity, as opposed to a sequence of acts, is completely present at each time at which it is going on and why the acts which constitute it can be different from what they actually are (*ibid*).

In line with this description, I want to suggest that when we are speaking about the activity of thinking, the "particular acts" with which its "variable embodiment" should be identified are inner speech acts. In the case of mind wandering, the speech acts might not qualify as actions (since they are not in the service of any specific intention) and so the activity involved might be classified as 'thinking$_r$'. When your mind wanders, it simple responds (reacts) to what preceded it in a way that, by definition, has no particular purpose. By contrast, when you are engaged in what I have called an episode of deliberation, the speech acts qualify as actions because they are caused, sustained, and guided by an intention – trying to reach a conclusion, make a decision, work something out. These utterances will be variable in the sense that there is no determinate set of utterances by which the intention must be achieved. Just as there may be many different kinds of step – fast, slow, limping, etc – by which a stroll could be realised, so there might be many different series of inner speech acts which would allow the subject to achieve her goal of coming to a decision. Call this activity 'thinking$_a$'.

Section 4. Intention

My account of a token thought$_a$ as a species of intentional action is at odds with much of the literature on intention. In this section I explain why that is, and I propose a modest revision

[86] Note the similarity here between Fine's/Stroud's characterisation of 'activity', and Hornsby's (as discussed in Chapter 6, Section 9.) On her view, an activity, unlike an event, is not a particular; activity is a noun but not a count-noun. Other non-count nouns are names of stuff – bronze, for example.

to the way intention should be understood. When the term is used in ordinary discourse, 'an intention' typically refers to a decision or a commitment that has been arrived at following some form of deliberation or reflection. In the extensive philosophical literature on intention the same idea persists. This, from Bratman (1987), is typical:

> Intentions are conduct-controlling pro-attitudes, ones which we are disposed to retain without *reconsideration*, and which play a significant role as inputs into reasoning to yet further intentions (p. 20, emphasis added).

It is clear from the use of the term "*re*consideration" (and from many other places in his book), that while we might experience what Bratman calls "fleeting" intentions (whatever that means), a *genuine* intentional mental state only qualifies as such if it has been arrived at following rational consideration. This conception of an intention obviously won't do for my account. Assuming an episode of practical reasoning is constituted by a series of distinct, albeit connected thoughts, and if, as I have argued, thoughts are (intentional) actions, then an intention can't *always* be the product of "consideration" (i.e., practical reasoning) on pain of an infinite regress. To address this problem we need to take a closer look at the nature of intention.

Pacherie (2008) distinguishes between intentions which exist before the intention is acted on (distal intentions, or 'D-intention' in Pacherie's terminology), and intentions which exist during the fulfilment of the intention (proximal intentions, or 'P-intentions').[87] She writes:

> My notion of D-intentions is very close in certain respects to Bratman's notion of future-directed intentions (Bratman, 1987). Following his lead, we may stress three functions of D-intentions: as terminators of practical reasoning about ends, prompters of practical reasoning about means and plans, and intra- and interpersonal coordinators (p. 182)

[87] This is common among those who write about intention. For example: Searle (1983) distinguishes between prior intentions and intentions-in-action, Bratman (1987) between future-directed and present-directed intentions, Brand (1984) between prospective and immediate intentions, and Mele (1992) between distal and proximal intentions. Unlike these other theorists, ien's (2008) framework includes a third kind of intention: motor intentions (M-intentions). We encountered these in Chapter 8, in the context of skilled action theory.

Pacherie's notion of D-intention therefore generates the same problem as Bratman's notion of future-directed intentions: If D-intentions are "terminators of practical reasoning", and if practical reasoning is (a) something we do intentionally and, (b) involves a series of connected thoughts then, if thoughts are themselves intentional actions, as I have argued, how can a thought be intentional without an infinite regress being generated?

How does an intention to engage in deliberation about something get going? I will illustrate the problem with an example. Suppose I reflect on what time I should leave the house tomorrow morning to ensure that I will get to a meeting on time. The intention (to work out when I should leave the house tomorrow morning) is enough to cause (sustain and guide) the series of questions and answers I utter to myself as I think it through. Suppose, in conclusion, I decide to leave at 11.00. Call that decision a D-intention, one formed "as a terminator of practical reasoning about ends". But what caused me to reflect, in the first place, on when I should leave the house? (In Pacherie's terminology: what "prompted" this episode of "practical reasoning about means and plans"?) Presumably I initiated the episode of deliberation by saying to myself something like "When shall I leave the house tomorrow?". But what caused me to do that? One option, and the one I endorse, is to say that *that* thought (the event of my asking myself that question) was not caused by any prior intention at all; it was an event and not an action. If we take this option, then if someone were to ask me, "Why did you start to reflect on when to leave the house tomorrow?" I would be committed to answering, "I didn't *intend* to reflect on when to leave the house tomorrow, I just started to do it."

It might be thought that there is an alternative option, which is to say that a *higher-level* prior intention was responsible for causing the episode of reflection. Perhaps I had previously formed the intention to always be on time for meetings. Or maybe I made a New Year's resolution to always reflect on my schedule a day in advance, as part of a drive to be more organised. If we take this higher-level option then the event which occurred at the start of my episode of reflection – the event of my asking myself "When shall I leave the house tomorrow?" – qualifies as an action after all, because producing that utterance was an event caused by a prior intention. But the problem with this option is that we can now ask: When I formed my intention to always be on time for meetings, by thinking (i.e., saying

to myself), something like, “This year I am going to try to be on time for all my meetings”, was *that* a mere event, or was it an action, itself caused by *another* prior intention? At some point the sequence of explanations must come to an end and originate in the *event* of the formation of an intention that was not itself the product of deliberation.[88]

Pacherie is right when she says that “D-intentions *can* be the outcome of a conscious process of practical reasoning” (p. 197, emphasis added). But what she (and Bratman and others) has failed to emphasise is that they can’t *all* be. The utterance which signals the onset of an episode of practical reasoning (such as “When shall I leave the house tomorrow?”) falls into the category of a *reactive inner speech utterance* (see Chapter 2) *if* the event (of making the utterance) is not caused by an intention, and therefore it doesn’t qualify as an action. Of course, this is not to say that these events are inexplicable. They can be explained by reference to a subject’s *world view* (see Chapter 6, Section 8 for more on this). A subject’s world view, in combination with the particular circumstances in which she finds herself, explains *all* of a subject’s reactive inner speech utterances, not just the ones that initiate an episode of deliberation. But the explanation doesn’t involve the subject forming an intention as a conclusion to an episode of deliberation. What this means, ontologically, is that at the onset of *any* action, mental or bodily, there is, ultimately, a *non*-qua object: a mental state which represents an intention to do something the generation of which was not itself intentional.

(This returns us to a subject which came up at the end of Section 2, Fine’s idea that where a qua object is constituted from a series of other qua objects, there is some object at the base or start of that sequence which is not itself a qua object. He describes this as “the *ultimate basis* or *core*” of the object; he calls this core basis ‘*c*’. He said, “Whether *c* is itself an act is

[88] I should note here that the idea that intentions can form without the conscious participation, or even awareness of the subject, is orthodoxy among cognitive scientists. For example, Bargh and Fergusson (2000) write, “… the auto-motive [i.e. automatic motivation] model assumes that external events can trigger goals directly, without an explicit conscious choice, and that they then operate without the person knowing of it” (p. 934). They cite the experimental evidence for this claim on pages 934-937. The “auto-motive model” referred to in this quote is the idea, roughly, that “the environment itself can activate a person's goal within a situation, as part of the preconscious analysis of that situation, and this goal then operates in the same manner (without the individual knowing it) as when put into play consciously” (p. 933).

not so clear…. [I]t is perhaps most plausibly taken in the present context not to be an act but some kind of 'act-neutral' event" [REF]. My proposal was that in the case of a thought$_a$ the event (i.e., *c*) is the a-intentional formation of an intention.)

If that's right, then one of the implications of my proposal is that theorists will have to be more careful in future about how they characterise the nature of intention. They should not only emphasise that the formation of an intention *can* be the product of an episode of practical reasoning, as they currently do, but they should also emphasise that they *can't all be*. All intentions must ultimately originate with an "act-neutral event" of intention formation.

Section 5. Summary

In this chapter I have attempted to bolster my claims that thinking should be understood as a species of action, analogous to bodily action in some ways, by invoking Fine's Theory of Embodiment. Fine argues that human actions are made from other, distinct human acts; I argue that thoughts$_a$ qualify as actions in virtue of the fact that they are made (partly constituted) from the act of producing inner speech utterances. Fine accepts that two physical acts are distinct if they have different properties, and that being intentional is a property that one act might have and another physically identical act might not have. I argue that when an inner speech utterance is unintentional (or better, a-intentional) it is an instance of what I call 'reactive' inner speech, and this act expresses a thought which is likewise not intentional but merely reactive – I call this a thought$_r$. On the other hand, the very same utterance might be made actively, intentionally, as part of an episode of deliberation, and the thought it expresses is an action – call this a thought$_a$. In short, on my account, thoughts are a kind of 'embodied act' in line with Fine's theory of embodiment. I also argued that the role that intention plays in my application of Fine's theory to the activity of thinking forces a re-appraisal of some standard theories of intention. Not all intentions can be formed as a result of conscious deliberation or reflection if, as I claim, intentions form a constitutive role in reasoning. Some intentions must form spontaneously, sub-personally, and cause, sustain and guide the process of deliberation.

Chapter 9. *The Mind's Construction* by Matt Soteriou

Section 1. Introduction

One of the most sophisticated and comprehensive recent attempts to describe an ontology of mind and mental action is Matt Soteriou's book *The Mind's Construction* (2013). In this chapter I want to attempt an overview of one part of Soteriou's account, and to contrast it with my own. The book is far broader in its scope and ambition than this thesis and I will focus entirely on Chapter 10, entitled *The Ontology of Conscious Thinking*. What makes Chapter 10 so interesting from the perspective of this thesis is that there appear, on one reading, to be several points of common ground between us. For example, he addresses in depth the role that speaking plays in thinking, both thinking out loud and thinking silently, and he mounts a subtle defence against those who warn against pushing too far the analogy between thinking and speaking. He also argues, in opposition to those who reject the idea of a stream of consciousness, for the existence of what he calls 'occurrent mental states' – conscious mental states with temporal extension. (Recall that my account also needs these. An example is the state of 'being alive to what you are doing' when engaged in a serial cognitive activity such as reasoning or deliberating about something – see Chapter 6, Section 6.) And, in his account of 'thinking out loud', he invokes the notion of a basic, non-reducible action, not unlike the way in which I argue that speaking can be the basic action by which a non-basic action is performed. But these points of apparently common ground hide a deep divide between our two accounts.

My account, by comparison to Soteriou's, is very simple. On my account, whether we are asserting that p, or judging that p out loud, or judging that p silently, *speaking* (or inner speaking in the case of judging that p silently) plays essentially the same (triple) role. The action of making the utterance, (1) is the basic action by which a non-basic (intentional) action is performed; (2) makes a constitutive contribution to determining the meaning of what is uttered; and (3) makes the thought being expressed conscious to the thinker/speaker. In a slogan: the speaking *is* the thinking. The difference between the act of asserting and the act of judging (silently or overtly), on my account, is informed by the

difference in the intention which motivates the performance of the speech act; in the former case, it is an intention to *assert* that p, in the latter case it is an intention to *judge* that p (or judge *whether* p – see below for more on this).

It follows from my account that judging that p (silently or aloud) takes time – the time it takes to generate and perform the speech act. Soteriou, on the other hand, characterises judging that p as an achievement i.e., an instantaneous event, and this principle plays a key role in his ontology of the nature of conscious thinking. It obliges him, I will argue, to provide different accounts for each of the three cases mentioned above: asserting that p, judging that p out loud, and judging that p silently. In short, our respective accounts of conscious thought depend very significantly on which of us is right about the nature of judging. Does it take time, as my account implies, or is it instantaneous, as Soteriou claims? The stakes for both of us are very high. If consciously judging that p is instantaneous, then the thinking-as-speaking thesis is (perhaps seriously) flawed. But if consciously judging that p is not instantaneous then the same might be said of Soteriou's account.

Soteriou's position (that judgment is instantaneous) is, arguably, the orthodox one. Indeed he quotes several other writers who have said the same thing. But in Section 2 I will show that the arguments offered for this position are far from compelling. In the rest of the chapter I show his commitment to this questionable principle plays a key role in the three different accounts he offers for each of three different kinds of mental events: asserting that p (Section 3), judging that p out loud (Section 4), or judging that p silently (Section 5). Section 6 concludes.

Section 2. Consciously judging that p

In this section, I examine Soteriou's commitment to the principle that consciously judging that p is instantaneous, and I argue that it is not well motivated. In fact, Soteriou is committed to two principles, both of which he attributes to Geach (1969): (1) a token thought is individuated by its propositional content, and (2) a token thought is an instantaneous event. I have already offered arguments for thinking that (1) is mistaken (Chapter 6, Section 3). To recap my argument: I agree that thoughts are individuated by

their content (broadly construed), but since thoughts, according to my account, are *made determinate* via the process of performing an inner speech act, it follows that the *propositional content* of a thought (as that term is conventionally understood) is not enough to individuate it. Since it takes time to perform an inner speech act, it follows that, on my account, (2) is also mistaken, and that thoughts are not instantaneous events. I will say no more here about principle (1).

What reasons does Soteriou give for his commitment to the claim that a token thought is an instantaneous event? As I see it, Soteriou offers no *new* argument for the claim; what he does instead is to quote the arguments of others who have claimed the same thing. What do they say? The argument all these writers offer for the claim that judging is an achievement (and instantaneous) is broadly the same: if it wasn't – that is, if a judging took time – then certain locutions would be permitted which (a) we never in fact perform, and (b) sound distinctly odd. To be clear, I don't disagree with these observations, but I don't consider them to be nearly adequate to establish the metaphysical point they are recruited to support.

To make his point, Soteriou begins by quoting Geach (who in turn cites Malcolm):

> In his essay 'What do we think with?', Geach writes:
>
> > I think Norman Malcolm was right when he said that a mental image could be before one's mind's eye for just as long as a beetle took to crawl across a table ... but I think it would be nonsense to say that I 'was thinking' a given thought for the period of a beetle's crawl – the continuous past of 'think' has no such use. (The White Knight 'was thinking' of a plan in that he thought certain thoughts successively; and for each individual thought 'was thinking' would have no application.) (1969: 64) (p. 233).

Well, saying something is nonsense, as Geach does, is not much of an argument. (To recap: on my account, it is *not* nonsense to say that a given thought took time, since the "successive" thoughts that comprise an episode of conscious thinking all involve individual speech acts with temporal duration – see Chapter 6, Section 9.) Furthermore, even if you

were to agree with Geach that "the continuous past of 'think' has no such use", you might wonder why we should allow a particular linguistic convention to underpin our metaphysical commitments. In fact, Soteriou himself wonders the same thing. Immediately after quoting Geach he writes,

> Why should a consideration as to whether the continuous past of the verb has any such use be relevant to the question of whether the mental act of judging is an event with temporal extension? (ibid.).

This strikes me as a very, very good question.[89] But Soteriou doesn't answer it directly. He reminds us of Vendler's (1957) classification of verbs into the four categories of (1) state, (2) achievement, (3) activity, and (4) accomplishment, and of the distinction between them: "Crudely, states are non-dynamic situations, such as be happy or believe; activities are open ended processes, such as run; achievements are near-instantaneous events which are over as soon as they have begun, such as notice; and accomplishments are processes which have a natural endpoint, such as read the book" (p. 6).[90] With these categories in place, Soteriou proceeds to classify the act of making a judgement (e.g., judging that p) as an achievement – "an instantaneous event that lacks duration" (p. 234). This simply restates Geach's claim in different terminology; it doesn't answer the question he posed for himself about the significance of linguistic conventions governing the use of a verb to the ontology of mental acts. Next, to be fair, he considers the alternative position, and asks: Why not think of the event (of judging that p) as an accomplishment, where "the subject is doing something X with the intention of reaching a certain kind of terminus"? (ibid). (This, of course, is exactly how I would classify this sort of event: the subject is doing something we call judging, by performing a speech act, with the intention of coming to a judgement). Soteriou rejects his alternative position on the following grounds:

---

[89] In Chapter 6, Section 9, I made a similar point, saying that what achievement verbs do [e.g. arrive, judge, understand] is to *report* the product, upshot or outcome of some piece of activity, and that this linguistic fact should not be relied on to determine the metaphysical nature of the event they refer to. (I am leaving aside here the worry that no events can *literally* be instantaneous when they involve human animals. Consider, for example, the entry on "The Representation of Time in Agency" in a recent *A Companion to the Philosophy of Time* starts: "Our doings as agents in the world are irreducibly temporally extended" (Anderson, 2013, p. 470).)

[90] Strictly speaking Soteriou is quoting Rothstein's (2004, p.6) distinction between these different verbs. Recall that we encountered these distinctions before, in Chapter 6, Section 9, when we looked at Hornsby's (2012) taxonomy of action verbs.

> If judging were an accomplishment, it should be possible to stop [the subject] S halfway through her act of judging. It should be possible for there to be a situation in which it was not yet true that the subject had judged that *p*, but in which it was true that she had already begun her act of judging that *p*. However, here we might ask, what could the subject have done that counted as having already started the act of judging, and what else would she have needed to do in order to finish it?

With this passage, Soteriou answers his earlier question with another one, the rhetorical point of which is to imply that no possible answer to it makes any sense; but he doesn't explain why that is. Instead, to buttress the implication, Soteriou invokes the authority of Mouton (1969), who asserts the same implication (that no answer to the question would make sense) in slightly different terms:

> It is impossible for one to get halfway through a thought and stop. This is because thoughts are individuated by their content and every such content which comes before one's mind is a complete thought. There is, therefore, no such thing as a partial thought (1969: 65).

But once again, this is not so much an argument as a repetition of the same claim.

On my account, of course, it *is* possible to "get halfway through a thought and stop" because it is possible to get halfway through the performance of an inner speech act and stop. But note that this is not the same thing as saying that there can be such a thing as "a partial thought" if, by that expression, a *whole* thought is conceived of as an object. My account has it that there can be intentions to act (in this case to make a judgement) where the action is initiated and then abandoned before it is completed. On my account, a successful judging that p is an accomplishment, not an achievement, and a non-complete thought is not a *partial thought*, but an *incomplete action*.

It might be objected (against my account) that it is conceptually impossible for a subject to *intend* to judge that p, for the usual (well-rehearsed) reasons. But recall Chapter 5, Section 4. In that section I show how Peacocke (2023b) argues, convincingly in my view, that even if

we accept, for the sake of argument, that we can't intentionally judge *that* p, we can intentionally *judge whether p*, or *judge which thing is F*, or *judge wh-* for all sorts of *wh-*questions. This is all I need to counter the objection. Consider the example that Soteriou uses as an example of a judgment: John is tall. My account does not say that a subject has the explicit intention to judge *that* John is tall; my account says that a subject has the intention to judge *whether* John is tall (or to judge *what* is the best way to characterise John's height, or judge *why* it is that John sticks out from the crowd, or some other *wh-*question, depending on the context). You might worry that I have merely side-stepped the objection, and that I have not provided an account of judging *that* p (the target of Soteriou's comments) but only an account of judging *whether* p. But recall how Peacocke appeals to action theory to explain how one event might also constitute another: "[An] event that executes your intention to *judge whether p*, or your intention to *judge which thing is F*, might indeed constitute a *judgment that p* without thereby constituting an *intentional judgment that p*" (2023b, p. 6. original emphasis).

Furthermore, Peacocke's argument gives us the resources to make a positive argument for why judging that p *should* be understood as an event which takes time, and not an instantaneous one. When a subject has the (non-basic) intention to *judge whether John is tall* (or judge *(what is) John's height?* or any other relevant *wh-* question), she executes that intention by performing the basic action of uttering (silently) 'John is tall', which takes time. Not only that, but because this is a speech act it provides the doxastic commitment we demand of a judgment – the performance of the utterance *constitutes* the judgment, and therefore forms the subject's commitment to its truth. The performance of the utterance constitutes a *judgement*, and not some other attitude which might not involve doxastic commitment (such as a question, or a guess, or a supposition, or a hypothesis), because of the intention which motivated it; it was an intention to *judge*, and not to do anything else.

Meanwhile, we still have no compelling argument for the claim that a judgement that p is an instantaneous event; all we have is protestations to the effect that the contrary assumption would permit locutions which would sound odd, if we ever made them, which we don't. But even granting that point, we have not been given any argument for why linguistic conventions should play such a deciding role in the ontology of conscious thought. And yet

this assumption, I will argue in the following sections, is a fundamental pillar of Soteriou's ontology of conscious thinking.

Section 3. Asserting that p

Before he addresses the phenomenon of *judging* that p Soteriou considers the nature of *asserting* that p, in order to explore the difference. He explains: "An act of asserting is an accomplishment, but an act of judging is not" (p. 240). One reason for the difference is that an act of asserting involves making an (out loud) utterance, which has temporal extension, and the act of judging, by hypothesis, does not. In this section I want to consider how Soteriou characterises an act of assertion, before comparing it, in later sections, with the way he characterises an act of judging. In the case of assertion, what we want is an explanation of the relationship between the act of performing the utterance (out loud), and the content of what is asserted by that act. Soteriou begins with a quote from Geach concerning the utterance of a sentence:

> [U]nless the whole complex content [of the sentence uttered] is grasped all together–unless the Ideas ... are all simultaneously present–the thought or judgement just does not exist at all. (Geach 1957: 104) (p. 238).

To explore what this means, Soteriou applies this line of thought to the act of making an assertion (his example is, 'John is tall'). He makes the point that while the words spoken occur successively, over time, the propositional content of what is asserted has no temporal parts. For example, if I start to assert 'John is tall' and get as far as saying 'John' and then stop, I have not asserted some small part of the propositional content 'that John is tall'. After all, I might have continued by asserting '…took my pen', and of course 'John took my pen' is a completely different assertion from 'John is tall', with different propositional content.

Soteriou suggests that we should think of the temporal event of making the assertion (e.g. saying the words 'John is tall') as a *vehicle* which represents the propositional content being spoken. And he reminds us that we are very familiar with the idea that vehicles of

representation can have different properties from the properties of the content they represent. So, the fact that the act of asserting 'John is tall' has temporal properties, but the content represented by that act (i.e. *that John is tall*) does not, is not problematic. (Later on he offers the analogy of starting a race with a gun shot. The gun shot has temporal extension, but the start of the race (which is what the gun shot represents) does not. It would be odd ask the question, he says, "How long did the starting of the race go on for?". It would be odd because we conceive of the thing *represented* by the gun shot as something instantaneous.) On the other hand, he acknowledges that, even if we accept that the propositional content we invoke in individuating a given event does not have "temporal parts and successive phases", it does not follow from this that the event we thereby individuate (i.e., the event of judging) does *not* have temporal parts and successive phases. It might do:

> [E]stablishing that an event of judging is individuated in terms of such a propositional content does not in itself suffice to show that the event of judging lacks temporal parts and successive phases (p. 239-40).

So how can we be sure that the event of judging that is individuated in terms of propositional content *doesn't* have temporal parts and successive phases? Soteriou's response to this question is to refer us again to Geach, and his argument that it would not make sense to say that a subject has started judging but not yet finished – the topic of Section 2 above. Needless to say, I do not consider this an adequate answer to the question.

On my account, the very act of producing the assertion (the speech act) makes a constitutive contribution to the content (broadly construed) of what is being expressed. Now, it *can* seem as though Soteriou is saying something similar:

> Certain properties of the utterance, e.g. the kinds of sounds that are uttered and the order in which those sounds occur, are relevant to the question of the kind of assertion that is performed when the utterance is made. Those properties of the utterance (e.g. the kinds of

> sounds that are uttered and the order in which they occur) contribute to determining, at least in part, the propositional content that is expressed by that utterance (p. 239).

But in fact Soteriou is saying something different. He is operating with the orthodox conception of propositional content, which has it that the propositional content of asserting that p and judging that p is the same. On his account, when someone makes an assertion, "the kinds of sounds that are uttered and the order in which those sounds occur" makes a difference to the *utterance*, but not to the propositional content which that utterance expresses. Soteriou's key point is that *even though* the event of making the assertion is individuated by the content it expresses, and *even though* the vehicle by which that content is asserted and made conscious has temporal extension, nevertheless the content that is represented by the assertion has no temporal extension. And of course this is different to my account, which has it that an assertion and a judgment are two different speech acts with different content (where content is understood as the meaning of the performance of the speech act).

In summary, in this section we looked at Soteriou's account of asserting that p, and some of the things which make it different from judging that p. Notably: "An act of asserting is an accomplishment, but an act of judging is not" (p. 240). One reason for the difference is that an act of assertion involves making an out loud utterance, which has temporal extension, and the act of judging, by hypothesis, does not. The assertion is the vehicle for the propositional content expressed by the utterance, and we are familiar with the idea that vehicles of representation should have properties which are not shared by whatever it is they represent. However, Soteriou also acknowledges that, even if we accept that the propositional content we invoke in individuating a given event (such as judging that p) does not have "temporal parts and successive phases", it does not follow from this that the event we thereby individuate (i.e., the event of judging) does *not* have temporal parts and successive phases. But he claims that we can be sure that it doesn't because, if it did, we would be permitted to make statements about those events which we don't in fact make and would sound odd if we did. This is an argument I attempted to undermine in Section 2.

Section 4. Silently judging that p[91]

In the previous section we looked at the first of the three kinds of mental event explored by Soteriou: asserting that p, silently judging that p, and judging that p out loud. We saw that what distinguishes the case of asserting that p from judging that p is that it involves an out loud utterance which is the vehicle of the propositional content expressed by the utterance. In the case of asserting that p, the vehicle not only represents the propositional content being expressed, the event of making the utterance also makes the subject *conscious* of what she is asserting. This raises a question for the case of *silently* judging that p: since no utterance is involved, and since the event is, by hypothesis, an achievement, and therefore instantaneous, what makes it the case that the subject is *conscious* of what they are doing when they do it (i.e., when they silently judge that p).

To answer this question Soteriou invokes the notion of an 'occurrent mental state', by which he means something very particular:

> We might think of the conscious mental act of judging as involving the occurrence of a conscious event that has temporal extension— a conscious event that is not reducible to an event/process that is a change in mental states. This would be a conscious mental occurrence with the temporal profile of those aspects of mind that O'Shaughnessy labels 'experiences'— i.e. mental occurrences with the temporal profile of aspects of mind that feature in the stream of consciousness (p. 242).

Soteriou's proposal is that, as with the case of asserting, the content of the mental act of judging has a vehicle of representation, but unlike the case of asserting, the vehicle is not a linguistic utterance, silent or otherwise, but rather an 'conscious occurrent mental state' of the kind described above. Specifically, in the case of silently judging that p, the vehicle for the event is the conscious mental event of *believing that one is judging that p.* As with the case of assertion, he exploits the principle that the vehicle of representation need not share the same properties as the content that is represented.

[91] To help make my point more clearly, I address this issue first, and the issue of judging that p out loud second. Soteriou addresses these issues in the reverse order.

> On this view, the conscious event [i.e. the event of believing that one is judging that p] that is the vehicle of the conscious mental act of judging [i.e. judging that p] has temporal extension and hence temporal parts—but the temporal parts of that conscious event do not signify temporal parts of the content judged. When a subject consciously judges that *p*, that agent does something that we think of as an achievement [i.e. judging that p] *in*, or *by*, ϕ-ing [i.e. believing that one is judging that p], where ϕ-ing is not an achievement (p. 243. Square brackets are inserts by me.)

On this account, two events occur concurrently, the instantaneous act of judging that p and the occurrent event of believing that one is judging that p:

> [W]hat makes it the case that the relevant act of ϕ-ing [i.e. believing that one is judging that p] is the vehicle of a mental act of judging that *p* is the fact that the state that obtains when the act of ϕ-ing occurs is one of believing that one is judging that *p*. On this view, the occurrence of an event of the kind *conscious judging that p* depends upon the *concurrent* obtaining of some mental state of the subject—namely her belief that she is judging that p (p. 248. Again, square brackets are inserts by me.)

I admit to finding this hard to follow. I have a worry about how an event with temporal extension and with the content 'believing that one is judging that p' can be the vehicle for an *instantaneous* conscious event with the content 'judging that p'. To be clear, the worry is *not* about how a vehicle of representation can have a temporal profile different from the propositional content represented by the vehicle. That principle is granted. My worry is about how one could be in a state of continuously believing that one is doing something when the thing one believes one is doing is something which occurs instantaneously. If I form the judgment that the coffee is too hot (to drink), and if that judgment is (by hypothesis) instantaneous, how can I be in an occurrent conscious state of continuously believing I am judging that the coffee is too hot?

To be fair, this worry (or something like it) is anticipated by Soteriou himself, who addresses it in a footnote:

> So the occurrent mental state that is one's belief that one is judging that *p* has temporal extension. But as it is an occurrent state, it does not automatically follow that it is homogeneous down to instants. So although one's belief that one is judging that *p* obtains over an interval of time, it may not be correct to claim that it continues to obtain throughout the interval of time over which it obtains. So it may not be correct to claim that one *continues* to believe that one is judging that *p throughout* an interval of time. Note that this is in keeping with the suggestion that we think of the act of judging as falling under the category of achievement, and not under the category of activity (p. 250, FN 15).

I am not sure what to make of this. On the face of it, either the conscious mental act of silently judging that p involves the occurrence of a conscious event that has temporal extension (i.e., believing that one is judging that p) or it doesn't. If it isn't a single, continuous experience, homogenous down to instants, how can it be said to be obtaining occurrently over an interval of time?

These comments are not intended to be conclusive. What I mainly take away from Soteriou's account of 'silently judging that p' is that, once again, it is forced on him by his commitment to the principle that a judgment is an achievement, and not an activity, and the price he pays for this commitment is, at best, considerable complexity.

Section 4. Judging that p out loud

In the last section we looked at Soteriou's account of the act of judging that p silently. In this section we look at his account of judging that p out loud, and in particular at how this event is different from asserting that p, which also involves an out loud utterance. He rightly rejects the idea of treating the case of 'judging that p out loud' as two distinct events: (1) an achievement (the instantaneous *mental* event of judging, e.g., 'John is tall') plus (2) an accomplishment (the non-instantaneous *bodily* event of asserting 'John is tall'). He explains why he rejects this idea as follows:

> In the case of thinking out loud (e.g. calculating whether *p* out loud), it seems wrong to regard the out-loud utterances as overt actions that merely accompany, and that are separate from, the real mental activity of calculating whether *p*. For then we would not seem to have a

> genuine case of calculating whether *p out loud*, but rather a case of the agent reporting out loud what he or she had just done, or simply being engaged in two distinct activities that happen to be going on at the same time (p. 240).

This seems to me exactly right. But notice that he has made two subtle changes in terminology. The first change is that he has moved from talking about the act of "*judging*" (that p) to talking about the act of "*calculating*". The second change is that he has moved from talking about the act of judging "*that"* to the act of calculating "*whether*" (my emphasis). Now, 'calculating whether p' is a very different mental activity to the activity of 'judging that p'. Here are a couple of examples of calculating whether (out loud):

- He calculated (out loud) whether he had time for another drink before leaving to catch his train.
- She calculated (out loud) whether it was cheaper to take the train or to fly

After doing these various acts of calculating whether, the subject might *conclude* by judging that, yes, he had time for another drink before his train, or she might *conclude* by judging that, yes, it was cheaper to fly than to take the train. The 'judging that', in both cases, is the conclusion to the activity of 'calculating whether'. 'Judging that p', according to Soteriou, is an achievement, but 'calculating whether p' seems, on the face of it, more like an activity, one which, were it to be successful, would result in an accomplishment. This ambiguity finds its way into his proposal for how we should characterise the activity of thinking out loud, which he does as follows:

> [T]he activity of thinking out loud is a non-reducible basic activity—an activity of a basic, non-reducible type, that we might call *mental activity with an overt-bodily-action vehicle* (in this case, calculating whether *p* out loud). According to this way of regarding the case, the verbal utterance instantiates two kinds of activity—overt bodily (talking out loud), and mental (calculating whether p)—in virtue of the fact that it instantiates a third, basic, non-reducible kind of activity, namely a mental activity with an overt-bodily-action vehicle (in this case, calculating whether p out loud). On this view, an event of one's verbal utterance can instantiate two types of act, one's saying something out loud and one's judging that p, *because* it instantiates a third, basic, non-reducible type of act, namely one's judging that p out loud (pp. 240-1, original emphasis).

Notice how, at the end of this characterisation of the activity of thinking out loud, he moves back from the language of 'calculating whether p out loud' to the language of 'judging that p out loud'. But I suggest that, while the above passage contains a good characterisation of an activity which might result in an accomplishment, it is a poor characterisation of an achievement. That's because, as Soteriou himself says, "a subject's out-loud utterance counts as a case of judging that *p* out loud only if the subject believes that she is judging that *p*" (p. 242). Which means that a proper characterisation of judging that p out loud would accommodate all *three* of the elements involved, not just the two mentioned in this characterisation of thinking out loud, namely: judging that p, saying that p out loud, and believing that one is judging that p. The third element seems to be entirely missing from the above characterisation of the activity of thinking out loud. (Certainly, it is no longer needed as a vehicle for the mental activity of 'judging that p', that role having been taken by the overt bodily action of talking out loud.) In short, we don't actually have, from Soteriou, a characterisation of judging that p out loud that explains how all three elements work together, such that it can be clearly distinguished from either asserting that p (out loud) or judging that p (silently). And once again, the difficulties have been generated because of the commitment to the dubious principle that 'judging that p' is an achievement.

Section 5. Summary

Soteriou's (2013) account of the ontology of conscious thought is one of the most sophisticated and comprehensive of recent years. It has the great virtue, in my eyes, of taking seriously the role that language plays in thinking and thought, and of attempting to show how, contrary to the arguments of Strawson (2003), thinking is a mental activity and thoughts are actions. But I have tried to make two points in this chapter. The first is that Soteriou's account is built on a commitment to a principle which is questionable: the idea that judging that p is an achievement, and therefore instantaneous. The only arguments offered for this principle is that if judging that p were *not* instantaneous it would make sense to say certain things which sound odd, or even nonsensical. But even granting this, there is no argument for why we should build our ontology of mental action on this particular linguistic convention.

The second point I have tried to make is that Soteriou's commitment to this principle forces him to construct an account which raises some worries. His account of asserting that p implicitly endorses the orthodox view of propositional content and emphasises the fact that a representational vehicle (such as an out loud assertion of p) might have properties (in this case temporal properties) not possessed by the content by which the assertion is individuated (namely, the propositional content). But from my perspective this ignores the idea that an assertion is a speech act and includes illocutionary force, and therefore, even if it shares propositional content with other speech acts, such as judging that p, its content, construed as its meaning, is different. In the case of silently judging that p, Soteriou proposes that the state of 'believing that one is judging that p' is the vehicle for the act of 'judging that p'. This is puzzling: How can one be in a state of continuously believing that one is doing something when the thing one believes one is doing is something which occurs instantaneously?

Finally, Soteriou proposes an account of 'judging that p out loud' and, once again, his commitment to the principle that judging that p is an instantaneous event prevents him from adopting a more straightforward approach. A more straightforward approach would be to say that judging that p out loud is the same as asserting that p (out loud) but merely motivated by a different intention (the intention to *judge whether p* rather than *assert that p*). Instead he posits a *sui generis* type of activity – *mental activity with an overt-bodily-action vehicle*, and he argues that judging that p is a token instance of this kind of activity. But the way he characterises the activity of thinking out loud is in terms of the activity 'calculating whether p' and this is a very different activity from that of 'judging that p', which must accommodate the occurrent mental state of believing that one is judging that p. Soteriou's account of *thinking* out loud does not, in practice, marry with his account of *judging* out loud, because it fails to accommodate the idea that, by hypothesis, judging that p out loud also involves the subject *believing* that she is judging that p.

Chapter 10. Summary

Executive summary

It is widely assumed by philosophers of mind that (1) thinking is a purely mental activity, and (2) a token thought is an object of some kind 'in the head' of the thinker – either a sentence in the language of thought, or some other kind of representation of the content of the thought. This thesis challenges both assumptions. The first assumption is challenged by the following set of arguments: inner speech qualifies, at least sometimes, as a kind of intentional action; inner speech is an internalisation of, and continuous with, outer speech; speaking is a form of practical knowledge, and exercising that knowledge by making an utterance is sometimes the basic action by which a non-basic action is executed; thinking silently and thinking out loud are not fundamentally different activities. Taken together, these arguments strongly suggest that thinking is dependent in some deep way on the ability of a subject to use a natural language, and that even the silent use of that ability is a physical, or at least a quasi-physical, activity. The second assumption is challenged by the following set of arguments: inner speech utterances are very often inner speech acts – they are an amalgam of illocutionary force and propositional content; the very act of generating and performing an inner speech utterance makes a constitutive contribution to the content of what is expressed by it; the skill involved in thinking has something in common with skilled bodily actions – the skill of the execution makes a constitutive contribution to the final performance. Taken together, these arguments suggest that a token thought is not an object of any kind, but rather a type of action – specifically the action of performing an inner speech act. In short, some kinds of thinking should be understood as a kind of embodied activity.

Chapter summaries

The rest of this chapter reproduces the summaries of each of the preceding chapters (excluding Chapter 1, the Introduction.)

Chapter 2. Inner speech: what are we talking about?

This chapter has been an introduction to the phenomenon of inner speech, focusing on the latest scientific and psychological understanding of the phenomenon. There is widespread support for the Vygotskian idea that inner speech should be understood as an internalised form of outer speech. There is also strong evidence for the motor-sensory view of inner speech – the idea that it is embodied, and that it involves physical processes that unfold over time. I described the different stages of Levelt *et al*'s (1999) high level theory of speech production, which is neutral on the question of whether inner speech is abstract or motor-sensory in nature, but then I offered evidence (from Loevenbruck *et al,* 2018), for the motor-sensory view. To explain how we come to experience an inner voice at all I offered a detailed description of the predictive control account of inner speech, according to which our inner speech is, in effect, the experience of an audio image of a prediction (or forecast) of what our utterances, as generated by the language production process, *would* sound like *were* they to be voiced. The instruction to voice them is inhibited, so the words are not spoken but nevertheless 'heard' by one's 'inner ear'. I explained how the predictive control account, when applied to the language production process, is thought to be responsible for a sense of ownership of our thoughts, and therefore of our sense of agency. I ended by registering how this observation puts the phenomenon of inner speech at the heart of current debates about mental action.

Chapter 3. An inconvenient fact: the phenomenon of unsymbolised thinking

Anyone who wants to explore the relationship between language and thought has to address an inconvenient fact: research suggests that quite a large percentage of people experience conceptual thought without experiencing language. How can this be? And how is such a finding compatible with the sensory motor view of inner speech, which has it that inner speech is the experience of 'hearing' an auditory image which is a prediction, in effect, of what the speech would sound like were it to be vocalised? I offer an explanation, from Vicente and Jorba (2019), which proposes that, before the utterance is submitted to the motor planning system (which is responsible for generating an efference copy of the motor plan), the utterance is available to the language comprehension system, and its meaning made available to the subject. If, immediately after that, the intention to make the utterance is abandoned, and the motor planning process aborted, the subject will have the

experience of the meaning of what they were thinking without the experience of the language that determined the meaning.

Chapter 4. Action and reaction: the two voices of inner speech

Some inner speech utterances are reactive: they are spontaneous, they require no effort, and we are not in control of their occurring. These inner speech utterances fail to meet the standard criteria for qualifying as intentional actions. But some inner speech utterances are genuine actions, performed deliberately, effortfully and with as much control as any other intentional action. For example, when we engage in an episode of deliberation, we are performing the non-basic action of trying to bring it about that we achieve some cognitive goal – coming to a decision, reaching a conclusion, solving a problem, etc. The action of trying to bring this goal about is achieved by performing basic actions – making inner speech utterances. An inner speech utterance, when made in this context, stands in the same relation to the intentional action of trying to (say) reach a decision, as raising an arm has to the intentional action of trying to (say) attract a waiter. That is, it stands in the relation of a basic action to a non-basic action.

Chapter 5. Is reasoning intentional?

Reasoning is an intentional activity which is at least partly constituted by the action of making inner speech utterances. These utterances are the basic actions by which the non-basic action of reasoning-with-a-view-to-deciding-something is executed and as such they qualify as intentional. The inner speech utterances are intentional whether they take the form of questions or answers. *Pace* Frankish (2018) reasoning is wholly intentional. During an episode of reasoning the inner speech utterances which result in the subject coming to a decision have the same characteristics as other exploratory actions – they involve an interrogative attitude and a kind of expectant attention. These features give reasoning its characteristic phenomenology, and this is responsible for the misleading impression that some of the inner speech utterances involved in reasoning are actions (albeit basic ones), and some are not when, in fact, they all are.

Chapter 6. A New Argument for "Thinking-as-Speaking"

Some philosophers are attracted to what I have called the *thinking-as-speaking* thesis – the idea that some thoughts *just are* inner speech utterances, and vice versa. Standing in the way of this thesis is the popular idea that the content of a thought is fully determined before it is translated from a non-linguistic format into the thinker's natural language. I argued against this by appealing to speech act theory. I argued that many inner speech utterances are inner speech *acts*. If that's right, then inner speech utterances, like other speech acts, are individuated by their meaning in the fullest sense, an amalgam of propositional content and illocutionary force. I argued that the determinate content of a token thought only comes into being with the production of a natural language expression, and therefore could not be represented beforehand by a mental state. That being so, the thought being expressed by the inner speech utterance has a different content from the content of any mental state which might have preceded the performance of the utterance. Since a thought is said to be individuated by its content it follows that the thought expressed by a speech act is not the same as whatever content/information existed (in a non-linguistic format) before the performance of the speech act. I suggested that, based on this analysis, we should reject the assumption that a thought is any kind of object, such as a mental state or an expression in a natural language, and we should recognise that a thought is a kind of action – the action of performing a speech act. I buttressed this proposal by arguing that thinking is a type of activity, with the same ontological status as other (bodily) activities, such that individual actions should be conceived of as the product, upshot or outcome of a type of activity. My proposal casts doubt on the assumption which informs much thinking about thinking – that a token thought is individuated by its content. That's because if a token thought is understood as a token thinking, and if a token thinking is understood as the performance of an inner speech act, and if the *content* of a token thought is understood as what the performance of the relevant speech act means, then, because the meaning of speech acts is indeterminate, so too is the content of the thought.

Chapter 7. An objection to *thinking-as-speaking*, and a response

According to what I have called the *strict hybrid* view, skills have two components: a cognitive component, and a motor component. The cognitive component involves propositional knowledge, and the motor component involves automatic, low-level causal processes, acquired through brute repetition. All the intelligence associated with the skill is

associated with the cognitive component; the motor component is dumb – little more than reflex. The hybrid view can be used to pose an objection to the thinking-as-speaking thesis, by arguing that thinking is the cognitive component of this activity, speaking is the motor component, and all the intelligence involved is supplied by the cognitive component. On this account, contrary to the claims of Chapter 6, the speaking component would make no contribution to the content of what is said by an utterance; the speaking component would merely execute the production of a meaningful expression already fully specified in a non-linguistic format.

I offered two arguments against this objection. The first is that the hybrid view is wrong in the case of skilled bodily actions. Skilled bodily actions show none of the characteristics one would expect them to show if the strict hybrid view was right. The second is to pose the objector a dilemma. If the objection is confined to the last stage of the language production process, then it misses its mark, since by then the language production process has produced a meaningful utterance and made its constitutive contribution to the meaning of that utterance. If, on the other hand, the objection extends to the whole of the language production process, the implication is that all the meaning of a natural language utterance is entirely determined before the language production process has even begun. For various reasons, this is tremendously implausible.

Chapter 8. Thoughts as embodied acts

In this chapter I have attempted to bolster my claims that thinking should be understood as a species of action, analogous to bodily action in some ways, by invoking Fine's Theory of Embodiment. Fine argues that human actions are made from other, distinct human acts; I argue that thoughts$_a$ qualify as actions in virtue of the fact that they are made (partly constituted) from the act of producing inner speech utterances. Fine accepts that two physical acts are distinct if they have different properties, and that being intentional is a property that one act might have and another physically identical act might not have. I argue that when an inner speech utterance is unintentional (or better, a-intentional) it is an instance of what I call 'reactive' inner speech, and this act expresses a thought which is likewise not intentional but merely reactive – I call this a thought$_r$. On the other hand, the very same utterance might be made actively, intentionally, as part of an episode of

deliberation, and the thought it expresses is an action – call this a thought$_a$. In short, on my account, thoughts are a kind of 'embodied act' in line with Fine's theory of embodiment. I also argued that the role that intention plays in my application of Fine's theory to the activity of thinking forces a re-appraisal of some standard theories of intention. Not all intentions can be formed as a result of conscious deliberation or reflection if, as I claim, intentions form a constitutive role in reasoning. Some intentions must form spontaneously, sub-personally, and cause, sustain and guide the process of deliberation

Chapter 9. *The Mind's Construction* by Matt Soteriou

Soteriou's (2013) account of the ontology of conscious thought is one of the most sophisticated and comprehensive of recent years. It has the great virtue, in my eyes, of taking seriously the role that language plays in thinking and thought, and of attempting to show how, contrary to the arguments of Strawson (2003), thinking is a mental activity and thoughts are actions. But I have tried to make two points in this chapter. The first is that Soteriou's account is built on a commitment to a principle which is questionable: the idea that judging that p is an achievement, and therefore instantaneous. The only arguments offered for this principle is that if judging that p were *not* instantaneous it would make sense to say certain things which sound odd, or even nonsensical. But even granting this, there is no argument for why we should build our ontology of mental action on this particular linguistic convention.

The second point I have tried to make is that Soteriou's commitment to this principle forces him to construct an account which raises some worries. His account of asserting that p implicitly endorses the orthodox view of propositional content and emphasises the fact that a representational vehicle (such as an out loud assertion of p) might have properties (in this case temporal properties) not possessed by the content by which the assertion is individuated (namely, the propositional content). But from my perspective this ignores the idea that an assertion is a speech act and includes illocutionary force, and therefore, even if it shares propositional content with other speech acts, such as judging that p, its content, construed as its meaning, is different. In the case of silently judging that p, Soteriou proposes that the state of 'believing that one is judging that p' is the vehicle for the act of 'judging that p'. This is puzzling: How can one be in a state of continuously believing that

one is doing something when the thing one believes one is doing is something which occurs instantaneously?

Finally, Soteriou proposes an account of 'judging that p out loud' and, once again, his commitment to the principle that judging that p is an instantaneous event prevents him from adopting a more straightforward approach. A more straightforward approach would be to say that judging that p out loud is the same as asserting that p (out loud) but merely motivated by a different intention (the intention to *judge whether p* rather than *assert that p*). Instead he posits a *sui generis* type of activity – *mental activity with an overt-bodily-action vehicle*, and he argues that judging that p is a token instance of this kind of activity. But the way he characterises the activity of thinking out loud is in terms of the activity 'calculating whether p' and this is a very different activity from that of 'judging that p', which must accommodate the occurrent mental state of believing that one is judging that p. Soteriou's account of *thinking* out loud does not, in practice, marry with his account of *judging* out loud, because it fails to accommodate the idea that, by hypothesis, judging that p out loud also involves the subject *believing* that she is judging that p.

Acknowledgements

I want to thank my supervisor, Alex Grzankowski, for his support, interest and encouragement over the years I have worked on this thesis. He has been a great teacher.

I also want to thank the other philosophy teachers at Birkbeck, and the other graduate students. Collectively they create a remarkable ethos – supportive and encouraging, but also rigorous and challenging. Long may it last.

Reference List

Aflalo, T. et al. (2022) 'Implicit mechanisms of intention', *Current Biology* 32, 2051–2060.

Alderson-Day, B. & Fernyhough, C. (2015) 'Inner Speech: Development, Cognitive Functions, Phenomenology, and Neurobiology', *Psychological Bulletin*, Vol. 141, No. 5, 931-965.

Amaya, S. (2017) 'Basic actions reloaded', *Philosophy Compass*, Volume 12, Issue 9, 12:e12435.

Amaya, S. (2016 'Slip-Proof Actions', in Roman Altshuler & Michael J. Sigrist (eds.), *Time and the Philosophy of Action,* Routledge. pp. 21-36. 2016.

Anderson, H. (2013) 'The Representation of Time in Agency', in *A Companion to the Philosophy of Time*, Bardon, A. and Dyke, H. (eds), John Wiley & Sons Incorporated. *ProQuest Ebook Central*.

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

Armstrong, J & Pavese, C. (2024) 'Language as Skill'. Manuscript, available on *PhilPapers.*

Austin, J. L. (1962) *How to Do Things with Words*. Oxford: Oxford University Press.

Baars, B. J. (2003) 'How brain reveals mind: Neural studies support the fundamental role of conscious experience'. *Journal of Consciousness Studies*, 10, 100-14.

Bach, K. (1980) 'Actions are not Events', *Mind*, New Series, Vol. 89, No. 353, pp. 114-120.

Balcarras, D. (2023) 'What is it to have a language?', *Pacific Philosophical Quarterly*, DOI: 10.1111/papq.12446.

Bargh, A. & Ferguson, M. (2000). 'Beyond Behaviorism: On the Automaticity of Higher Mental Processes', *Psychological Bulletin*, Vol. 126, No. 6, 925-945.

Bayne, T. J. & Levy, N. (2006) 'The feeling of doing: Deconstructing the Phenomenology of Agency'. In N. Sebanz & W. Prinz (Eds.), *Disorders of volition*. Cambridge, Mass: MIT Press.

Bayne, T. J. & Montague, M. (2011) *Cognitive Phenomenology*, OUP, Oxford.

Bermudez, J. L. (2000) 'Personal and Sub-personal; a difference without a distinction', *Philosophical Explorations*, 3:1, 63-82.

Borg, E. (2022) 'In Defence of Individual Rationality', *Proceedings of the Aristotelian Society*, 142nd Session - 2021/22 - Volume CXXII.

Brand, M. (1984) *Intending and acting*. Cambridge, MA: MIT Press.

Bratman, M. E. (1987) *Intention, plans, and practical reason*. Cambridge, MA: Cambridge University Press. [Page references in my text refer to 'Bratman, M. (1999). *Intention, plans, and practical reason*. CSLI Publications, USA.']

Brozzo, C. (2017) 'Motor Intentions: How Intentions and Motor Representations Come Together', *Mind & Language*, Vol. 32, No. 2 April, pp. 231–256

Camp, E. (2007) 'Thinking with Maps', *Philosophical Perspectives, 21, Philosophy of Mind.*

Camp, E. (2009) 'A language of baboon thought?' in *Philosophy of Animal Minds*, R. Lurz (ed). Cambridge: Cambridge University Press.

Campbell, J. (1999) 'Schizophrenia, the Space of Reasons, and Thinking as a Motor Process', *The Monist*, Vol. 82, No. 4, Cognitive Theories of Mental Illness (OCTOBER), pp. 609-625.

Carruthers, P. (2007) 'The Illusion of Conscious Will', *Synthese*, Nov., Vol. 159, No. 2, pp. 197-213.

Carruthers, P. (2009) 'How we know our own minds: The relationship between mindreading and metacognition', *Behavioral and Brain Sciences*, 32, 121–182.

Carruthers, P. (2011) *The Opacity of Mind: An Integrative Theory of Self-Knowledge*, Oxford: Oxford University Press.

Carruthers, P. (2018) 'The Causes and Contents of Inner Speech', in P. Langland-Hassan & A. Vicente (eds.) *Inner Speech New Voices.* Oxford University Press.

Carruthers, P & Veillet, B (2011) 'The Case Against Cognitive Phenomenology', in Bayne, & Montague, (eds.) *Cognitive Phenomenology*. OUP, Oxford.

Cassam, Q. (2011) 'Knowing What You Believe', *Proceedings of the Aristotelian Society*, 111: 1–23.

Chater, N. (2018) *The Mind is Flat: The Illusion of Mental Depth and the Impoverished Mind.* Penguin Random House: UK.

Chomsky, N. (2000) *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.

Christensen, W., Sutton, J. and McIlwain, D.J.F. (2016) 'Cognition in Skilled Action: Meshed Control and the Varieties of Skill Experience', *Mind & Language*, Vol. 31, No. 1, February, pp. 37-66.

Cooper, J. M., & Hutchinson, D. S. (1997) *Plato: complete works*. Hackett Publishing.
Crane, T. (2017) 'The Unity of Unconsciousness', *Proceedings of the Aristotelian Society*, Issue No 1, Volume CXVII.

Crowther, T. (2011) 'The Matter of Events'. *The Review of Metaphysics*, Vol. 65, No. 1, pp. 3-39 .

Davidson, D. (1963) 'Actions, Reasons and Causes', *Journal of Philosophy*, 60, pp. 685-700.

Davidson, D. (1973) 'On the Very Idea of a Conceptual Scheme.' *Proceedings and Addresses of the American Philosophical Association*, Vol. 47, pp. 5-20.

Davidson, D**. (**1971) 'Agency', in *Agent, Action and Reason*. In R. Binkley, R. Bronaugh, A. Marras (eds.), Toronto: University of Toronto Press, 3–37; reprinted in Davidson (1980), 163–80.

Deamer, F. (2021) 'Why do we talk to ourselves?'. *Review of Philosophy and Psychology* 12:425–433.

Dickie, I. (2012) 'Skill Before Knowledge', *Philosophy and Phenomenological Research*, Vol. LXXXV No. 3. Pp. 737-745.

Dummett, M. (1959) 'Truth', *Proceedings of the Aristotelian Society* 59:141-62.

Endicott, R. (Forthcoming in *Frontiers in Psychology*) 'Inner Speech and the Body Error Theory.'

Evans, J. St. B.T. (2003) 'In two minds: dual-process accounts of reasoning', *TRENDS in Cognitive Sciences* Vol.7 No.10 October.

Evans, J. St. B. T. and Stanovich, K. (2013) 'Dual-Process Theories of Higher Cognition: Advancing the Debate', *Perspectives on Psychological Science* 8(3) 223-241.

Evnine, Simon J. (2016) *Making Objects and Events: A Hylomorphic Theory of Artifacts, Actions, and Organisms*. Oxford, online edition, Oxford Academic.

Feinberg, I. (1978) 'Efference Copy and Corollory Discharge: Implications for Thinking and its Disorders.' *Schizophrenia Bulletin*, 4: 636-640.

Fernyhough, C. (1996) 'The Dialogic Mind: A Dialogic Approach to the Higher Mental Functions'. *New Ideas in Psychology,* Vol. 14, No. 1, pp.47-52.

Fernyhough, C. (2008) 'Getting Vygotskian about theory of mind: Mediation, dialogue, and the development of social understanding'. *Developmental Review* 28 (2008) 225–262.

Ferretti, G. (2016) 'Pictures, action properties and motor related effects'. *Synthese, Special Issue: Neuro-science and Its Philosophy, 193*(12), 3787–3817. https://doi.org/10.1007/s11229-016-1097-x.

Ferretti, G. (2017) 'Two visual systems in molyneux subjects'. *Phenomenology and the Cognitive Sciences*. https://doi.org/10.1007/s11097-017-9533-z.

Ferretti, G. (2020) 'Anti-intellectualist motor knowledge', *Synthese*, https://doi.org/10.1007/s11229-020-02750-9.

Fine, K. (1982) 'Acts, Events and Things', in Leinfellner, W. (ed.), *Language and Ontology: proceedings of the Sixth International Wittgenstein Symposium 23rd to 30th August 1981, Kirchberg am Wechsel*, 97–105. Austria, Holder-Pichler-Tempsky, Vienna.

Fine, K**. (**1991) 'The Identity of Material Objects'. In Albertazzi, L and Poli, R (eds.), *Topics in Philosophy and Artificial Intelligence*, 33–38. Insituto Mitteleuropeo di Cultura, Bozen.

Fine, K**. (**1999) 'Things and Their Parts'. *Midwest Studies in Philosophy*, 23(1): 61–74. DOI: https://doi. org/10.1111/1475-4975.00004.

Fine, K**. (**2000) 'A Counter-Example to Locke's Thesis', *Monist*, 83(3): 357–361. DOI: https://doi.org/10.5840/ monist200083315.

Fine, K. (2003) 'The Non-Identity of a Material Thing and Its Matter', *Mind*, April, Vol. 112, No. 446, pp. 195-234.

Fine, K. **(**2007). 'Arguing for Non-identity: A Response to King and Frances', *Mind,* 115(460): 1059–1082. DOI: https://doi.org/10.1093/mind/fzl1059

Fine, K. (2022) 'Acts and Embodiment', *Metaphysics*, 5(1), pp. 14–28. DOI: https://doi. org/10.5334/met.98.

Fogal, D., Harris, D. W., Moss, M. (eds). (2018) *New Work on Speech Acts*. Oxford: Oxford University Press.

Forster, E.M. (1927) *Aspects of the Novel.* (reprinted Harmondsworth: Penguin 1976).

Frankfort, T. (2022) 'Action and Reaction: The Two Voices of Inner Speech', *Teorema*, Vol. XLI/1, pp. 51-69.

Frankfort, T. (Forthcoming) 'A new argument for 'Thinking-as-Speaking', *Philosophical Explorations.* (Published online: 12 Jan 2024.)

Frankfurt, H. (1978) 'The Problem of Action', *American Philosophical Quarterly*, 15, pp. 157-162.

Frankfurt, H. (1988) 'Identification and Externality', in *The Importance of What We Care About: Philosophical Essays*, pp. 58–68. Cambridge: Cambridge University Press.

Frankish, K. (2009) 'Systems and levels: Dual-system theories and the personal– subpersonal distinction**'**, in J. Evans and K. Frankish (eds.), *In Two Minds: Dual Processes and Beyond*. Oxford: Oxford University Press.

Frankish, K. (2010) 'Dual-Process and Dual-System Theories of Reasoning', *Philosophy Compass* 5/10: 914–926.

Frankish, K. (2018) 'Inner Speech and Outer Thought', in P. Langland- Hassan & A. Vicente (eds.). *Inner Speech New Voices.* Oxford University Press.

Fridland, E. (2017) 'Skill and motor control: intelligence all the way down', *Philosophical Studies*, Vol. 174, No. 6, pp. 1539-1560.

Fridland, E. (2019) 'Intention at the Interface', *Review of Philosophy and Psychology*, 12:481–505.

Fridland, E. (2019) 'Longer, smaller, faster, stronger: On skills and intelligence', *Philosophical Psychology*, Vol. 32, No. 5, 760-784.

Fridland, E. (2020) 'The Nature of Skill', in Fridland, E., & Pavese, C. (Eds.) *The Routledge Handbook of Philosophy of Skill and Expertise* (1st ed.). Routledge. https://doi-org.ezproxy.lib.bbk.ac.uk/10.4324/9781315180809

Friedman, J. (2019) 'Inquiry and Belief', *Noûs*, 53: 296- 315. https://doi.org/10.1111/nous.12222.

Frith, C. (2012) Explaining delusions of control: the comparator model 20 years on. *Consciousness and Cognition* 21, 52–54.

Gallese, V., & Metzinger, T. (2003) 'Motor ontology. The representational reality of goals, actions and selves', *Philosophical Psychology, 16*(3), 365–388.

Gauker, C. (2018) 'Inner Speech as the Internalization of Outer Speech', in P. Langland-Hassan & A. Vicente (eds.) *Inner Speech New Voices.* Oxford University Press.

Geach, P. (1957) *Mental Acts*. London: Routledge and Kegan Paul.

Geach, P. (1969) *God and the Soul*. London: Routledge and Kegan Paul.

Geurts, B. (2018) 'Making Sense of Self Talk', *Review of Philosophy and Psychology*. 9: 271–285.

Godfrey-Smith, P. (2017) *Other Minds*. London: William Collins.

Goldwasser, S. (2022) 'Memory as Skill'. *Review of Philosophy and Psychology*, published online: 04 January, https://doi.org/10.1007/s13164-021-00605-x.

Goodin, R. E. (2024) 'Consent as an act of commitment', *European Journal of Philosophy*, 32: 194-209, DOI: 10.1111/ejop.12852.

Grandchamp R, *et al* (2019) 'The ConDialInt Model: Condensation, Dialogality, and Intentionality Dimensions of Inner Speech Within a Hierarchical Predictive Control Framework', *Front. Psychol. 10:2019. doi: 10.3389/fpsyg.2019.02019.*

Gregory, D. (2020) 'Are Inner Speech Utterances Actions?' *teorema*, Vol. XXXIX/3, pp. 55-78.

Hanks, P. W. (2007) 'The Content-Force Distinction', *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, May, Vol. 134, No. 2, pp. 141-164.

Hanks, P. W. (2015) *Propositional Content*. Oxford: Oxford University Press.

Harris, D. (forthcoming) 'Semantics without Semantic Content', *Mind and Language.*

Hanks, P. W. (2018) 'Types of Speech Acts', in *New Work on Speech Acts*, Fogal, D., Harris, D. W., & Moss, M. (eds.), Oxford: Oxford University Press.

Heavey, C. L. & Hurlburt, R. T. (2008) 'The phenomena of inner experience', *Consciousness and Cognition,* 17, 798–810.

Heck, R.G. (2014) 'Semantics and Context-Dependence: Towards a Strawsonian Account', in Burgess, A. & Sherman, B. (eds.), *Metasemantics: New Essays on the Foundations of Meaning*, 327-364, Oxford: Oxford University Press.

Hinshelwood, A. (2023) 'Representation in Action', *European Journal of Philosophy*, DOI: 10.1111/ejop.12932.

Hornsby, J. (1980) *Actions*. London: Routledge and Kegan Paul.

Hornsby, J. (2005) 'Semantic knowledge and practical knowledge', *Proceedings of the Aristotelian Society, Supplementary Volumes*, 107-145.

Hornsby, J. (2009) 'Speech acts and performatives', in *The Oxford Handbook of Philosophy of Language*, Lepore, E and Smith, B.C. (eds), Oxford Handbooks Online, https://doi.org/10.1093/oxfordhb/9780199552238.003.0035

Hornsby, J. (2012) 'Actions and Activity', *Philosophical Issues*, Vol. 22, Action Theory, pp. 233-254.

Hornsby, J. (2013) 'Basic Activity', *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 87, pp. 1-18, Oxford: Oxford University Press.

Hurlburt, R. T. (1990) *Sampling normal and schizophrenic inner experience*. New York: Plenum Press.

Hurlburt, R. T. (1993) *Sampling inner experience in disturbed affect*. New York: Plenum Press.

Hurlburt, R. T., & Akhter, S. A. (2006) 'The Descriptive Experience Sampling method', *Phenomenology and the Cognitive Sciences*, 5, 271–301.

Hurlburt, R. T., & Heavey, C. L. (2001) 'Telling what we know: Describing inner experience', *Trends in Cognitive Sciences*, 5, 400–403.

Hurlburt, R. T., & Heavey, C. L. (2002) 'Interobserver reliability of Descriptive Experience Sampling', *Cognitive Therapy and Research*, 26, 135–142.

Hurlburt, R. T., & Heavey, C. L. (2006) *Exploring inner experience: The Descriptive Experience Sampling method*. Amsterdam/Philadelphia: John Benjamins.

Hurlburt, R. T., Heavey, C. L., & Seibert, T. (2006) 'Psychological science's prescription for accurate reports about inner experience', in Hurlburt & Heavey, *Exploring inner experience: The Descriptive Experience Sampling method* (pp. 41–60). Amsterdam/Philadelphia: John Benjamins.

Hurlburt, R. T., & Schwitzgebel, E. (2007) *Describing inner experience? Proponent meets sceptic*. Cambridge, MA: MIT Press.

Hurlburt, R. T. & Akhter, S. A. (2008) 'Unsymbolized Thinking*', Consciousness and Cognition* 17: 1364–1374.

Hurlburt, R. T., Heavey, C. L. & Kelsey, J. M. (2013) 'Toward a Phenomenology of Inner Speaking', *Consciousness and Cognition*, 22: 1477–1494.

Hurlbert, R. T. & Heavey, C. L. (2018) 'Inner Speaking as Pristine Inner Experience', in P. Langland- Hassan & A. Vicente (eds.) *Inner Speech New Voices.* Oxford: Oxford University Press.

Jackson, J. H. (1866/1932) 'The Selected Writings of John Hughlings Jackson. Volume 1. On Epilepsy and Epileptiform Convulsions', *Archives of Neurology and Psychiatry* 27(3):757.

Jeannerod, M. (2006) *Motor cognition: What actions tell the self* . Oxford: Oxford University Press.

Jenkins, D. (2021) 'Reasoning and Its Limits', *Synthese*, https://doi.org/10.1007/s11229-021-03211-7.

Jones, S. R., and Fernyhough, C. (2007) 'Thought as action: Inner speech, self-monitoring, and auditory verbal hallucinations', *Consciousness and Cognition*, 16, 391-399.

Kahneman, D. (2011) *Thinking, fast and slow*. London: Allen Lane.

Klinger, E., & Cox, W. M. (1987-1988) 'Dimensions of thought flow in everyday life', *Imagination, Cognition and Personality, 7*(2), 105–128.

Korsgaard, C. M. (2009) 'The Activity of Reason', *Proceedings and Addresses of the American Philosophical Association*, Vol. 83, No. 2, pp. 23-43.

Kompa, N. A. (2023) 'Inner Speech and 'Pure' Thought – Do we Think in Language?**',** *Review of Philosophy and Psychology.* DOI:org/10.1007/s13164-023-00678-w.

Kompa, N. A. & Mueller, J. L. (2022) 'Inner speech as a cognitive tool—or what is the point of talking to oneself?', *Philosophical Psychology*. DOI: 10.1080/09515089.2022.2112164.

Langland-Hassan, P. and Vicente, A. (eds). (2018) *Inner Speech: New Voices.* Oxford, Oxford University Press.

Levinson, S. (2003) 'Language and mind: Let's get the issues straight', in D. Gentner & S. Goldin-Meadow (eds.), *Language in Mind: Advances in the Study of Language and Thought*, 25-46. Bradford Book.

Levelt, W. J. M. *et al* (1999) 'A theory of lexical access in speech production', *Behavioral and brain sciences*, 22, 1–38.

Levy, N. (2017) 'Embodied savoir-faire: knowledge-how requires motor representations', *Synthese*, Vol. 194, No. 2, Special Issue, pp. 511-530.

Levy, Y. (2019) 'What is 'mental action'?'. *Philosophical Psychology*, Vol. 32, No. 6. 971-993.

Loevenbruck *et al* (2018) 'A cognitive neuroscience view of inner language', in Langland-Hassan P. & Vicente A. (eds.) *Inner Speech New Voices*. Oxford: Oxford University Press.

Machery, E. (2018) 'Know Thyself: Beliefs vs. Desires in Inner Speech', in Langland-Hassan P. and Vicente A. (eds) *Inner Speech: New Voices*. Oxford: Oxford University Press.

Martínez-Manrique, F. and Vicente, A. (2010)., 'What the ... !' The role of inner speech in conscious thought', *Journal of Consciousness Studies*, 17(9–10), 141–67.

Martínez-Manrique, F. & Vicente, A. (2015) 'The activity view of inner speech', *Frontiers in Psychology*, *6*, 232.

Mele, A. R. (1992) *Springs of action*. Oxford: Oxford University Press.

Mele, A. (2009) 'Mental Action: A Case Study', in O'Brian, L. and Soteriou, M. (eds), *Mental Actions.* Oxford: Oxford University Press.

Mendelovici, A. (2018) *The Phenomenal Basis of Intentionality*. Oxford: OUP.

Mourelatos, A.P.D. (1978) 'Events, Processes, and States', *Linguistics and Philosophy* 2: 415-34.

Mouton, D.M. (1969) 'Thinking and Time', *Mind* 78:60-76.

Mylopoulos, M. & Pacherie, E. (2018) 'Intentions: The dynamic hierarchical model revisited', *WIREs Cognitive Science*; e1481. https://doi.org/10.1002/wcs.1481

Mylopoulos, M. (2020) 'The Intelligence of Motor Control' in Fridland, E., & Pavese, C. (eds.) *The Routledge Handbook of Philosophy of Skill and Expertise* (1st ed.), 258-268. Routledge. https://doi-org.ezproxy.lib.bbk.ac.uk/10.4324/9781315180809.

O'Brian, L. (2024) 'I, myself, move', *European Journal of Philosophy'*, 1-14, DOI: 10.1111/ejop.12944

Oppenheim, G. and Dell, G. (2010) 'Motor movement matters: The flexible abstractness of inner speech', *Memory & Cognition*, 38 (8), 1147-1160.

O'Shaughnessy, B. (1973) 'Trying (as the Mental 'Pineal Gland')', *Journal of Philosophy*, 70, pp. 365-386.

Pacherie, E. (2007) 'The Sense of Control and the Sense of Agency', *Psyche*: volume 13 issue 1.

Pacherie, E. (2008) 'The phenomenology of action: A conceptual framework', *Cognition*, 107, 1:179-217.

Papineau, D. (2013) 'In the zone.' *Royal Institute of Philosophy Supplement*, 73, 175-196.

Pavese, C. (2022) 'Knowledge How', *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/fall2022/entries/knowledge-how/>.

Peacocke, A. (2023a) 'How to think several thoughts at once: plurality in mental action', in Brent M. & Miracchi Titus, L. (eds.), *Mental Action and the Conscious Mind.* New York, NY: Routledge. pp. 31-60.

Peacocke, A. (2023b) 'How to judge intentionally', *Philosophical Perspectives* 37 (1):330-339.

Pickering, M. & Garrod, S. (2013) 'An integrated theory of language production and comprehension'. *Behavioral and Brain Sciences* 36, 329–347.

Pinker, S. (1994) *The Language Instinct*, (New York: Morrow).

Plato. (1921) *Plato in Twelve Volumes*, Vol 12; translated by Harold N. Fowler. Cambridge, MA; Harvard University Press.

Quine, W. V. (1985) *From Stimulus to Science*. Cambridge, Mass: Harvard University Press.

Recanati, F. (2013) 'Content, Mood, and Force', *Philosophy Compass*, Wiley, 8, pp.622-632.

Rizzolatti, G., & Sinigaglia, C. (2008) *Mirrors in the brain, How our minds share actions and emotions*. Oxford: Oxford University Press.

Roessler, J. (2015) 'Thinking, Inner Speech, and Self-Awareness', *Review of Philosophy and Psychology*, 7 (3): 541-557.

Rothstein, S. (2004) *Structuring Events*. Oxford: Blackwell.

Rupert, R.D. (Manuscript) 'Naturalism Meets the Personal Level: How Mixed Modelling Flattens the Mind', available at: https://philpapers.org/rec/RUPNMT.

Ryle, G. (1949) *The concept of mind*. Chicago: The University of Chicago Press.

Sandis, C. (2010) 'Basic Actions and Individuation', in O'Connor T. and Sandis C. (eds.), *A Companion to the Philosophy of Action*, pp. 10-17. Oxford: Wiley-Blackwell.

Searle, J. R. (1969) *Speech acts: An essay in the philosophy of language*. Cambridge: Cambridge University Press.

Searle, J.R. (1976) 'A Classification of Illocutionary Acts', *Language in Society*, Apr., Vol. 5, No. 1, pp. 1-23.

Searle, J. (1983) *Intentionality*. Cambridge: Cambridge University Press.

Shepherd, J. (2015) 'Deciding as Intentional Action: Control over Decisions', *Australasian Journal of Philosophy*, 93:2, 335-351, DOI: 10.1080/00048402.2014.971035.

Shepherd, J. (2019) 'Skilled Action and the Double Life of Intention', *Philosophy and Phenomenological Research*, Vol. XCVIII No. 2, March.

Sinigaglia, C., & Butterfill, S. A. (2015) 'On a puzzle about relations between thought, experience and the motoric', *Synthese*, Vol. 192, 6, pp. 1923-1936.

Slobin, D. (1987) 'Thinking for Speaking', *Proceedings of the Thirteenth Annual Meeting of the Berkeley Linguistics Society*, pp. 435-445. DOI: https://doi.org/10.3765/bls.v13i0.1826

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

Stanley, J. (2011) 'Knowing (How),' *Noûs*,, 45:2, pp. 207-238.

Stanley, J. (2012) Replies to Dickie, Schroeder and Stalnaker. *Philosophy and Phenomenological Research, LXXXV* (3).

Stanley, J. & Krakauer, J. (2013) 'Motor skill depends on knowledge of facts', *Frontiers of Human Neuroscience*. DOI: 10.3389/fnhum.2013.0050.

Stanley, J. & Williamson, T. (2001) 'Knowing how', *Journal of Philosophy*, 98, pp. 411-444.

Steward, H. (1997) *The Ontology of Mind: Events, Processes, and States*. Oxford: Oxford University Press.

Steward, H. (2012a) 'Actions as Processes', *Philosophical Perspectives*, Vol. 26, Philosophy of Mind, pp. 373-388.

Steward, H. (2012b) *A metaphysics for freedom*. Oxford University Press.

Stout, R**. (**1997) 'Processes', *Philosophy*, 72(279): 19–27. DOI:https://doi.org/10.1037/13698-003

Strawson, G. (2003) 'Mental ballistics *or* the involuntariness of spontaneity', *Proceedings of the Aristotelian Society*, 103, pp. 227–56.

Todorov, E., & Jordan, M. I. (2002) 'Optimal feedback control as a theory of motor coordination', *Nature Neuroscience*, 5(11), 1226-1235.

Vendler, Z. (1957) 'Verbs and Times', *Philosophical Review,* 66: 143-60.

Vicente, A. (2022) 'Speaking for Thinking: 'Thinking for Speaking' Reconsidered' in Pablo Fossa (ed.), *Inner Speech, Culture & Education*. Springer.

Vicente, A. & Jorba, M. (2019) 'The Linguistic Determination of conscious thought contents.' *Noûs*, 53(3), 737–759. https://doi.org/10.1111/nous.12239.

Vygotsky, L. S. (1987) 'Thinking and speech.' In R.W. Rieber & A.S. Carton (eds.), *The Collected Works of L.S. Vygotsky, Volume 1: Problems of General Psychology* (pp. 39–285). New York: Plenum Press. (Original work published 1934.)

Wegner, D. (2003) 'The mind's best trick: how we experience conscious will', *TRENDS in Cognitive Sciences* Vol.7 No.2 February: 65-69.

Wilkinson, S. (2020) 'The Agentive Role of Inner Speech in Self-Knowledge', *Teorema: Revista Internacional de Filosofia* 39 (2): 7-26.

Wilkinson, S. and Fernyhough, C. (2018) 'When Inner Speech Misleads', in Llangland-Hassan, P. & Vicente, A. (eds), *Inner Speech: New Voices.* Oxford, Oxford University Press.

Williams, B. (2002) *Truth and Truthfulness*. Princeton, NJ: Princeton University Press.

Wilson, M. (2002) 'Six views of embodied cognition', *Psychonomic Bulletin & Review*, 9 (4), 625-636.

Wollheim, Richard (1999) *On the Emotions.* New Haven: Yale University Press.

Wu, W. (2016) ‘Experts and Deviants: The Story of Agentive Control’, *Philosophy and Phenomenological Research,* Vol. XCIII No. 1, July.