REVEALING THE LANGUAGE OF THOUGHT

A Thesis

submitted in partial fulfilment

of the requirements for the Degree

of

Master of Arts in Philosophy

by

Brent Silby

University of Canterbury

2000

Contents

Abstract.....2

Chapter 1: Introduction 3

Chapter 2: Thinking Sentences.....11

1. Preliminary Thoughts
2. The Language of Thought Hypothesis
3. The Map Alternative
4. Problems with Mentalese

Chapter 3: Installing New Technology: Natural Language and the Mind.....33

1. Introduction
2. Language... what's it for?
3. Natural Language as the Language of Thought
4. What can we make of the evidence?

Chapter 4: The Last Stand... Don't Replace The Old Code Yet.....61

1. The Fight for Mentalese
2. Pinker's Resistance
3. Pinker's Continued Resistance
4. A Concluding Thought about Thought

Chapter 5: A Direction for Future Thought.....88

1. The Review
2. The Conclusion
3. Expanding the mind beyond the confines of the biological brain

References.....96

Acknowledgments.....101

Abstract

Language of thought theories fall primarily into two views. The first view sees the language of thought as an innate language known as mentalese, which is hypothesized to operate at a level below conscious awareness while at the same time operating at a higher level than the neural events in the brain. The second view supposes that the language of thought is not innate. Rather, the language of thought is natural language. So, as an English speaker, my language of thought would be English.

My goal is to defend the second view. My methodology will see the project broken down into three major areas. First I will show that human thinking requires a language of thought, after which I will highlight some problems with assuming that this language is innate and hidden. Included in this section will be a small introduction to the compatibility problem. The compatibility problem offers some obvious difficulties for mentalese theories and these will be discussed. The next stage of the project will focus on evidence that can be put forward in support of the claim that natural language is the language of thought. Our most direct source of evidence comes from introspection, and this will play a dominant role in the discussion. The final part of the thesis will involve an examination of the principle arguments that have been put forward against the idea that natural language is the language of thought. My goal will be to show that these arguments do not entail the existence of mentalese, nor do they show that natural language is not the language of thought. I will provide answers to the arguments, and will explain the phenomena they point to in terms of natural language being the language of thought.

1

Introduction

Language of thought theorists claim that the mind is a device that operates according to strict rules concerning the manipulation of symbols. Such theorists claim that the mind is some sort of digital processor that runs on the highly parallel neural structure of the brain. This line of thinking draws an analogy between the mind and digital computation, and offers significant explanatory power. There is, however, some tension between two rival language of thought theories. One theory claims that the best candidate for the language of thought is an innate, behind-the-scenes language known as *mentalese*. The other theory states that a person's language of thought is their native natural language -- for example, English for English speakers, French for French speakers, or Japanese for Japanese speakers. Philosophers such as Jerry Fodor (1975) and Steven Pinker (1994) support the existence of mentalese. According to Pinker (1994), the common-sense view is that thought is independent of natural language. On this view natural language would not in itself shape the human mind in any fundamental way, although the internal mentalese thoughts being represented by the natural language sentences would. The alternative view, which can be found in the work of Sapir and Whorf on linguistic determinism, and in Wittgenstein's work on meaning and representation, suggests that our thoughts are constructed from sentences of natural language. Sapir said:

> *...thought may be no more conceivable, in its genesis and daily practice, without speech than is mathematical reasoning practicable without the*

lever of an appropriate mathematical symbolism (Sapir 1921: pg 14).

This passage shows the strength of Sapir's conviction that thoughts are dependent on natural language. The claim is not only that speech (natural language) shapes our thoughts, but that our thoughts themselves are not possible without natural language. Benjamin Whorf had similar beliefs. For Whorf, our experience of the world is one that has been organized by the linguistic systems at work in our minds. Through the use of language, the world is organized into the concepts and ideas that are prevalent throughout the linguistic community to which we belong. This line of thinking looks plausible. After all, it certainly *seems* as if our thinking consists largely in sentences of natural language. Whether we are thinking aloud or internally, the process almost always seems to involve a voice -- and this voice always talks in a natural language. Of course, merely *seeming* as if something is the case is not enough to show that something *is* the case. It is possible that the ideas of Sapir and Whorf are placing too much significance on the role that natural language plays in thought. Pinker's view may in fact be correct. Our thoughts may exist independently from natural language. After all, people who are born deaf and cannot speak can obviously think. And what about animals such as dogs? They cannot speak, yet some would claim that their behaviour demonstrates fairly well developed thought processes. Don't dogs have minds? In "The Language Instinct", Steven Pinker argues that his "common-sense" view of thought is correct. Indeed, he believes that he "can afford to be smug about common sense being true." Pinker claims that linguistic determinism is a "conventional absurdity" and that thought *is* different from natural language (Pinker 1994: pg 67).

I will argue that theorists such as Pinker have been too hasty in dismissing the crucial role natural language plays in thought processes. I will argue that we must draw a line between creatures that use natural language, and those that do not. If we can show that language users employ a different type of thought process, it will be possible to account for animal minds and the minds of humans who cannot speak, while still keeping the importance of natural language intact. In my view, natural language *is* essential for human thinking and for the structure of the human mind. The argument for this thesis is roughly as follows:

1. Human thinking exhibits productivity and systematicity

2. Productive and systematic thoughts require syntactic structure

3. The required syntactic structure can be explained by *either* mentalese *or* natural language

4. We should favour economy in our ontology (this is to say: we should avoid positing the existence of something that we cannot locate in the empirical world if the phenomenon in question can be explained by something that we can locate in the empirical world).

5. We cannot locate mentalese or the brain mechanism responsible for mentalese

6. We can locate natural language and there is good evidence to support its role in thinking

7. Therefore, natural language should be used to explain the syntactic manipulation required for producing productive and systematic thinking

8. Therefore, our thoughts should be explained by our use of natural language

This argument does not imply that animals and prelinguistic humans do not think. I am attempting to account for fully developed human thought, which seems to be far more sophisticated and flexible than the thoughts of animals. I think that natural language serves the purpose of restructuring the mind in such a way as to give humans a new way of thinking.

In chapter 2, I will introduce the mentalese hypothesis and will describe some important features of language. Part of my goal in chapter 2 will be to show how language (regardless of whether it is mentalese or natural language) can account for the *systematicity* and *productivity* of human thought. I will explain what is meant by these terms in chapter 2, but briefly, to say that thought is *productive* is to say that by applying certain rules to existing thoughts, a person can generate new thoughts. And to say that thought is *systematic* is to say that thoughts are composed of symbols that can be systematically rearranged in order to form different thoughts. So, if a person can entertain the thought 'Romeo loves Juliet', then by rearranging the constituent parts of that thought, the person is also capable of thinking 'Juliet loves Romeo'. Chapter 2 will also point to some problems with mentalese. These problems arise mainly because mentalese is supposed to be innate, behind-the-scenes and empirically difficult to find. For comparison, I will look at an alternative to the language of thought theories known as the map alternative, which views human thinking as involving a map-like system of representation. I will show that the map alternative is not as flexible as language of thought theories and conclude that thinking involves the manipulation of sentential structures. The required sentential manipulation can be explained either by mentalese or natural language, hence premise 3 above.

The most obvious way for a mentalese supporter to attack the above argument is to deny premise 3 and part of premise 6. They could say that natural language cannot be used to explain the syntactic structure required for thinking and that there is insufficient evidence to support its role in thought. I will show that the mentalese theorist cannot make this move. In chapter 3, I will provide evidence that gives us good reason to suppose that natural language is crucially involved in thinking. Most of this evidence is based upon introspection, and although the mentalese theorist may deny the reliability of such evidence, I will attempt to show that it should be taken seriously. When we introspect, we discover that our thinking is carried out in sentences of natural language, which manifest themselves as an internal monologue, or *inner speech*. Now, the mentalese supporter may claim that the inner speech we experience must be translated into mentalese in order for thinking to take place, but my claim is that such a view leads to an unnecessarily messy and complicated picture of the mind. Why should a mind have to devote so much time and resources to translation activities when thinking could be done in natural language? To back up the suggestion that thinking is done in natural language, I will consider the introspective testimony of Helen Keller, who could only communicate through a unique language of touch, and who claimed to think in terms of tactile sensations. The interesting feature of Keller's testimony is that she could remember her languageless existence, and claimed to be a non-thinking entity during that time. The implications of this claim are profound and point to the important role natural language plays in thought.

In the next section of chapter 3, I will look again at the productive and systematic nature of human thinking. I will provide evidence that should lead us to

believe that human thoughts do not become productive and systematic until natural language is acquired by the individual. There seems to be a specific stage in development at which time an infant acquires the ability to think symbolically. It is at this stage that the infant starts to assimilate natural language and starts to think in a distinctly human fashion. This will lead to chapter 4 in which I look at attempts to show that mentalese can explain certain phenomena of thought that natural language cannot explain. For example, a mentalese supporter may explain the familiar 'tip-of-the-tongue' sensation by saying that there is difficulty translating a mentalese thought into natural language. The idea here is that when you experience the phenomenon, you feel like you know what you want to say but cannot find the words to express the thought. The thought itself is perfectly formed and encoded in something other than natural language and that is why you feel like you know what you want to say. Mentalese theorists claim that if thinking is done in natural language, then there would not be any translation difficulties and so the tip-of-the-tongue phenomenon would not exist. I will show that the mentalese theorists are mistaken and that such phenomena can be explained in terms of natural language. It could, for example, be the case that we experience the tip-of-the-tongue sensation when we simply cannot form a thought. The feeling that we know what we want to say, but cannot find the words, may be the illusory result of simply not being able to figure out what we want to say.

Steven Pinker (1994) highlights several other reasons to suppose that thinking must be done in mentalese. These reasons include: 1) accounting for the fact that we often remember the 'gist' of a sentence rather than the sentence itself; 2) the fact that new words are coined to express new thoughts; and 3) the fact that many sentences of natural language can be ambiguous or logically inexplicit while, presumably, the

thoughts they represent are not. I will provide answers to these and other such suggestions after which I will turn my attention to the nature of Einstein's thoughts. Einstein claimed to think in the manipulation of visual imagery and mentalese supporters such as Pinker believe that his testimony refutes the claim that thinking is done in natural language. I accept that Einstein's thoughts involved visual imagery, but I think that those images must have been augmented by sentences of natural language. Sentences must have been attached to the images to provide information that cannot be captured visually, like information concerning temporal concepts such as 'tomorrow', or abstract concepts like 'perhaps'.

The final part of chapter 4 will deal with the evidence from aphasia. People who suffer from aphasia lose some (or all) of their natural language ability. But despite their loss of language, aphasia sufferers exhibit intelligent, distinctly human behaviour. This seems to show that human thought is possible without natural language, and may point to the existence of mentalese. I will look at the case of Brother John, a periodic aphasia sufferer, and will attempt to explain his behaviour in terms of natural language. I will attempt to show the possibility that his mind was operating on natural language structures below the level of consciousness. So, even though Brother John was unable to consciously comprehend or use language, his behaviour was still guided by natural language structures that existed below his level of awareness. To back up this claim, I will appeal to Chomsky's (1995) model of the human language system. Chomsky proposes that the language system is comprised of several self-contained modules. There are modules responsible for sentence comprehension and production, and there is a linguistic database that is built during the lifetime of an individual. In cases of aphasia, it is possible that the linguistic

comprehension and production modules have been disrupted, while the linguistic database is still operating behind the scenes, thus producing intelligent behaviour despite the supposed absence of language.

This thesis will not show that mentalese does not exist, but it will give us reason to conclude that natural language is a more suitable candidate for the language of thought. My conclusion will be that since mentalese and natural language each have significant explanatory power, the issue should be decided by looking for the simplest explanation. Mentalese is empirically difficult to find, and if it exists the brain must expend a significant amount of resources supporting both it and natural language. Because both these languages offer the same explanatory power, it is easier and cheaper to suppose that natural language is the language of thought. Once this possibility is accepted, researchers can devote more time to discovering the extent to which natural language shapes the mind and thought processes.

2

Thinking Sentences

1: Preliminary thoughts

Before embarking on a debate between *natural language* and *mentalese* as the language of thought, we have to be clear about the differences between these two languages. We also have to understand what it is about language that leads us to believe that it is required for human thinking. Clarifying these preliminary questions is one of the goals of this chapter.

I will begin this chapter by outlining the language of thought hypothesis, while at the same time describing the features of language that make it a useful tool for explaining human thought. One of the aims of this chapter will be to show that the distinctive features of human thought can best be explained by claiming that thinking involves the manipulation of sentences. I will show that sentential manipulation is vastly more flexible than other systems of representation. Another aim of this chapter will be to introduce *mentalese* as a possible language of thought. Mentalese, as I will explain, is hypothesized to be an innate language that operates behind-the-scenes, and as such, we are not consciously aware of its existence.

A further goal of this chapter will be to consider three problems with the idea that the language of thought is innate and exists behind-the-scenes. These problems relate to: **1)** The fact that we cannot locate mentalese, or the brain mechanisms responsible for its implementation, while at the same time we have natural language staring us in the face boasting its ability to account for the features of human thought

that we wish to explain. **2)** The problem of accounting for the meaning of mentalese sentences. **3)** The problem of accounting for the compatibility among the brain's distinct modules. Or, in other words, how is it possible that the brain, which evolved a piece at a time, came to be comprised of modules that share an ability to communicate to each other through a common language of thought? As I consider each of these problems, I will mention the possibility that they could be overcome if we consider natural language to be the language of thought.

My conclusion will be that thinking must involve operations on sentential structures, and that the problems associated with supposing that we posses an innate language of thought give us reason to think that natural language is the language of thought. This will provide a path to chapter 3 in which I offer evidence and defend the idea that natural language is the language of thought.

2: The Language of Thought Hypothesis

The language of thought hypothesis claims that when a person has a thought such as the thought 'grass is green', the content of that thought is represented in that person's mind by a sentence. However, according to Jerry Fodor (1975), this sentence is not a natural language sentence like English or Japanese. It is a sentence of an entirely different language -- the innate language of thought. The name often given to this language of thought is *mentalese.* An important feature of mentalese is that it is not a language that we have to learn through experience -- rather, we are born with it (Fodor 1975: pg 70). Furthermore, even though we are not directly aware of the language, it underlies all of our thought processes and has a similar structure to that of any of the natural languages. This is to say that mentalese, like English or French, is structured

according to certain rules of *syntax*, which determine how sentences are to be formed in order to give them a *semantic* (or meaningful) content.

The language of thought hypothesis draws an analogy between thought and computation. Mentalese is equivalent to the computational language of a digital computer, while higher level cognitive functions are achieved through the construction and manipulation of mentalese sentences. Within a digital computer, input received from the outside world (from a video camera say) is converted into strings of symbols, which represent the input data. These symbols constitute the computer's internal language and carry meaning when structured in certain ways. According to Fodor's language of thought hypothesis, the same is true of the mind. Input from the environment is converted into strings of mentalese symbols, which can then be operated upon by processes in the brain. An innate set of rules determines how sentences are to be structured, and how they are to be manipulated. So, the content of my thought 'grass is green' is written in my brain as a string of mentalese symbols -- or in other words, a mentalese sentence. Furthermore, the basic symbols that make up the mentalese sentence 'grass is green', are like the words of a natural language in that their meanings remain constant (Braddon-Mitchell and Jackson 1996: pg 164). The mentalese symbol that stands for 'grass' will always stand for 'grass', while the mentalese symbol that stands for 'green' will always stand for 'green'. Now, the symbols that comprise a computer's language of thought are implemented upon the physical on/off switches contained in its circuitry. The brain, on the other hand, does not simply contain on/off switches and as such, mentalese (if it exists) must somehow be implemented upon a different type of system. Discovering the

brain's symbol system is not a straight-forward task; but according to Fodor, one thing is certain -- the brain must contain, or at least instantiate a symbol system.

> *There must be mental symbols because ... only symbols have syntax, and ... the only available theory of mental processes ... needs the picture of the mind as a syntax-driven machine* (Fodor 1990a: pg 23).

Here, Fodor is claiming that mental processes must involve the manipulation of symbols. For Fodor, a rational mental life is best explained in terms of symbol manipulation according to a set of well defined rules. Alternatives, such as *associationism*, do not explain our ability to reason. Associationism is a habitual system rather than a rule following system. It sees our thoughts as arising because of their association with other thoughts or external stimuli. On an associationist account, my decision to drink a glass of water, say, may have been brought about by my having seen a picture of a hot desert, which reminded me of water. Fodor does not want to deny that a certain amount of association goes on in mental process, but he thinks that our thoughts are more often reasoned. Purely associative thinking would give rise to unpredictable human behaviour. For each thought there is a multitude of possible associated thoughts, which would give rise to disorganized thinking unless they were somehow brought under control by a rule following process.

Supposing that the mind operates by transforming symbols gives us a way to explain how our reasoning takes place. Fodor suggests that mental symbols have shapes that can be latched on to by mechanisms (or rules) that are sensitive to those shapes. These mechanisms facilitate the alteration of strings (or sentences) of the symbols giving rise to new strings that can be operated upon by other mechanisms. According to Fodor's view, a decision such as my decision to drink water, may have

came about because I was thirsty, and I chose water rather than coffee because I had already had several coffees and did not want to give myself insomnia. In short, the decision came about through reasoning. It is possible that the symbols involved in mentalese sentences are instantiated in the states of neurons and neural pathways that exist in the brain. If mentalese exists, we can suppose that future cognitive scientists will discover how the basic symbols (or words) that comprise mentalese sentences are implemented in the neural circuitry of the brain.

I have mentioned the fact that mentalese sentences must be structured according to rules. This is to say that mentalese sentences must have a *syntactic* structure, which determines how words are to be strung together in order to give a sentence a correct grammatical structure. We do not know what mentalese words look like, but we can understand the importance of syntax by considering sentences of natural language. For example, the English sentence 'The tree was tall and green,' is correctly structured, while the sentence 'Tall and tree the was green,' is not. Of course, syntax is not all that is required of mentalese sentences. Mentalese sentences, like natural language sentences, must also have a *semantic* content. That is to say they have to be meaningful. For example, consider the following English sentences...

The tree was tall and green.

And...

The noise was tall and green.

Now, although these two sentences are structured according to the same rules, the first one has a meaningful content while the second one is meaningless. This is to say that the two sentences have an identical syntactic structure, but only the first sentence has a semantic content. This is because one of the words used in the second sentence carried a meaning that did not belong in the context of that sentence. The use of the word 'noise' made the sentence meaningless.

An important feature of sentences is that they get their semantic content not only from the meanings of the words used to construct them, but from the syntactic rules used to connect the words together (Crane 1995: pg 139). The sentences 'Romeo loves both Juliet and his pet dog,' and 'Juliet loves both Romeo and his pet dog,' have different meanings despite the fact that they contain the same words. The differing meanings of these two sentences is brought about by the way in which their constituent words have been combined. This feature of sentences points to what Fodor calls the *systematicity* of language. In general, if I can understand the sentence **a** loves **b**, I will be able to understand the sentence **b** loves **a**. Both sentences are meaningful even though their constituent parts have been systematically rearranged.

Another feature of language is its *productivity.* According to Fodor, it is possible to create new sentences with different meanings by simply using a syntactic rule to add another word to an existing sentence. In fact, given the finite set of English words and the finite set of English syntactical rules, I could (in principle) generate an infinite number of meaningful sentences. For example, I could form the sentence:

I think that Romeo loves both Juliet and his pet dog.

Then I might form the sentence:

> John knows that I think that Romeo loves both Juliet and his pet dog.

I might decide that these statements are not worth worrying about:

> Nobody cares that John knows that I think that Romeo loves both Juliet and his pet dog.

It is possible to continue this process indefinitely. In principle, any number of sentences can be generated and understood, as long as they are built from words that are part of the language being used, and are constructed according to the syntactical rules of that language. The point here is that if the language of thought hypothesis is correct, it will explain the diversity of human thought and the human capacity to create new thoughts. The language of thought theorist can correctly claim that sentential representation is a form of expression and representation that is more versatile than rival systems of representation. To clarify this point, consider a different sort of representational system that consists entirely of pictures such as road signs. One sign may contain a picture of a car sliding across the road and carry the meaning: WARNING! ICE ON ROAD. Another sign may contain a picture of two children, which means: BE AWARE! CHILDREN CROSS HERE. Now, with no syntactical rules, it would be impossible to combine these two signs in such a way as to carry the meaning: CHILDREN CROSS ON ICE, or BE AWARE! ICE ON

CHILDREN (see Crane 1995: pg 140 for a similar example). A mind that contained a 'road sign' system of representation like this would not exhibit the productivity and systematicity of thought that our minds exhibit. Such a mind would be rigid and inflexible, and could only entertain thoughts with very specific contents. In other words, the mind could only think thoughts whose contents were represented by various pictures in a strict system of signs. On the other hand, a mind that operates according to a sentential system of representation has the ability to construct novel thoughts. The rules of syntax give minds the ability to generate a potentially infinite number of thoughts by directing the combination and recombination of a finite set of words.

Postulating the existence of a language of thought can also help us to explain the apparent way in which thought *evolves* through the process of *reasoning*. For example, I may make the following inference about black holes:

If black holes exist, then Roger Penrose was correct.

Black holes exist.

Therefore: Roger Penrose was correct.

Fodor believes that many of our thoughts evolve through this kind of systematic reasoning and compares the process to the inferences that occur within computational devices (Fodor 1990a: pg 22). The idea here is that the symbols contained in the first two sentences are operated upon by a formal rule that gives rise to the third sentence. In logic, this rule is known as *modus ponens* and takes the form:

If P, then Q.

P.

Therefore: Q.

According to the language of thought theorist, our thought processes can best be explained by the existence of this sort of process; and furthermore, the fact that this process is successfully employed within digital computers proves beyond doubt that this type of system works.

It is true that the language of thought hypothesis can explain the reasoning that occurs in thought. It is also true that the productivity and systematicity of human thought can be accounted for by supposing that thinking consists in operations upon sentential structures. But before continuing, I will consider an alternative to sentential theories of thought. The so-called 'map story' is a plausible theory that has the capacity to explain the features of thought that we are interested in.

3: The Map Alternative

Sentential structures offer us a tidy, effective way of explaining thought but we should remember that there are many other structured systems of representation. Some of these systems do not have the scope to explain the nature of human thought (like the road sign example earlier). On the other hand, the 'map-story' is a rival explanation of thought that offers a challenge to sentential theories. According to the map-alternative, the roles that would have been played by mentalese sentences are played by mental maps. Of course, this is not to say that there literally are maps drawn in the brain; it is simply to claim that mental states represent in much the same way as maps

represent, rather than the way sentences represent. Map-like representations can contain a lot of information, and they can be used to represent any feature of the empirical world (Braddon-Mitchell and Jackson 1996: pg 172). The way maps represent is different to the way sentences represent. By using a sentence, I could represent a very specific piece of information such as 'Auckland city is in the North Island of New Zealand'. But with a map, I could not isolate such a small piece of information. Consider geographical maps. I could not look at a map of New Zealand and extract information about the location of Auckland without simultaneously extracting other information related to Auckland. When I look at the location of Auckland, not only do I discover its position relative to other cities, I am also instantly aware of its approximate size, its shape, the fact that it sits above sea-level, and so on and so forth. All this information is connected and present in a very small portion of a geographical map and is instantly available if required.

When we consider the amount of information that can be contained within a two-dimensional map, we can imagine that adding further dimensions to the map will significantly increase the information carried. For example, a four dimensional map with a sufficient resolution would offer information that connects events from the past to the present, and to possible future events. I should note here that mental maps should not be considered to be as accurate as geographic maps. As Armstrong (1973: pg 3) points out; mental maps will contain errors, fantasies, contradictions and many blank spaces. This is because our mental maps are formed as we move through the world during our life time. Presumably, as we encounter novelty and gain experience our mental maps grow in sophistication but continue to represent erroneous data until they are corrected.

The map alternative is not as intuitive as the language of thought hypothesis. This could be because it is difficult to see how map-like structures can be used to represent all the features of our experience. However, just because something is counter-intuitive does not mean that it should be dismissed. The map alternative can, in fact, offer us an explanation for the versatility of human thought. For example, the potential for a human to produce an unlimited number of distinct thoughts can be explained by the map story if we suppose that there is no limit to the number of possible arrangements of elements within the map. This could be achieved if the map wrapped around upon itself like a map of Earth printed on a globe. The map theory would then allow for unlimited thoughts in this way: I could scan the surface of the global map one time and see Auckland, Brisbane, and then Perth. Alternatively, I could scan the surface a couple of times and see Auckland, Perth, Auckland (again), and then Brisbane. Given enough time, I could scan the globe an infinite number of times and see these cities in an infinite number of combinations. Furthermore, I could also carry out any one of these scans in reverse.

The map story can also be used to account for the way in which thoughts seem to evolve over time. This process can occur by blending two maps together, or by updating an existing map as new information arrives. A useful way to imagine this process is to imagine superimposing a meteorological map over the top of a geographic map, and then continually updating the meteorological map as new weather information arrives. A map that behaves in this way would evolve according to new data and a set of rules that defines how the new data should be incorporated into the existing map. Map **M2** at time **T2** is a product that has arisen from map **M1** + input data **D1** at time **T1** according to updating rules. Since this process is

hypothesized to occur according to specific formal rules, we can use it as a possible explanation for our ability to reason. We could suppose that there are rules that correspond to the formal rules of logic that we use to operate on sentential structures. So in a map-like system of thought, there is a rule that is analogous to modus ponens and is responsible for the appropriate transformations of mental maps.

There are some difficulties with the map story of representation. Consider, for instance, the way in which the elements of a map are all connected and related. Because of the interconnectedness of information contained in a map, it seems that we cannot reduce the complex meaning of a map to simpler elements with simpler meanings. Consider a very small map such as a map of my office desk. This contains, among other things, information concerning the number of objects on the desk, and the amount of free space on the desk. But this information comes along together and cannot be separated into its simpler elements. I cannot isolate information about the number of objects on my desk without simultaneously representing information about how much free space there is on my desk. It's all tied together. With sentences, however, I can think about the number of objects on my desk without thinking about how much free space there is. This seems to be more like the way we think.

Another problem with the map story is that it is difficult to see how a map-like system could represent complex thoughts, or thoughts that contain a negation. Consider, for example, the English sentence:

It is not raining.

How could one represent this sentence as a map? A weather map may depict the fact that it is not raining by having no 'rain' pictures over the city, but this method would not capture the meaning of the English sentence 'It is not raining'. The map represents the absence of rain by the absence of a 'rain' symbol, but to carry the meaning 'It is not raining', the map must somehow represent 'rain' and attach a negation to that representation. I think that this could only be achieved through the combination of a map and a sentence. A map that does not contain any 'rain' symbols is not representing the proposition that 'It is not raining'; it is merely representing a city with no extra features. Nothing is included that tells us anything about rain.

The major objection to the map theory is that it fails to account for our systems of belief and the fact that we often fail to connect our beliefs together correctly. Consider the following example: Jane believes that Fred has two children but does not have the belief that the number of Fred's children is the smallest prime (Braddon-Mitchell and Jackson 1996: pg 191). According to the map story, however, all the information relating to the number two is connected, so Jane *should* believe that the number of Fred's children is the smallest prime. Because of the connectedness of information in maps, the map-story requires that if Jane believes that **P**, and if **P** implies **Q**, then Jane believes that **Q**. Sentential systems of thought, on the other hand, do not have this requirement and can account for Jane's not having the belief that the number of Fred's children is the smallest prime. The sentence 'Fred has two children' can be written in Jane's head without the sentence 'The number of Fred's children is the smallest prime' also being in there (Braddon-Mitchell and Jackson 1996: pg 193). Of course, with all the appropriate information, Jane could easily generate a sentence, and hence a belief, about the number of Fred's children being the

smallest prime, but this does not have to happen. Here's another example: Suppose Jane believes that Fred lives in Christchurch city. Now, it happens to be the case that if Fred lives in Christchurch he will have to travel north to visit Auckland city. If Jane represents her beliefs in a map-like way she should believe that Fred must travel north to go to Auckland. This is because all the appropriate information is contained in a mental map, and is available simultaneously. But clearly this is not the way humans form beliefs. It is possible that Jane could believe that Fred lives in Christchurch and even believe that Auckland is north of Christchurch without ever forming the belief that Fred must travel north to get to Auckland. Her beliefs offer the *possibility* of making that inference, but the inference may never be made.

Computer scientists are experimenting with systems of representations that are more like maps than sentences (Braddon-Mitchell and Jackson 1996: pg 173), and I think that map systems offer promising new ways of designing computers -- especially robots that need to navigate around changing environments. However, it seems to be the case that certain aspects of human thinking cannot be adequately explained by using map theories. Sentential operations still offer the best way of accounting for all the features of human thought.

Having identified the plausibility of there being a language of thought, I will now point to some problems with the idea that the language of thought operates behind-the-scenes and is innate. In chapter 3, I will claim that the ability to assimilate and use natural language is all that the human mind requires to produce complex thoughts. After all, natural language contains sentential structures and operates in much the same way as mentalese is hypothesized to do so.

4: Problems with Mentalese

(1) *Where is the language of thought?*

The first problem I shall address relates to the notion of syntax. This problem is quite simple and can be expressed by one question. Where are the mentalese words? When we think in sentences of natural language we can point to words and see very clearly how these words connect to other words in order to yield meaningful sentences. Mentalese words, on the other hand, are not so easy to track down. For a start, they cannot be discovered through introspection because mentalese is a 'behind the scenes' language of thought. When we introspect, the only language we are aware of is natural language. The constant chatter of our internal monologue is with us virtually all the time, and while we can certainly identify syntax within its sentence structures, this internal speech is as removed from mentalese as the English words on my computer screen are removed from the computer's internal programming language. So, we cannot look for mentalese through introspection. Of course, this fact should not bother a mentalese supporter. Mentalese, like many of the brain's functions, is hidden from conscious awareness and the fact that we are only consciously aware of natural language does not have any impact on the existence (or not) of mentalese. In keeping with the computer analogy, we have to remember that we cannot look at a word processor display screen in order to discover a computer's internal language of thought. But, on the other hand, we *can* look *inside* a computer's hardware to discover its binary language of thought. We can do this by examining the multitude of physical on/off switches contained within its circuitry. A digital computer's circuitry is a structured system that has regularities. In other words, it is possible to point to its internal states and identify the way in which syntax is implemented in

those states. The same does not seem to be true of the brain. The brain is a messy collection of interconnected neurons that do not appear to have a syntactic structure at all. The point I am trying to establish here is that so far it has not been possible to identify a mechanism by which mentalese syntax is implemented in the low level neural structure of the brain. It is worth noting here that a mentalese supporter could reply to this objection by pointing out that natural language cannot be specifically located in the neural structure of the brain, but that it clearly exists. Their point would be that we should not dismiss the existence of mentalese just because neuroscience cannot fully describe its implementation in the brain. I think the mentalese supporter would be right in making this comment. Mentalese, after all, is supposed to be a high level language which is an abstraction away from the neurological events that underpin its workings, while at the same time remaining somewhere at the subpersonal level. So the question as to whether a mentalese implementation mechanism exists should remain open until some evidence can be found to give us a firm answer.

However, in the next chapter I will suggest that since we have evidence for the existence of natural language, and since natural language can be used to explain our thinking, it may not be necessary to evoke an 'invisible' mentalese. When we consider our introspective experience, it certainly seems to be the case that our thinking is being done in natural language. It could be that natural language offers us the simplest explanation of human thought. In light of this possibility I think we should postpone our search for mentalese and concentrate on discovering the way in which natural language is implemented in the brain.

(2) *Where does the meaning come from?*

Let us suppose for now that syntactic structures *do* exist in the brain. The next question we need to ask relates to the meaning of mentalese sentences. Unfortunately this question leads to a problem. When we ask where the meaning of a mentalese sentence comes from, we find that explanations are either regressive or head in an unnecessary direction. This is to say that on the one hand, an explanation of where mentalese gets its meaning might appeal to some other language that itself needs to be explained, or on the other hand, any non-regressive explanation of where mentalese gets its meaning could also work for natural language, thus making mentalese unnecessary. This is an important problem, which requires some exploration. Innate language of thought theorists may claim that sentences of natural language get their meaning from the mentalese sentences that they were translated from, but this still leaves open the question as to where the mentalese sentences get their meaning (Dennett 1996: pg 51). Some theorists have claimed that natural language sentences get their meanings from the way in which they are consciously used and understood by individuals (Crane 1995: pg 149). But this idea cannot explain where mentalese sentences get their meanings. If we were using internal mentalese sentences, and were conscious of the way in which those sentences were being used, it would seem that we would have an awareness of mentalese. But mentalese is hypothesized to be a behind-the-scenes operational language of which we are not aware. If we could become aware of mentalese through introspection, there would be no need to question its existence.

Perhaps mentalese sentences get their meaning from the way they are used and understood by something else -- something deeper. We could suggest that spoken or written sentences of natural language get their meanings from the way in which we use them, and the mentalese sentences that those thoughts were originally formed in have their meanings in virtue of being used and understood by some other system deep within the brain. This other system may operate according to its own internal language; say 'sub-mentalese'. But this leads to a problem. That is, we must now discover where the meaning of the sub-mentalese sentences lies. What part of the brain understands sub-mentalese? It would be pointless to appeal to yet another system that uses its own inner language such as 'deep-mentalese', because that would simply push the problem down to another level. Continuing this process would lead to a regress of systems and languages, and would not answer any questions. If 'deep-mentalese' has meaning, then why can it not be the case that mentalese has meaning (Sterelny 1990: pg 32). Why can we not simply state that mentalese *is* the most basic language, and that a multitude of mentalese sentences constitutes thought and higher-level meaning? Consider the way in which Fodor explains how mentalese symbols are used. Fodor supposes that the 'shape' of mentalese symbols determines how they might causally interact with each other (Fodor 1987: pg 18). For Fodor, the mind is a computational device and its syntactic operations can be explained in the same way that a computer's internal operations should be understood. Such operations consist in the transformations of strings of symbols. The processes involved in manipulating symbols are concerned only with the 'shape' of the symbols, not the meaning. In effect, these processes can be understood in much the same way as we understand how the shape of a key determines which lock it will open (Fodor 1987: pg 19). A

mentalese symbol may only be operated upon by processes that can 'latch on' to its shape. According to this account there is no question as to what it is in the brain that understands mentalese, or where the meaning of mentalese sentences come from. Instead we view mental operation as a brute mechanical process in which the problem of meaning does not exist. Viewing mental operation in this way can get us passed the problem of meaning, but it leaves open the following question: if Fodor is right, and if it is possible for us to understand meaning and thought in terms of mentalese symbols being operated upon by processes that can lock on to their shapes, then why can we not simply explain thought in terms of natural language? Natural language contains symbols that can be manipulated by mechanical processes, and if it is possible to explain thought and meaning in terms of brute mechanical symbol manipulation, then natural language should offer as much explanatory power as mentalese.

(3) *The compatibility problem.*

When engineers design a computer system, they typically have a completed end product in mind. This is to say that the computer is designed to meet certain specifications, and once it has been successfully built, it will carry out its specified tasks with no intervention from the design team. Of course, the design team will usually include a facility by which new hardware can be easily added to the system just in-case future tasks require additional tools. This new hardware will be designed by engineers who have knowledge of the original system design so that the additional components can be integrated into the system without the need of changing the original design specifications. Computer engineers will ensure that the original

system and its peripheral add-ons all share the ability to communicate with each other through a common language. In other words, the hardware will be compatible. This engineering strategy works very well for computers, but the brain was not designed in this way. In fact the brain was not *designed* at all and it certainly did not appear as a complete unit. According to theorists such as Cosmides (1989), the modern human brain had its origins as a much smaller device, which was a collection of self contained modules, each responsible for its own specific task. For example, there may have been a module that was dedicated to regulating cardiovascular activity, while another module may have been dedicated to processing visual stimuli. These modules were only able to carry out their specified tasks and could not be utilized for any new tasks. As time passed, the process of natural selection gave rise to new 'purpose built' modules that were added to the existing brain. These add-ons served new purposes and may have had the ability to communicate information to other parts of the brain.

Now, if this modular theory is correct we find ourselves with a problem. If the brain operates according to an internal language of thought such as mentalese, it would seem that each of its modules would need to operate according to the same language in order to be compatible with the rest of the system. But these modules appeared through the process of evolution by natural selection and as such were not designed in order to be easily 'plugged in' to an existing system. The process that gave rise to new modules is a process that does not move towards end goals. It is not possible that evolution by natural selection could have considered an integrated brain or a pre-existing language of thought. So, how can it be that all these distinct modules are compatible with each other? I call this the *compatibility problem,* and I believe

that it gives us reason to consider alternatives to the innate language of thought theory. Of course, it is possible that evolution gave rise to modules which all happened (by chance) to operate according to the same language of thought. Or, at the very least, evolution may have given rise to modules that had the ability to 'learn' and become compatible with the existing system. Natural selection would, of course, favour these modules. However, the probability of this occurring is infinitesimal. But then again, the probability of life arising on Earth is itself infinitesimal, and yet it happened. The compatibility problem, then, does not seriously challenge the possibility of there being an innate language of thought. It does, however, give us sufficient reason to keep our options open and search for a better theory that does not suffer the same problem. Such a theory might suggest that the human brain did not acquire a language of thought until it had reached a certain stage of evolutionary development -- perhaps when it became a modern human brain. We could claim that once nature had completed the construction of the brain, a global language of thought was installed. Once installed, this language provided a new way of thinking and gave rise to a rich and flexible mind. In the next chapter, I will show that this language was not necessarily mentalese. It could simply have been a natural language.

5: Where to next?

I have shown that the innate language of thought hypothesis has difficulties that are not easily solved. Such difficulties occur when we attempt to locate the syntax and semantics of thought. There seems to be no way of locating mentalese (or its implementation) in the brain. On top of those difficulties, we have the question of

how the brain, which evolved gradually over millions of years, could end up with a system-wide, innate operational language of thought.

Despite the difficulties with mentalese, I believe that Fodor is correct in claiming that the features of thought can best be explained by viewing thinking as sentence manipulation. Thinking is productive, systematic, creative, and evolves through processes of inference and reasoning. Possible alternative systems of thought and representation (such as the road sign system I mentioned earlier) do not have these features. Minds that use such systems would be inflexible and non-creative. The map alternative offers a way to account for the productive, systematic and creative aspects of thinking but cannot account for other features such as simple negation; nor can it account for the fact that we often fail to connect our beliefs correctly.

Having agreed with Fodor's claim that our thinking must involve sentential operations I will, in the next chapter, consider the possibility that natural language, rather than mentalese, is the language of thought. After examining evidence for this possibility, I will move on, in chapter 4, to discuss objections and arguments that have been put forward against the claim that natural language is the language of thought.

3

Installing New Technology:

Natural Language and The Mind

1: Introduction

In Chapter 2, I described the language of thought hypothesis. I also exposed some of the problems with supposing that the language of thought is innate and hidden. These problems do not show that there cannot exist an innate language of thought such as mentalese, but they do weaken its position.

One of the aims of chapter 2 was to establish some of the features of language. Whether we are talking about mentalese, machine-code, Japanese, or English, the structural features of language are very important and I agree with Fodor that these features *can* be used to account for the flexibility of the human mind. As I stated in chapter 2, I do not believe in the existence of mentalese but I do believe that there is a language of thought. In this chapter, I will propose that natural language is more than just a communication device, and that its implementation on the connectionist network of the brain allows it to serve the function of the language of thought. On this account, a thought is not formed by some other means and then translated into natural language to be spoken externally (or internally) -- rather, the thought is formed *in* natural language. In other words, natural language is the medium in which human thinking occurs. Of course, there will be resistance to this idea. In defending the existence of mentalese, Steven Pinker (1994) argues that thinking is not done in natural language -- rather, thinking requires an innate language of thought. For

Pinker, there is sufficient evidence to show that thinking does occur in the absence of natural language. Such evidence includes examples of deaf, languageless people who behave as if they think; examples of thought in pre-linguistic humans; and evidence that monkeys think. For example, Pinker points to the case of Ildefonso, a deaf Mexican immigrant who seemingly has no language. Pinker cites Susan Schaller's (1991) observations of Ildefonso and believes that there is evidence to show that Ildefonso thinks despite his absence of natural language. Ildefonso exhibits seemingly intelligent behaviour and is curious about the activities of people. He has managed to work regularly pruning trees and picking crops on a farm, and is eager to describe his daily activities through mime. Pinker thinks the existence of this sort of case points to the existence of an innate language of thought, which provides the ability to think even in the absence of natural language. However, in his haste to make this point, Pinker overlooks the fact that people such as Ildefonso usually mix with other people and spontaneously develop a system of signs with which to communicate[1]. So, we cannot be entirely certain that people like Ildefonso do, in fact, lack language. Their systems of signs may not conform to any of the standard sign languages, but they seem to be languages nonetheless. These languages involve a consistent syntax and can be learned by people who are exposed to them.

Interestingly, Fodor -- who believes in the existence of mentalese -- acknowledges the importance natural language plays in shaping our world view.

> *...I am not committed to asserting that an articulate organism has no cognitive advantage over an inarticulate one. Nor, for that matter is there any need to deny the Whorfian point that the kinds of concepts one has may be profoundly*

[1] This point was made by Oliver Sacks in his BBC television documentary "The Mind Traveller" and there is evidence to support the existence of this system of signs in the case of Ildefonso.

> *determined by the character of the natural language that one speaks.* (Fodor 1975: pg 85.)

This is a nice way of expressing the importance of natural language. Unfortunately, Fodor's theories require that the concepts determined by natural language be translated into mentalese in order to be integrated into the subject's thoughts. I do not think that any such translation takes place.

Showing that natural language is the language of thought is a big topic and will have to span two chapters. The goal of this chapter is to establish that thinking is done in natural language. To start, I will spend some time highlighting a distinction between the two roles that language plays -- its role in communication and its role in thinking. This will be followed by a three part defense of the claim that thinking is done in natural language. First, I will consider introspection, which provides us with very strong evidence in support of the claim that natural language is crucial for thought. Second, I will consider evidence from cases of psychological development in the absence of language. In particular, I will focus on Helen Keller, whose disability meant that she could not acquire language until she was about seven years old. The third part of my defense will appeal to the nature of human thought. Evidence shows that human thought has features that animal thoughts do not have. Furthermore, these features can only be explained in terms of language. The work in this chapter will lead to chapter 4, in which I elucidate and examine arguments that have been put forward to show that thinking is independent of natural language.

2: Language... what's it for?

Peter Carruthers (1996) categorizes two distinct roles that language plays -- the *communicative* role and the *cognitive* role. This is a useful distinction because it allows us to isolate the specific function of language that we wish to talk about. Carruthers believes that the most important role language plays is its cognitive role -- that is, the function it plays in thought. As I will make clear in this chapter, I agree with Carruthers that language plays an important role in thinking, but at the same time, I do not think the cognitive role could exist without the communicative role. This is to say that language cannot serve its function as a cognitive device without simultaneously serving the function of a communicative device. I will say more about this later, but first I will briefly describe these two conceptions of the function of language.

The Communicative Role.

This is the role that language plays in communication. To say that language has a communicative function is to say that it has an interpersonal use. Language is used for the transmission and reception of information among individuals in the language using community. This activity allows groups of language users to exchange thoughts and work collectively. It is important, however, not to confuse the complex communication that language enables with the basic signaling systems that are in use by other animals. Animals have been communicating simple messages to each other for millions of years but natural language is quite a recent development. It emerged sometime within the last 2 million years. Before this time, the vocalizations of the primates were nothing more than simple signaling devices, which were used to mark

territory, warn others of danger, attract mates, and so on. There was no flexibility in the cries of the early hominids, and as such, novel events could not be identified by new vocalizations or by constructing combinations of the existing signals. Michael Corballis (1999) conjectures that the remnants of our ancestral signaling systems still exist in our behaviour today. For Corballis, laughing, screaming and crying are some examples of signals that may once have been used to convey very specific information.

Natural language is a more complex system of communication that may have had its origins in something other than these primitive signaling systems. One popular view claims that language had its origins in manual gestures and it was only very recently that language made the transition to a vocal system. This is not a new idea. The French philosopher Etienne Condillac made this suggestion as far back as the 17th-century (Corballis 1999), and even though there is little hard evidence to back up this claim, it has continued to gain support. When we observe the behaviour of humans, we can clearly see the way in which gesture and language are tied together. Furthermore, neurological scans reveal that speech and gesture both originate from the same place in the brain. William Calvin (1998: pg 48) draws our attention to a 'sequencing' system that appears to be common to both hand movements and language. This sequencer seems to be centered around the left lateral brain -- an area that is thought to be crucial to language -- and damage to this area not only gives rise to a condition known as aphasia, which affects language, it also affects a persons ability to execute novel sequences of hand and arm movements. Merlin Donald (1991) has suggested that early communication involved the very complex gestural combinations involved in mime. Following this premise, Corballis (1999)

claims that mime was the precursor of spoken language. This is plausible. Acting out scenarios through mime is a very powerful way of communicating complex ideas, and it still seems to exist in our behaviour today. People often (unconsciously) accompany their spoken sentences with elaborate hand gestures. Presumably, as our ancestors became skilled at mime, their movements began to be accompanied by vocal sounds. Eventually the vocal sounds came to replace the movements of the mime, giving humans the ability to transmit thoughts over larger distances or in the dark, while also freeing up their hands for other tasks.

Regardless of its origins, the role that language plays in communication is obvious. But is communication all that natural language is used for? Is there any reason to suppose that natural language also plays a role in thinking? If no evidence can be found to support the claim that natural language plays such a role, we will have to consider the possibility that natural language is only used for the communication of pre-formed thoughts. Theorists such as Fodor and Pinker are committed to this position. On this account, we end up with a picture that has our thoughts being produced by some internal language (perhaps mentalese) before being translated into natural language for communication. One of the obvious problems with this view (which I will address later) is the fact that many humans talk out-loud to themselves. Why communicate a thought if there is no-one around to communicate it to?

The Cognitive Role.

Philosophers such as Wittgenstein have suggested a cognitive role for language. This is to say that language is far more than a mere communication device. On this account, sentences of natural language are not external encodings of hidden internal

thoughts. Instead, sentences of natural language actually *are* the thoughts. It is important to note that this claim encompasses inner speech, external speech, and even written sentences. Carruthers (1996) suggests that when we write we are carrying out much the same activity as when we speak aloud or think internally. We are not forming thoughts and then translating those thoughts into natural language to be spoken or written -- rather, *the actual thoughts* are being spoken or written. When we are engaged in a stream of thought, we can quite easily switch between an external and internal mode of operation. The reason this is important is because it shows that there is not much difference between entertaining a thought in inner speech or external speech -- it is always done in natural language. As Wittgenstein (1958: pg 43) put it, "the experience of thinking *may* be just the experience of saying". So, human thinking and reasoning consists in sequences of natural language sentences that are spoken internally or externally. Communication does not involve the translation and transmission of hidden internal thoughts; it involves the *direct communication of the actual thoughts.* Users of the same language will understand transmitted thoughts in virtue of the fact that for them, the utterances constitute virtually the same thought (Carruthers 1996: pg 2).

I stated above that both the communicative and cognitive functions of language are important. Moreover, I do not think that one can exist without the other. If language first emerged as a system of communication, the restructuring that occurred in the minds of those who were exposed to language may have eventuated in their thoughts being carried out according to linguistic rules. In "Consciousness Explained", Daniel Dennett proposes a mechanism by which language's initial communicative function could have given rise to a cognitive function. Dennett claims

that language may have facilitated the communication between previously unconnected modules of the brain thus enabling more complex thought processes. This effect occurred when our ancestors accidentally started communicating to each other through primitive grunts. If an individual needed help making a stone knife, he or she would have uttered a sound signifying that assistance was required. The sound may have been heard by someone else and, upon entering their brain, triggered a 'knife making' mechanism, which caused that person to carry out knife making activities. If no-one else was around, the cry for help would only have been heard by the individual who uttered it. The sound would enter that human's own brain and trigger its own 'knife making' mechanism. Dennett suggests that as the ages passed, a neural link emerged that connected the vocal production center directly to the auditory system. This gave rise to the inner speech that we all experience, and was responsible for the exponential increase in human cognitive capacities.

Peter Carruthers (1996: pg 192) thinks that Dennett's proposal does not accurately capture the phenomenology of inner speech. He thinks that if our inner speech were being used solely for the communication between the brain's otherwise unconnected modules, our inner speech would simply consist in something like a continual question and answer dialogue. Clearly, says Carruthers, this is not the case. Despite this objection, I think Dennett's suggestion is plausible. We often find our inner speech slipping out into the world when we think out loud. This occurs especially when we are carrying out complex tasks, or thinking through difficult problems. We guide ourselves through certain tasks by means of an external (but more often internal) monologue. My answer to Carruthers' objection is that questions are not all that is required to illicit responses from other modules in the brain. A

single word can trigger a multitude of comments from within the brain giving rise to a rich internal monologue.

Another possibility that can connect the communicative and cognitive roles of language takes an entirely different approach to the origin of language. Language's obvious role in communication may lead us to believe that its first, and primary use was that of a communicative tool. But what if language had its origins as a cognitive device? Consider the idea that language had its beginnings in gesture. Complex gestures gave rise to mimes, by which individuals could re-enact events, or plan new events. We assume that these mimes were intended for an audience, but what if it were the case that the mimes were not used in this way? I think it is possible that our ancestors were not miming for an audience at all -- rather, they were *rehearsing* complex physical events for their own benefit. After all, practice makes perfect. They may also have replayed events for themselves in order to iron out the bugs in their day to day activities. It could have been nothing more than an accident that other members of the group saw these rehearsals and recognized something that they themselves might do. Perhaps new hunting techniques were passed on in this manner. So, it is possible that a system that was *initially* beneficial only to individuals *became* a way of communicating ideas to others.

3: Natural Language as the Language of Thought -- Where's the Proof?

The rest of this chapter will focus on the role that language plays in cognition. I have already suggested that much of our thinking is carried out in natural language. Moreover, I think that for modern humans, one of the crucial roles of natural language is the role it plays as a medium for thinking. I will now launch a defense of this idea.

(1) *Introspective evidence*

The most immediate evidence that suggests a role for natural language in thinking comes from introspection. When we turn our attention inwards, we become aware that our thoughts are being formed in sentences of natural language, which manifest themselves as an internal monologue, or *inner speech.* This inner speech may consist in hypothetical conversations with people, or replays of past conversations. It can be used to plan future events or solve everyday problems. Consider, for example, the role inner speech plays in the following situation: The other day I arrived at my office and discovered that I had run out of coffee. Needless to say I was panic-stricken, but suddenly it occurred to me that I had put some new coffee in my bag the previous night, and I had brought my bag with me to the office. My stream of thought took the form of natural language sentences which went something like this "...oh no, there's no coffee. What will I do? Where will I get some? Damn! What a hassle, I'll have to get some. **But I put some in my bag last night**...". The final part of this thought was accompanied by a visual image of me placing coffee in my bag. It is difficult to determine whether the sentence caused the image, or whether the image caused the sentence; but I think that since the first part of the monologue was entirely vocal (no images), it is reasonable to assume that language played a dominant role in the thought process. Carruthers (1996: pg 36) suggests the possibility that images are involved in thinking. He does, however, state that the images must be augmented by sentences of natural language. His reason for this is that sometimes images do not carry enough information and sometimes they carry too much information, and that information needs to somehow be isolated. Consider, for example, the possibility that

in my stream of thought, the image of me placing coffee in my bag was *not* accompanied by a sentence. If this was the case, I would have had no way of knowing whether that image related temporally to an event that took place the previous night, or a few weeks ago. It was the sentence that completed the thought and gave it meaning. I will say more about this in chapter 4 when I discuss the nature of Einstein's thoughts.

Sometimes our inner speech slips out into the environment when we think aloud. We see this in children who usually accompany their play with a spoken monologue. Observations of children provides evidence that speaking is not always for communication. Young children engage in what Jean Piaget (1932) calls *egocentric* speech.

> *Egocentric functions are the more immature functions, and tend to dominate the verbal productions of children 3-7 years of age, and, to a lesser extent, children 7-12 years. In this form of speech, a child does not bother to know to whom he is speaking nor whether he is being listened to. He talks either for himself or for the pleasure of associating anyone who happens to be there with the activity of the moment. This talk is egocentric, partly because the child speaks only about himself, but chiefly because he does not attempt to place himself at the point of view of his hearer. Anyone who happens to be there will serve as an audience.* (Piaget 1932: pg 9).

Piaget observes that children talk regardless of whether or not there is an audience. Adults also chat to themselves when no-one else is around. The constant issuing of natural language sentences (whether internal or external) in the absence of an audience shows that these sentences are rarely intended for communication. In fact, these sentences are often fragmentary -- switching rapidly from one stream of thought to another -- and would not make much sense if another person did happen to hear them. So, if these sentences are not used for communication, we have to consider the

possibility that they are used for something else -- for thinking. It would seem to be a puzzling feature of human design if this rich verbalization turned out to serve no purpose. Introspection leads us to believe that our inner verbalization serves the purpose of thinking. We reason by running through streams of natural language sentences. Consider thought processes such as this: 'If I leave home now, I should get to Fred's house by about 3pm -- unless there is too much traffic, which would make it a bit later. There's usually alot of traffic at this time of day, so I'd better call and let him know that I might be late.' We reason like this all the time and the process is always conducted in sentences of natural language. Our awareness of this process being carried out in natural language would be wasted if it was all being carried out behind the scenes in mentalese. Furthermore, if natural language was not being used for our reasoning, and if mentalese processes were completely non-conscious, then wouldn't we simply find ourselves carrying out actions that had been decided by behind the scenes processes? Obviously, the simple answer to this last point is that sometimes we *do* simply find ourselves carrying out actions that have been deployed by non-conscious processes. Just go for a walk and you will notice that a great deal of complex actions are being set in motion by non-conscious activity. However, we have to remember that the existence of such activity does not entail the existence of mentalese. Operations such as walking may involve non-sentential mechanisms while our intellectual, or reasoning abilities involve natural language. Situations in which we find ourselves carrying out intelligent, reasoned activities without any awareness of the reasoning involved may involve natural language structures that are either not present in consciousness, or are not stored in memory.

Our introspective evidence gives us to reason to believe that our inner speech is used for thinking and that our thoughts are constructed from sentences of natural language. However, we must remember that the existence of inner speech does not *settle the question* about our thinking being carried out in natural language. There are other ways to account for the existence of inner speech, which should not be overlooked. We could account for our experience of inner speech by suggesting that the brain's modules operate in mentalese and send data to the speech center, which acts as a junction box and sends the data on to other parts of the brain. When information arrives at the speech center, it is automatically translated into natural language -- ready to be spoken if required. Thinking of the language module as a junction box could help explain the way in which the brain's modules link together in order to work collectively and form a unified mind. I think, however, the idea of constant translation between mentalese and natural language leads to an unnecessarily complex view of the mind. Why should a mind have to spend so much time on translation tasks? It would seem to be a waste of resources, especially when we consider the possibility that the entire process could be carried out in either mentalese or natural language without the need of any translation. Furthermore, consider the speed at which we apprehend the utterances of others. If the speech of others needed to be translated into mentalese before being understood, we might expect some delays to occur when parsing certain long sentences. But there are no such delays. We understand spoken sentences of our native language in real time -- *as we hear them.* We also seem to be able to speak our thoughts aloud *as we think them*; and of course we appear to understand our own thoughts *as we speak them* (internally or externally) with no delays. If thinking was done in mentalese and then translated into natural

language to be either heard as an internal monologue or spoken aloud, before being translated back into mentalese to be examined and understood by other parts of the brain, we would expect delays to occur in thinking. But our thinking seems to occur in real time without any lag behind.

Of course, the possibility of our thought process being carried out entirely in mentalese or natural language leads to its own problems. This is because it implies that each of the mind's modules is capable of using either mentalese or natural language, and we do not want a picture of the mind like that. It is difficult to imagine that each and every one of the brain's sub-systems has a complete understanding of whatever language the language of thought happens to be. A better picture is one that posits a global language of thought that is not fully used or understood by any of the brain's individual sub-systems. It could, for example, be the case that the brain's modules are connectionist networks and that their connection to the language module allows them to communicate to each other. Perhaps the language module acts as a junction box facilitating the communication between these otherwise unconnected networks. This could explain our inner speech and the role of natural language without requiring the existence of mentalese. Another solution is to think of the mind's modules as functional abstracts and not delve too deeply into their physical origins. In this way, we can speak of a system wide language of thought without having to worry about the physical implementation problems. This would leave us with the question as to whether mentalese or natural language is the language of thought and bring us back to the evidence of introspection, which leads us to believe that we think in natural language.

It should also be mentioned that it is entirely plausible that the 'inner speech' phenomenon is an *accidental* result of the way in which the brain is wired up. Our thought processes may be carried out in mentalese, and it is an accident that some of it slips through to the speech production center. Once there, it is automatically translated into natural language, thus making us aware of thought processes that would otherwise have remained hidden. The phenomenology of certain experiences could be used to support this idea. People occasionally report that the solution to a problem just 'pops' into their minds fully formed -- sometimes this happens several days after they originally started thinking about the problem. Such experiences may be put forward to show that the mind works on problems behind the scenes and lets the solutions slip through to the speech production module once the problems are solved. In fact, a mentalese supporter might suggest that while our conscious, occurrent thoughts are being carried out in natural language, all the non-conscious, behind the scenes activity is being carried out in mentalese. It is interesting to note that Carruthers does not rule out this possibility. Throughout his book, Carruthers makes it clear that he is interested only in our *conscious* thoughts. His account allows for the possibility that non-conscious thinking may involve mentalese. I should make clear that I am not suggesting that we do not have any non-conscious thoughts. I don't doubt that a lot of creativity goes on below our level of awareness; I actually think it is possible that a huge amount of non-conscious thinking takes place. But there is no compelling reason to suppose that this non-conscious thinking must be carried out in mentalese. Our non-conscious thinking may be non-sentential; or as I will suggest in chapter 4, it is possible that non-conscious thinking can involve the manipulation of natural language sentence structures below the level of awareness.

Without any evidence to support the possibility that unconscious thinking is being carried out in mentalese, we should continue to take the evidence of introspection seriously and favour the hypothesis that thinking is done in natural language. Moreover, we have to consider simplicity and economy in our theory of the mind. That is, we should not postulate an extra language, such as mentalese, without sufficient evidence -- especially when we consider the fact that natural language alone requires such a significant amount of the brain's resources, which are finite, and the possibility that natural language can do all the required work itself.

There are further reasons for believing that natural language plays a vital role in human thinking. Examples of humans who have been deprived of language offer us compelling evidence to support the claim that thinking is done in natural language.

(2) *Psychological development in the absence of language*

There are many documented cases of children who have been denied access to language. Malson refers to these children as "wolf children", and lists 53 recorded cases (pg 81-82 of his 1964). In most of these cases the psychological development of the child is severely impaired; and furthermore, in extreme cases the child's behaviour indicates thought processes that are no more sophisticated than those of non-linguistic animals (see Malson 1964 for several specific examples). Interestingly, however, these children often start to develop normally after they have been exposed to language. For example; a child named Genie was discovered in the 1970's when she was 13 years old. At that time she had a mental age of about 2 years and could not speak. She had spent her entire life locked in a small room with minimal human contact. In fact, the only time she ever saw another human was when her father or

brother brought her food -- and they never spoke to her. Despite her unfortunate upbringing, Genie (after being rescued) managed to acquire language and develop a nearly normal level of intelligence (Curtiss 1977). Does this show that language is essential for thought and psychological development? Does it show that language is simply a medium by which children can assimilate the concepts and culture of their community, and thus behave in a way that reveals the fact that they are thinking, while at the same time leaving open the question as to the nature of the process involved? Or is there some other explanation? In the case of Genie the conclusion we should draw from the evidence is far from clear. Genie was not only deprived of language, she was deprived of normal emotional contact with humans (Carruthers 1996: pg 42). She was not allowed to leave her room and spent a large amount of time restrained in a home made strait-jacket; meaning that she was not exposed to novel events and did not have a stimulating environment. These factors (in addition to her lack of exposure to language) may have played a part in inhibiting her mental growth. Furthermore, it is difficult to be certain as to whether Genie did in fact develop a normal level of intelligence after being rescued. She performed quite well in intelligence tests, but her language usage remained abnormal (Curtiss 1977: pg 204) and as such it is difficult to determine how sophisticated her thought processes actually became.

There is a wealth of material pertaining to such cases. I will consider the case of Helen Keller -- a well known example of a child who had an abnormal upbringing -- and I will use her testimony as support for my claim that thinking is carried out in sentences of natural language. The reason I want to refer to Keller's case is that she did not suffer the psychological abuse that has been associated with some of the other

cases. Furthermore, in her writings, Keller offers us a clear comparison between being a linguistic creature and being a non-linguistic creature.

Helen Keller lost her sight and hearing between the ages of 1 and 2 years. Being both blind and deaf, Keller was not able to acquire language in the normal fashion. She remained a languageless being until the age of 7, at which time she was taught to use a language of touch. This was achieved by her teacher repetitively exposing Keller to an object (water, for example), and then spelling the name of the object by tracing symbols onto her hand. Eventually Keller managed to master her language of touch and went on to write books about her experiences both before and after learning language. One of the points that becomes clear in Keller's writings is her belief that language plays an essential role in her thinking. She seems to believe that she did not exist as a thinking being before becoming a language user.

> *Before my teacher came to me, I did not know that I am. I lived in a world that was a no-world. I cannot hope to describe adequately that unconscious, yet conscious time of nothingness. I did not know that I knew aught, or that I lived or acted or desired. I had neither will nor intellect* (Keller 1909: pg141).

I should note here that Keller's introspective testimony is no more authoritative than our own introspection. If we can't settle the question of the role natural language plays in thought by appealing to our own introspection, then merely appealing to Keller's introspection should not settle the question either. However, Keller's testimony can be used to add strength to our own introspective evidence. This is because Keller can -- through introspection -- describe the difference between having language and not having language. Keller, unlike most of us, has memories of her pre-lingual existence. It is a comparison between those memories and her linguistic

state that lead her to make very interesting claims that highlight the importance of natural language in thought. Keller (1909: pg 143) claims that before acquiring language she had no inner life and that her behaviour was much like the instinctual behaviours of animals. Before language, Keller had "no power of thought", and "did not compare one mental state with another". She says that her sense of self identity did not exist until after she learned to use language, and it was not until then that she began to think (Keller 1909: pg 145). Interestingly, Keller claims that after acquiring language, she began to experience an "inner speech", which took the form of a feeling of words being spelled out into her hand. Thus, it seems that her private thinking was carried out in terms of her language of touch. Indeed, Keller says that if she were to construct a person, she would "put the brain and soul in his finger tips".

Keller's case is a striking example of someone who can articulate the introspectively perceived difference between being a language user and being languageless. In her writings we see the role that language appears, to her introspection, to play in thinking and the formation of a human mind. The fact that Keller perceived her thoughts as being constructed from words being spelled into her hand reasserts the cognitive function of language and emphasizes the role of inner speech. It would seem that our inner speech must be more than a mere 'side effect' or 'accident' because if it were just an accident, we would not expect Keller to experience it through touch. We can speculate that human brains are genetically determined to become wired in such a way as to connect the vocal output device to the auditory system, but it is more difficult to imagine that Keller's finger movement system was genetically determined to become linked to her tactile detection system. It seems reasonable to suggest that this type of neural feedback link between language

input and output devices serves an important role in thinking. It is possible that as humans develop after birth, their brain's connect themselves up in a way that enables such a link. This activity would, of course, depend upon early childhood experience and that is why Keller's brain connected itself up in a manner consistent with her language input/output systems being centered around her tactile movement and detection systems.

It is unfortunate that Keller does not describe the onset of her inner speech. Our inner speech seems to develop over time as we progress through our early childhood. It has its beginnings as an external monologue and becomes internalized as time goes by. It would be interesting to find out if Keller ever went through a developmental stage in which she physically traced words onto her own hand before developing the ability to do it all internally. If such a developmental stage existed, we could cite it as a further similarity between her inner speech and the inner speech that other humans experience.

Keller's case provides evidence to show that installing language onto a languageless brain gives that brain the ability to produce, assimilate, and support complex thoughts. Keller's description of her pre-linguistic state shows that her thoughts were reflexive and were not subject to an intelligent will -- meaning she responded to stimuli in an automatic, non-thinking way. There was no creativity, no sense of self, and no unified experience. It was acquiring a language that enabled Keller's mind to grow and start working in a distinctly human way. Without language her mind and thoughts would probably have remained no more sophisticated than those of the non-human animals.

Despite the intuitive force behind Helen Keller's case, I anticipate an objection that needs to be addressed. Consider Keller's claim that she had no sense of self before acquiring language. Is this because she had no language? Or is it because she cannot remember that far back? We must remember that Keller wrote about her early life experiences many years after they happened. It is possible that she simply does not remember what life was like before the age of seven; and if this is so, we cannot conclude that language was responsible for her developing a sense of self and a unified mind. I cannot remember what took place during the week of 15th October 1977 when I was seven years old, and I would probably describe that lack of experience in much the same way as Keller describes her lack of early life experience. My lack of memory may lead me to doubt the fact that I was a thinking entity during that time, and I would be forced to describe my life during that period as a "time of nothingness". But I undoubtedly *was* thinking during that period of time, and if someone was to have asked me to describe my state *at that time*, I would have given reports of a rich mental life. The point is that Keller's descriptions of her early life may be nothing more that a description of a lack of memory.

My reply to this is to emphasize the fact that Keller can recall *absolutely no experience* of being a *unified thinking thing* before acquiring language. I cannot remember events from the 15th of October 1977, but I can remember other experiences that go all the way back to when I was about 3 years old. Keller's experiences as a thinking thing only go back to her age of 7 years. What took place before that she describes as a "nothingness". However, she also describes her pre-linguistic behaviour as instinctual. I think Keller has retained some sort of *tactile memory* of events that took place before language, but she was not in control of those

events. Describing her behaviour as instinctual may be a result of remembering her pre-linguistic states in terms of language. Perhaps as she moved around and lived in a reflexive, non thinking way, languageless memories were formed. Accessing those memories later (with the benefit of language) gives Keller the unique insight into knowing what it is like to be running on 'automatic pilot' with no thought guiding her actions. So when Keller talks about a time of "nothingness", she is *remembering an absence* of linguistic thought. This is not the same as my failing to remember anything about my state in 1977.

A mentalese supporter might respond to Helen Keller's testimony by reminding us that mentalese does its work behind-the-scenes. It is possible that Keller's early life experiences are a description of what it is like to have mental processes being carried out entirely below conscious awareness -- in mentalese. After all, we have to remember that there is nothing in Keller's testimony that challenges the idea that mentalese exists behind the scenes. Keller's testimony may be nothing more than a description of what happens when one learns to describe ones mental states in terms of natural language. And, of course, her story shows the influence that natural language has in shaping the way one thinks about oneself. But a response such as this would be nothing more than an assertion on the part of the mentalese theorist. I could also make such assertions. I could say that mental maps do all the behind the scenes work. Or I could say that Keller's behaviour provides evidence to show that nothing more than stimulus/response reflexes went on before she learned her language of touch. But this is all conjecture. We have to focus on what we have that is not just conjecture -- namely, the fact that from her point of view (and the point of view of observers), Keller started thinking in a human way when she acquired

natural language. It seems to be the case that natural language gave rise to her productive and systematic thoughts, of which there appeared to be no evidence before natural language came along.

Our introspective evidence coupled with Helen Keller's testimony gives us good reason to suppose that our thinking is carried out in natural language. I will now back up this evidence by considering the nature of our thoughts. My claim will be that the features of our thoughts can only be explained if we suppose that thinking is done in natural language.

(3) *The nature of human thought requires language.*

Recall from Chapter 2 that our thoughts exhibit features of *productivity* and *systematicity*. Briefly, to say that thought is *productive* is to say that by applying certain rules to existing thoughts, a person can generate new thoughts. To say that thought is *systematic* is to say that thoughts are composed of symbols that can be systematically rearranged in order to form new thoughts with different meanings. So, if a person can entertain a thought of the form **a loves b**, then that person is also capable of thinking **b loves a**. Fodor believes that these features of thought should be explained by appealing to an innate language of thought, which he calls *mentalese*. This is because he believes that only language can account for the productivity and systematicity of thought. I agree with Fodor that language is the only way to explain the nature of our thoughts, but I do not agree that mentalese is the right language to use. This is because our thoughts do not become productive and systematic until we acquire a natural language. If Fodor's mentalese hypothesis was correct, we would expect to find that the thoughts of pre-linguistic humans share the productive and

systematic properties of fully developed human thought. Indeed, one of the reasons for positing mentalese rather than natural language as the language of thought was to account for the thoughts of infant humans and non-human animals (Fodor 1975: pp 56-58). Fodor thinks that productive and systematic thought is required for infants to learn natural language through forming hypotheses about the meanings of words and then testing those hypotheses. But experimental evidence shows that animal and infant human thoughts are not productive or systematic (Garfield 1997: pg 429). This is not to say that animals and infants do not think; it is simply to say that their thoughts are different. It could be that natural language provides a new way of thinking for humans once they reach a certain stage of development (Harman 1973). Jean Piaget's work with human children shows that there is a specific stage in development at which time children acquire a symbolic system of representation. This is important because (as we saw in chapter 2) symbols are required to produce productive and systematic thoughts. Piaget characterized four stages of cognitive development that all children go through. The first of these stages -- the *sensory-motor stage* -- is of relevance here.

During the *sensory-motor stage*, which lasts from birth until about the age of 2, the infant's cognitive system starts off as being limited to motor reflexes. The child's thoughts are primitive and consist in the simple coordination of sensory information with bodily movements. According to Piaget, as development continues, the child eventually makes its first big jump in cognitive ability and acquires *object permanence*. This is the realization that objects continue to exist even when they are moved out of sight. Piaget observed that before reaching this stage of development, children act as if objects do not exist when they are not being perceived (Piaget 1953:

pg 211 and Piaget 1983: pg 104). A child may be very interested in looking at a toy, but if the toy is then hidden behind a piece of paper (while the child watches), the child will not move the paper to find the toy. However, once the object permanence stage of development is reached (between 3 months and 1 year according to recent research; see Baillargeon 1991), the child understands that objects continue to exist and will seek out hidden objects. For Piaget, object permanence is a crucial stage of development because it marks the beginning of *symbolic* thought. It is at this time that infants can hold concepts in mind and can start to consistently use specific words to represent specific objects. This leads to an explosion of symbolic language usage and the subsequent restructuring of thought processes that give rise to productive and systematic thinking.

The lack of symbolic representation in infant humans gives us reason to suppose that their thought processes lack the productive and systematic features of linguistic human thought. But does this show that it is the assimilation of natural language that allows human thoughts to become structured in that way? Perhaps not. Mentalese supporters could claim that the acquisition of symbolic representation occurs at a certain stage of cognitive development and it is at this time that an infant's innate language of thought mechanisms come 'on-line'. Alternatively, it could be suggested that an infant reaches the object permanence stage and acquires symbolic representation simply as a result of an innate language of thought mechanism firing up. There is, after all, nothing in the mentalese hypothesis to suggest that mentalese must be 'up and running' from birth -- it could happen later as a result of developmental forces. As I stated earlier, I think this suggestion leads to a view of the mind that is more complex than it needs to be. It gives us two languages that occur as

the result of developmental forces -- mentalese and natural language. Both of these languages can account for the systematicity and productivity of thought, but so far only one of these languages can be empirically pointed to. We should, therefore, appeal (again) to considerations of simplicity and economy. This is to say, we should attempt to keep our ontology as simple and uncluttered as possible. If natural language can account for the development of productive and systematic thought, then we should not posit the existence of mentalese.

A further reinforcement of the possibility that it is natural language that structures our thoughts comes from connectionist models of language. An artificial neural network that has been designed to acquire language, but currently has no linguistic abilities, has none of the productive or systematic powers of language. However, it has been observed that as such a network acquires concepts (through training), those properties emerge (Garfield 1997: pg 429; see also Clark 1993 and Dennett 1991b). The evidence from infant development leads us to believe that the same is true of the human brain. Natural language is an external technology that is productive and systematic. After the mind acquires the concept of object permanence, it is ready to assimilate natural language. It is at this time that the brain starts to produce thoughts that are productive and systematic. Of course, the brain's structure is probably different to the structure of an artificial neural network -- it is almost certainly more complex. So, facts about artificial neural networks cannot be assumed to be facts about human brains. The two systems may differ in some fundamental way, and this difference may be related to differences in the acquisition of productive and systematic thoughts. However, the fact that artificial neural networks acquire productive and systematic abilities through training can, at the very least, be used to

show the *possibility* that the brain behaves in a similar way. This possibility can be used as further support for the claim that human thoughts become productive and systematic with the assimilation of natural language; thus reinforcing the possibility that natural language is the language of thought.

4: What can we make of the evidence?

In this chapter I have provided reasons for supposing that thinking is done in natural language. The evidence in support of the claim that thinking is done in natural language came primarily from three sources -- introspection, psychological development, and the nature of our thoughts. In the case of introspection, we find ourselves describing our thinking in terms of sentences of natural language. If someone asks us what we are thinking, our expression of the thought is the same as our inner experience of that thought. Whether we are thinking internally or externally, the thinking always involves natural language mechanisms. Helen Keller's introspective testimony also serves to show us that thinking is done in natural language. Her psychological development was severely impaired until she managed to acquire her language of touch. Furthermore, she reported that before learning to use language, she was not a thinking thing and experienced no will or intellect.

The evidence from the nature of our thoughts shows that the productive and systematic aspects of our thinking require symbolic representations. Young children do not have symbolic systems of representation and so do not have productive and systematic thoughts. However, when infants acquire the ability to symbolically represent things to themselves, they quickly learn language. The syntactic rules of language then re-organizes their minds and gives rise to productive and systematic

thoughts. Experiments with artificial neural networks show that this process works in practice. Networks with no productive or systematic abilities have been designed to acquire language. As these networks acquire a language, their operational processes become increasingly productive and systematic -- just like a human mind.

Although this chapter has provided good reasons for believing that natural language is the language of thought, the issue is not yet closed. In the next chapter I will consider arguments that have been put forward to deny the role natural language plays in thinking. Arguments against the role natural language plays in thinking are intuitive and are designed to convince us that mentalese is the required language of thought. It will be my task to show that these arguments can be answered in a manner consistent with the hypothesis that natural language is the language of thought.

4

The Last Stand...

Don't Replace The Old Code Yet!

1: The Fight for Mentalese

The defense of natural language as the language of thought in the previous chapter has moved us toward a better understanding of the relation between language and thought, and the nature of human minds. Despite this, however, we are not yet in a position to close the book on mentalese. Having identified evidence and reasons for believing that natural language is required for thinking, I will now consider arguments that have been put forward to refute this claim. Some of these arguments rely on intuition and familiar phenomena such as the 'tip-of-the-tongue' phenomenon. Other arguments ask us to consider cases in which the behaviour or testimony of people give us reason to believe that thinking is not carried out in natural language. Einstein, for example, claimed to think in visual images, while on the other hand, a man named Brother John continued to behave intelligently even though a neurological condition known as aphasia had disabled his linguistic abilities.

On the surface, the arguments against the cognitive role of natural language seem convincing and may tempt us back down the path of mentalese. However, looking carefully at the facts surrounding the basis of these arguments shows that they can be answered in terms of natural language being the language of thought. For example, I will suggest that it is possible to explain the 'tip-of-the-tongue' experience in terms of selective memory failure, or in terms of difficulties in thought formation.

In the cases of Einstein and Brother John, I will show that the evidence is not as clear cut as it seems. I will attempt to provide reasons for believing that natural language is essential for their thinking. My conclusion will be that attempts to convince us that mentalese must be posited as the language of thought are unconvincing and ultimately fail.

2: Pinker's Resistance

One of the major themes in the early chapters of Steven Pinker's book "The Language Instinct" is the battle between natural language and mentalese as the language of thought. Pinker, as I stated earlier, champions the thesis that mentalese is the language of thought and argues forcefully against the hypothesis that natural language is the language of thought. In the following passage, Pinker raises several problems for claim that thinking occurs in natural language:

> *We have all had the experience of uttering or writing a sentence, then stopping and realizing that it wasn't exactly what we meant to say. To have that feeling, there has to be a 'what we meant to say' that is different from what we said. Sometimes it is not easy to find* any *words that properly convey a thought. When we hear or read, we usually remember the gist, not the exact words, so there has to be such a thing as a gist that is not the same as a bunch of words. And if thoughts depend on words, how could a new word ever be coined?* (Pinker 1994: pg 57-58).

There are several good points contained in this passage, and collectively they pose a threat to the idea that natural language is the language of thought. Because of this threat, Pinker thinks that thinking cannot be done in natural language and sentences of natural language must be translations of thoughts that exist in some other medium, which is probably mentalese. Before proceeding, I must note that Pinker's threat does

nothing to show that thinking must be carried out in mentalese. In the above passage, Pinker is attempting to show us that thinking cannot be done in natural language. But even if he succeeds in convincing us that this is true, the problems he points to cannot convince us that thinking must therefore be carried out in mentalese. There could be some entirely different mental mechanism at work such as mental maps or visual imagery..

Continuing my defense of natural language as the language of thought, I will isolate Pinker's points and attempt to provide answers to them consistent with the idea that natural language is the language of thought. If I am successful, the answers I provide will show that Pinker's arguments do not entail the conclusion that thinking is not be carried out in natural language. I will then consider the case of Einstein and the case of Brother John, both of whom seem to be counter evidence to the claim that thinking is done in sentences of natural language.

(1) *Saying something other than what I wanted to say.*

This type of experience is familiar to most people. It can happen in conversations when you want to express a point but cannot find the right word to use in the sentence. In the end you settle for some other word that expresses the point in a slightly different way. Sometimes the feeling that there is a better word is so strong that you may experience the 'tip-of-the-tongue' phenomenon. You *know* there is a word that you need to finish expressing your thought, but you cannot access it. It is possible to induce the tip-of-the-tongue experience in people by asking them to provide a name for an object according to its definition. For example, if asked to name the following object: *an instrument with a graduated arc of 60 degrees used in navigation*

(especially at sea) and surveying for measuring the angular distance of objects by means of mirrors, most people will have difficulty recovering the name. Even if they can form a mental image of the object in question and know that under normal circumstances they would be able to name it, in this instance they cannot. Pinker believes that this type of experience shows that thought is independent of natural language. His position certainly captures the phenomenology involved in these cases. Such experiences make it seem as if the thought exists independently (perhaps in the form of mentalese) and there is a difficulty in translating the thought into natural language. Pinker thinks that if thinking is done in natural language we would not expect to have the feeling that there is a word we need to use in order to express an independently existing thought, yet cannot find that word. By the way, the object defined above was a sextant.

Carruthers (1996: pg 57) suggests another way of explaining the tip-of-the-tongue phenomenon. Rather than claiming that the thought exists in mentalese and there is a difficulty translating it into natural language, we could explain the phenomenon by stating that the thought is not fully formed and the difficulty lies in the formation of the thought. My experience of the phenomenon may be the result of my difficulty forming a thought, which subsequently forces me to form a slightly different thought. I might believe there is a better thought that would be more precise, but I cannot form that thought. But if I *could* form the thought, I know that I would recognize it as a better thought than the one I am currently entertaining. The obvious response to Carruthers' idea is to point to the difficulty in accounting for the belief that there may be a better thought than the one that gets uttered. Pinker might claim that the belief that there is a better thought comes about precisely because there *is* a

better thought, which is encoded in mentalese and cannot be adequately translated into natural language. The difficulty in translation subsequently results in the formation of another thought that is easily expressible in natural language. To respond to this, Carruthers must maintain that the belief that there is a better thought is not brought about by the existence of a mentalese encoding of that better thought, but by some other means. I do not see any problems with this position. In a tip-of-the-tongue experience, the belief that there is a better thought does not entail that the better thought is being entertained in mentalese. Consider Gauker's (1992) alternative ways to account for the tip-of-the-tongue phenomenon. Situations in which we think we cannot find the words to express a thought could, for example, be the result of a retrieval problem. We may have known exactly what we were going to say but then suddenly forgot. The frustrating feeling that follows may be brought about by our efforts to retrieve the original sentence. This may be similar to the tip-of-the-tongue sensation one experiences when trying to remember the name of a familiar face. In such cases one has forgotten the name of a familiar person, but also knows that in the past the name was always instantly available. It is the effort to remember the person's name that gives rise to the phenomenon. Another possibility is that these experiences may occur when one has a mental image that is not easy to describe in words; like an image from a disjointed illogical dream. Or, it could simply be the case that our supposed difficulty expressing a thought occurs because we actually have no thought to express. We may *want* to have an answer to a question or an opinion on a topic, but we don't, and we will not admit to ourselves that we don't. So we try to form a thought -- any thought -- that is relevant to the topic in question. Our failure to form such a thought results in the annoying tip-of-the-tongue experience. Possibilities such

as these show us that the tip-of-the-tongue phenomenon does not give us sufficient reason to deny that thinking is done in natural language. It certainly gives us no reason to suppose that mentalese is a required medium of thought.

A related phenomenon is the so called 'slip-of-the-tongue', which occurs frequently and often goes undetected. Here is a typical example: A few weeks ago I was watching a scene in a movie which involved an old woman. It was obvious to me that the woman was young and had been made up to look old. Her hair was dyed grey and she had wrinkles painted on her face. I commented on this by issuing the following statement: "Her hair isn't really that *dark*". I immediately realized my error and corrected my statement in this way: "I mean *light*". The question is, how is such a mistake possible? Introspection tells me that I wanted to say that the woman's hair was not light, and that is why I issued the correction. Does this show that the thought was formed in mentalese and a mistake occurred in translating the thought into natural language? It could be argued that my knowing that the sentence was not what I wanted to say shows that there must have been a thought that I did want to say, and that thought was independent of my natural language utterance.

Sigmund Freud tried to explain our slips of the tongue by suggesting that they are caused by our unconscious thoughts and desires. For Freud, our unconscious intentions, or wishes, somehow manage to push themselves to the surface to be revealed in normal conversation (Freud 1940: pg 284). I do not want to debate the plausibility of Freudian psychology, but I do think his account of the slip-of-the-tongue can help show that the phenomenon can be explained in terms of natural language, without the need of evoking mentalese. Notice that in most instances of slips of the tongue, the accidental word is related to the sentence. Freud explained this

by saying that the person is unconsciously thinking a related thought. In light of this suggestion, I think it is possible to explain the above case by suggesting that the thought I expressed was incorrect and I became aware of my inaccurate thought moments after I expressed it. It is possible that my thought 'that woman's hair isn't really that light', which is what I wanted to say, was somehow accompanied by, or caused by the related thought 'that woman's hair is really dark', which is also what I wanted to say. Somehow these thoughts merged during deployment and gave rise to the incorrect statement 'that woman's hair isn't really that dark'. Then, as I heard my expression, I noticed that it did not match the features contained in my visual field and so I issued a correction. It may be the case that instances of slips of the tongue point to multiple streams of thought, most of which remain outside of consciousness. But it is important to realize that the existence of non-conscious thinking does not entail the existence of mentalese. As I will suggest later, it is possible that our non-conscious thinking also involves natural language structures (see the section on Brother John).

We have seen that the 'tip-of-the-tongue' and 'slip-of-the-tongue' phenomena can be explained without the need of appealing to mentalese. I will now move on to the next part of Pinker's attack against natural language as the language of thought.

(2) *We usually remember the gist of what is said and not the exact words.*

For Pinker, the so-called 'gist' of a statement differs from sentences of natural language. This is because we usually do not remember the thoughts of other people exactly as they communicate them to us -- rather, we remember roughly what they said, or the 'gist' of what they said. Does this show that the 'gist' of a thought is

stored as a hidden sentence of mentalese? Or is there some other way to explain the 'gist' of what we hear and read?

It is true that we usually do not remember what other people have told us in exactly the same words that they have used, but this does not refute that claim that thinking is done in natural language. Pinker says that the gist of what is said must be something different from a bunch of words, but I do not agree. I think it is possible that the gist of a sentence is a sentence itself. If I think about the gist of a natural language sentence I have heard, the thinking of that gist is done in natural language. It may not be the same sentence that I originally heard, but it is still a sentence (perhaps with a slightly altered meaning). It could be that as I hear the utterances of other people, my mind alters the utterances in order to store them in the most efficient way possible. The result of this would be the formation of slightly different sentences that constitute roughly the same idea as the original sentences.

Pinker might object to this suggestion by claiming that I am only explaining why we store different sentences and not why we store a gist. He could say that the fact that these different sentences mean the same thing shows the existence of an independent gist, which may be a mentalese encoding of the natural language sentences. Indeed, one of Pinker's reasons for supporting mentalese comes from our ability to translate sentences from one language to another while keeping their meanings intact. For Pinker, this ability shows that our thoughts have a meaning independent of our natural language formulations of them. However, there is a problem associated with this line of reasoning (Cole 1999). Saying that two different natural language sentences have the same meaning in virtue of there being an underlying gist to them, leads to a regress. This occurs when we attempt to account

for the meaning of the gist. We may be led to believe that the gist is a mentalese sentence, but then we have to explain the likeness in meaning between these gists and their natural language counterparts. How should we explain that? Perhaps we could say that they both have the same 'vibe'; but then a regress ensues when we try to account for that vibe. We can avoid these problems by not thinking of the gist as a mysterious, hidden mentalese encoding of a natural language sentence. Instead, the gist of a sentence could more accurately be thought of as a different sentence that we store in an attempt to remember a thought in the most efficient way possible.

Alternatively, we could think of the gist of a sentence as a *recipe* for reconstructing the original sentence[2]. I find this a convincing solution to the problem. Consider the vast amount of information that we can recall. Now, the brain has a limited storage capacity and cannot be expected to record all of this information as it is presented. It is therefore possible that when we recall experiences or sentences, a certain amount of reconstruction occurs. This reconstruction may involve a process by which fragments of the original experiences are pieced together according to a recipe, while the brain fills in the missing details. So, when I read or hear a sentence that needs to be stored for future use, that sentence is deconstructed into a collection of fragments, and a recipe is formed determining how the sentence should reconstructed. On this account we can say that when Pinker talks about the gist of a sentence, he is talking about a recipe for recreating a sentence. Of course Pinker could agree with this analysis and suggest that talk of recipes amounts to the same thing as talk of mentalese encodings of sentences. He would be correct in making this claim if he could show that thinking involves operations upon, and manipulations of

[2] This was pointed out to me by my supervisor, Derek Browne.

recipes. But it would be difficult for Pinker to show this because recipes contain nothing more than fragments of information and instructions that show how the information should be reassembled. There is nothing in a recipe for a thought process to get a grip on. Moreover, if a mental process *could* manipulate the fragments, then the recipe specifying their reconstitution would become useless because the original fragments would no longer exist. The situation is analogous to attempting to slice a cake before it has been made. How can you perform such an operation upon a recipe and raw ingredients? The answer is: you can't. So, the existence of recipes (or gists) for reconstructing sentences does not refute the claim that thinking is done in natural language. It simply points to possible storage and retrieval system and does not suggest that thinking could be carried out by performing operations upon recipes. The sentences must be reconstituted in order to play a role in thinking. As I see it, the idea is similar to the practice of compressing data for efficient storage in a computer system. In its compressed state, computer data cannot be operated upon by the same mechanisms that are in place to operate on the data in its uncompressed state. The data must be decompressed in order to be used by the appropriate software. In keeping with this terminology, we can think of the gist of a sentence as a combination of compressed data and an algorithm specifying the reconstruction of that data. Once it is reconstructed, it can then play its role in thinking.

(3) *If thoughts depend on words, how could a new word ever be coined?*

Pinker seems to think that new words are invented in order to express new thoughts that are formed in mentalese. He thinks that if thought was dependent on words of

natural language, then new words would never appear because there would be no new thoughts that needed to be expressed with new words. All of our thoughts would already be formed in words and so new words would not be necessary.

I do not think this question requires us to believe that mentalese is the language of thought, nor does it pose any serious problem for the claim that thinking is done in natural language. The meanings of new words can always be expressed in terms of existing words. Even names for new inventions (such as 'transistors' for example) can be expressed with existing words that describe the function of the new invention. New words usually catch on when they form a complex meaning (or thought) in a quicker, more direct way. In other words, new words make thinking more efficient. Take the word 'eftpos' for example. This word was originally an acronym that stood for 'Electronic Funds Transfer at Point Of Sale'. Now, although many people do not know what 'eftpos' actually stands for, if asked to explain what they mean when they use the word 'eftpos', they will usually describe it in similar terms; for example, money coming straight out of their bank account when they make a purchase. It is not the case that a new thought arose that needed a word to express it. Rather, the new word appeared in order to form a complex thought that was originally constituted by a number of other words. Another example is the word 'scuba'. Most people know that this word has got something to do with ocean diving, but few people know its origins. 'Scuba' was an acronym that stood for 'Self Contained Underwater Breathing Apparatus'. The acronym 'caught on' for obvious reasons -- its easier and quicker to say. It makes talk and thought more efficient.

Sometimes words and catch phrases are coined simply because of the way they sound (Cole 1999). An example of such a phrase is *Rock 'n' Roll.* This catch phrase

is only remotely descriptive of the way in which people danced while listening to the contemporary music of the 1950's and 60's. The reason it caught on was because of its use of alliteration. It sounds good and its easy to say. It is difficult to see how a phrase like 'rock 'n' roll' could have been coined if thinking was done in mentalese. Mentalese does not contain the letter 'R' and so there seems little reason for a mentalese sentence that corresponds to 'rock 'n' roll' to form. Why would a mind put these two words together? If natural language sentences are just a translation of mentalese thoughts, we would not expect to find catch phrases that rely on alliteration. As far as the mentalese thought would be concerned, the phrase 'rock 'n' roll' could mean 'stone 'n' tumble', and it is hard to imagine that a mind would think this thought when thinking about a musical genre. The only reason the words 'rock 'n' roll' were put together was because of their appeal in natural language.

Now that the preliminary objections have been addressed, I will move to the next section in which I will consider Pinker's further objections.

3: Pinker's Continued Resistance

Pinker offers further reasons that are supposed to convince us that thinking cannot be done in natural language. These are (1) the problem of ambiguity, (2) the problem of logical inexplicitness, (3) the fact that Einstein thought in visual images rather than words, and (4) evidence from aphasia. I will tackle these one at a time and will show that they offer no serious threat to natural language theories of thought.

(1) *The problem of ambiguity*

This problem points to the fact that many words have more than one meaning, and sentences that contain such words are often ambiguous. An example that is often cited is the sentence: *Jane is at the bank.* The ambiguity in this sentence occurs because of the word 'bank'. We could take this sentence to mean one of two things -- either Jane is at a financial institution, or Jane is at the edge of a river. Pinker (1994: pg 79) lists several newspaper headlines that happen to contain an ambiguous word, and claims that the thoughts underlying these words were not ambiguous even though their natural language counterparts are. He then suggests that since there can be more than one thought corresponding to a specific sentence, thoughts cannot be sentences (Pinker 1994: pg 79). So, on Pinker's account, if Jane is adding money to her account, my thought that 'Jane is at the bank' has a definite meaning, which is determined by a mentalese sentence, and is distinct from my ambiguous natural language utterance 'Jane is at the bank'.

I do not find the problem of ambiguity a convincing objection to the claim that thinking is done in natural language. Pinker's claim is that certain sentences of natural language can correspond to several different sentences of mentalese. But if this is true, why do we find it so easy to understand exactly what someone is saying when they utter an ambiguous sentence like 'Jane is at the bank'? How does my mind know which mentalese sentence the utterance 'Jane is at the bank' is supposed to correspond to? Pinker's picture would have us struggling to fix the meaning of many normal, everyday statements. It is true that certain words have more than one meaning and can be used to construct sentences that are ambiguous, but these sentences are only ambiguous when viewed in isolation. Consider the last sentence in

the previous paragraph. I talked about Pinker's 'account' of ambiguous words, and I talked about Jane's bank 'account'. If we isolate small portions of these sentences, say 'Pinker's account', and 'Jane is adding money to her account', we end up with two instances of ambiguity. In the first fragment, we could be talking about Pinker's bank account or we could be talking about his account of natural language. In the second fragment, we could be talking about Jane depositing money into her bank account, or we could be talking about Jane including aspects of finance in her account of political corruption in some foreign government. However, when these fragments are included in the original sentence and are tied to the context of the present discussion, no ambiguity exists. I think Pinker makes the mistake of looking for ambiguities in small, isolated sentences without considering the context in which they belong. An analogous situation sometimes occurs when we view 2-dimensional drawings of 3-dimensional objects. For example, if someone shows me a drawing of a small set of stairs, I can see it either from the perspective of looking down upon the stairs from above, or from beneath looking at the underside of the stairs. But the ambiguity in this visual image only occurs when the set of stairs is isolated. If the stairs are included in a complete scene of the inside of a house, then they will not have this ambiguous appearance. The same is true of words of natural language. The word 'bank' has several meanings but when the word is used in a sentence and the sentence is heard as part of a larger context, the ambiguity disappears. If I am in a park near a river and I point towards the river and say 'Jane is at the bank', then my sentence is not ambiguous. Similarly, if I am on a city street in the center of town, waiting outside a bank and say 'Jane is at the bank', my sentence is not ambiguous. My disposition to utter the sentence came about through a collection of previous activities

that fit together building a scenario determining the context in which the sentence is being used. So the sentence is not uttered and understood in isolation. Instead, it is uttered and understood in relation to the situation it is being used in.

An obvious response from Pinker would be to agree with the preceding analysis and claim that an examination of context is the way in which the mind determines which mentalese sentence should correspond to the natural language sentence. So when someone tells me 'Jane is at the bank', my mind quickly reviews the history and any other information related to that utterance and decides which mentalese sentence to attach it to. In situations where the background information is incomplete, several mentalese sentences may be considered to be equal candidates and I find myself seeking more information. However, if Pinker attempts to make a move like this he will invite another response to his problem of ambiguity. This response will question the ability of mentalese to disambiguate sentences of natural language. If small sentences of natural language can be ambiguous in virtue of containing an ambiguous word, then it would seem possible that small sentences of mentalese could also be ambiguous for exactly the same reason. The only way a mentalese theorist could deny this possibility would be to say that mentalese has a separate word for each and every object and concept. So there is a mentalese word that only means 'financial institution' and a mentalese word that only means 'the edge of a river'. But we have to be skeptical of the idea that mentalese can have, built into it, a word for every concept and object because if it did, then it would seem that the ancient human mind contained concepts such as 'deep space travel', 'computer' or 'radio waves' long before these things were discovered or invented -- surely this is implausible.

(2) *Natural language is sometimes logically inexplicit*

Here Pinker claims that natural languages such as English sometimes lack logical explicitness. He shows this by drawing on an example from computer scientist Drew McDermott:

Ralph is an elephant.

Elephants live in Africa.

Elephants have tusks.

Now, we can deduce from these simple statements that Ralph lives in Africa and Ralph has tusks. We also realize that Ralph lives in the same Africa as the other elephants, and that Ralph has his own tusks, which he does not share with any other elephants. But, the problem is that there is nothing in the natural language statements above to imply that all elephants don't share one set of tusks. Nor is there anything in the natural language statements to imply that all elephants live in the same Africa.

Despite the logical inexplicitness of the above statements, I do not think that Pinker can use them as strong evidence to deny the claim that natural language is the language of thought. Mentalese, if it exists, would also contain sentences that are logically inexplicit for exactly the same reason as their natural language counterparts. Also, it seems that Pinker is forgetting to take into account our background knowledge and beliefs about elephant tusks and Africa. If I did not know what elephant tusks were, or did not know that Africa was a country, I might well believe that all elephants share one set of tusks or live in different Africas. But I *do* know what these

words mean and my knowledge is describable in natural language. Given my knowledge, I could easily reword the above sentences in the following way:

Elephants live in the one and only country called Africa.

Elephants each have their own set of tusks.

Worded in this way, these sentences are logically explicit and cannot be confused. The point I am trying to make here is that an isolated thought such as 'elephants have tusks' is connected in some way to a background pool of knowledge and beliefs, all of which involve natural language. Moreover, when I think the thought 'elephants have tusks', the thought is logically explicit because it induces a mental image of one or more elephants with tusks. As the image is mentally scrutinized, the tusks are noticed and labeled in natural language thus making the thought logically explicit. This is not to say that sentences are always accompanied by images, but it does offer one way in which a sentence can become logically explicit. In the section on Einstein coming up, I will show that there is no problem with allowing that thoughts can involve mental images with natural language sentences attached to them.

Having provided answers to the above problems, I will now consider the case of Einstein and the case of Brother John. Both of these people provide us with introspective testimony that could be construed as showing that thinking is possible without natural language. It will be my task to offer an alternative explanation of the evidence these cases provide.

(3) *What about Einstein's thoughts?*

Pinker (1994: pg 71) points out that Albert Einstein claimed to think in terms of visual images. In fact, Einstein's theories of general and special relativity have their origins in the manipulation of visual imagery. Einstein began his theorizing by imagining what things would look like if he was traveling on a beam of light. Pinker believes that Einstein's case offers us evidence that our thinking can be carried out without natural language; but, as I will explain, I do not think the evidence is as clear cut as he claims. Furthermore, I do not see how Einstein's thoughts can strengthen Pinker's position. Pinker wants to convince us that thinking is done in mentalese rather than natural language, but the case of Einstein does not show this. If Einstein's thoughts were entirely visual (non propositional), then mentalese would suffer as much a blow as natural language would. Einstein's claim that his thinking involved visual imagery does nothing to support the claim that thinking is carried out in a hidden, innate language of thought -- rather, it threatens the idea that thinking involves sentential structures of any sort, including mentalese (Cole 1999).

One way to disarm the threat from Einstein's testimony is to allow that his thinking involved visual imagery, but deny that his thoughts lacked sentences of natural language. It is true that Einstein employed visual imagery in his thought processes, but I have already agreed with Carruthers' (1996) suggestion that our thoughts can be a combination of words and images. I think that Einstein's thoughts involved such a combination. If I am right, Einstein's thought process was something like the following:

'If I were traveling at the speed of light, objects in front of me would look like this [IMAGE], and objects behind me would look like this [IMAGE].'

or,

'If I drop a coin while descending in a glass elevator, from my point of view it would look like [ANIMATED SEQUENCE OF IMAGES]. But from the point of view of a stationary person standing outside the glass elevator it would look like [ANIMATED SEQUENCE OF IMAGES].'

Some cognitive scientists believe we have pictorial representations in our heads that can be mentally manipulated and inspected. Crane (1995: pg 148) thinks this is plausible and suggests that if someone asks you the question: 'Do frogs have lips?', you will probably form an image of a frog and mentally examine it. Einstein's thinking could have involved the mental manipulation and inspection of images, which he then reasoned about through the use of sentences and linguistic operations upon those sentences.

Visual images must be augmented by sentences in order to have meaning. It is difficult to imagine how the content of a complex thought such as: 'Future cognitive scientists may one day discover how sentential structures are implemented upon the connectionist architecture of the brain', could be represented as an image -- or even as a series of images with no syntactic structure. To strengthen this point, consider the simple sentence: 'It *might* rain *tomorrow*'. It would not be possible to have this thought in non-propositional visual imagery. Visual imagery alone cannot enable us

to think temporal thoughts like 'tomorrow' or 'yesterday', nor can it be used to construct thoughts that involve concepts such as 'might', 'all', or 'necessarily'. We can form a visual image of rain but there can be nothing in the image to distinguish tomorrow from the next day, or the day after. A good way of illustrating this point was introduced by Wittgenstein (1953, section 332). He asked us to think a thought such as 'it might rain tomorrow', and then think the same thought again without any words, while leaving the meaning intact. It can't be done. If we try to do so with visual imagery, all we end up with is an image of rain. Now, before proceeding I should mention the possibility that deaf people who have not acquired the proper use of sign language, could have the ability to represent the passage of time. They could achieve their representation of time by moving their arms apart to represent a distance in the past. The distance apart they stretch their arms could be proportional to the distance back in time they wish to represent. If this behaviour was observed in deaf languageless people, it could be used as evidence for temporal thoughts in the absence of natural language. However, I think this conclusion would be too hasty. We have to remember the fact that these people would be representing the passage of time through the use of a mimetic sign. This sign may not conform to any of the standard sign languages, but it would be a sign nonetheless and should be classed as a linguistic tool. So it could not be concluded that languageless people can represent the passage of time. These people would not be entirely languageless.

It may be claimed that a deaf, languageless person could think a thought such as 'tomorrow' in visual imagery by constructing a visual sequence that shows the sun rise, fall, and then rise again. However, the problem with this is that without words or signs it would not be possible to determine exactly what the visual sequence was

supposed to represent. It could be a thought about 'change', 'movement', 'light and dark', 'astrophysics', or any number of different ideas that involve the motion of the Earth and sun. I think that some sort of sign would be required in order to precisely define what the visual imagery is supposed to be a thought about.

Continuing to highlight the superiority of natural language over visual imagery, consider abstract concepts such as 'equal', 'speed', or 'cause'. There are no images that can accurately express what we mean by these concepts. Images can only represent the appearance of things and since abstract concepts have no appearance, images cannot be used to represent them. So, although it is possible that thinking involves images, it seems that it must also employ some other system that can be used for thinking about abstractions and for reasoning. Positing a sentential system of representation is the best way to account for our ability to form thoughts about abstract concepts. Of course, it could be argued that the mind might use a suitable sequence of images to represent abstractions. But this would still require some sort of sentential type structure in order to string the images together in the correct way. I find it difficult to see how such a system would be distinct from natural language. In fact, a structured system of images could be considered to be a natural language much like the ancient Egyptian hieroglyphic languages. The images would presumably represent spoken words, just as written words represent spoken ones. If this were the case, we would probably maintain that thinking was done in sentences of natural language.

(4) *The case of Brother John.*

Pinker (1994: pp 46-48 and 67) believes the psychological condition known as aphasia is a threat to the notion that thinking is done in natural language. People who suffer from aphasia have lost some (or sometimes all) of their language abilities. Such people may once have been perfectly normal but then, perhaps as the result of a stroke, hemorrhage or head injury, they find some aspect of their language capacity impaired. There are varying degrees and types of aphasia. In some cases, the subject may lose their capacity to speak, while in other cases, such as instances of jargonaphasia, the subject's speech may be intact but their utterances are senseless. Occasionally a sufferer's ability to produce written sentences is the only feature of language affected. Severe cases can see a patient's language faculty removed altogether, leaving them unable to comprehend or utter sentences. Pinker thinks that the existence of aphasia shows that thinking can occur in the absence of natural language. This is because aphasia sufferers seem to retain much of their intelligent human behaviour. Merlin Donald (1991) seems to agree. Donald (1991: pg 82-86) points to the case of Brother John, who suffered periodic aphasia, and claims that his condition shows the independence of non-linguistic from linguistic thought. Brother John suffered from paroxysmal aphasia and was studied by Lecours and Joanette (1980). Unlike the cases of aphasia that were brought about as the result of a stroke or brain injury, Brother John only suffered from temporary episodes, which seem to have been related to focal seizures. During his aphasic spells, Brother John's language processing ability was totally shut down. He lost his ability to speak and understand language; and he even lost his inner speech -- providing important evidence for the connection of inner speech to the rest of the linguistic input/output system. Yet,

despite his lack of language, Brother John remained aware of his disability and could later recall events that took place during his seizures.

> *I hear noise. I hear people talking. I know it means something but I do not understand. I try hard to understand in my inner self but I cannot* (Brother John recounting an aphasic episode. Quoted in Lecours and Joanette 1980: pg 9).

The interesting feature of Brother John's case was his apparent ability to cope intelligently with his environment during his loss of language. Lecours and Joanette document an event in which Brother John suffered an aphasic seizure, which lasted several hours, while he was on vacation in Switzerland. With no explicit language ability whatsoever, Brother John managed to find a hotel, check into a room, have lunch and then go back to his room for a nap. He accomplished the requisite communication through a series of careful mimes and gestures. To book into the hotel, he pointed to his name on his passport so that the receptionist could fill in the appropriate forms. He then found his way to the hotel restaurant and ordered lunch by pointing to a random dish on the menu (though, unfortunately he inadvertently ordered fish, which he was not fond of). After lunch he returned to his room and slept through the rest of the seizure.

The case of Brother John is a compelling example of someone who could obviously think despite the apparent absence of any natural language ability. During his languageless spells, Brother John coped with his environment and handled situations in an intelligent and distinctively human fashion. He was even found, during some spells, to keep a radio nearby so that he could periodically listen to it and check to see if his language recognition ability was returning to normal. Given the

complexity of Brother John's behaviour, the question we must ask ourselves is this: does the case of Brother John refute the claim that thinking is done in natural language? The answer is no. The available evidence leads us to believe that Brother John had lost the ability to produce and comprehend language, but this does not show that natural language was not involved in his thought processes. During his life, Brother John's mind was structured through the use of natural language. It could be that during his seizures, Brother John's linguistic systems were still working on linguistic structures at an *unconscious* level, thus giving rise to his intelligent behaviour. This possibility is consistent with what we know about the case. We have two major sources of evidence for Brother John's language deficit:

(1) external observations of his behaviour;

and,

(2) his introspective recollections of his experience during the aphasic spells.

Now, the external observations of Brother John's behaviour show that he could not (or would not) speak, and that he appeared unable to understand the speech of others. Examining the observations made of Brother John is a useful way of assessing the extent of his public language deficit. But observations alone cannot tell us anything about what was going on in his mind, and so they cannot be used to decide whether or not natural language was involved in his thinking. Brother John's introspective testimony may be more authoritative than external observations, but it too cannot be used to rule out the possibility that natural language was involved in guiding his behaviour during his aphasic episodes. Introspection can only tell Brother John about

what he was conscious of. The fact that he could not *consciously* use or understand natural language does not show that natural language was not in use at a lower level, below his awareness. So, although it is *possible* that Brother John's aphasic episodes disabled every mental structure related to language, his introspective testimony does not provide strong evidence to support this conclusion. It is also possible that a system involved in vocal production and comprehension had been 'taken off-line' while the rest of the system was intact and influencing his behaviour. Chomsky's (1995) model of the language faculty can be used to support this second possibility. According to Chomsky, the language system is comprised of at least three separable sub-systems. There are performance systems, which include output production and input reception; and there is a cognitive system, which is like a linguistic database that is accessed by the input and output systems. Chomsky believes that cases in which input or output systems have been impaired while the cognitive system remains intact, reveal "the kind of modular structure expected in any complex biological system" (Chomsky 1995: pg 12). It is plausible to explain Brother John's behaviour in these terms. On Chomsky's model, Brother John's linguistic database (the thinking part of his brain) was still at work, making inferences and producing intelligent behaviour while his linguistic input/output module was malfunctioning.

Cases of aphasia have been used to show that thought is possible without natural language, but I am not convinced. Models such as Chomsky's offer us an alternative explanation which keeps the importance of language intact. Whether or not Chomsky's model is correct is a question that may one day be discovered as brain science progresses. In the meantime, it serves the purpose of showing that the

existence of aphasia does not provide very strong support for the independence of thought from natural language.

4: A Concluding Thought about Thought

In this chapter I have highlighted and considered arguments that have been put forward against the view that thinking is done in natural language. Some of these arguments rely on intuition, while others consider cases in which evidence seems to show that thought is possible in the absence of language. The intuitive cases point to features of our experience like the 'tip-of-the-tongue' phenomenon. Such experiences may make it seem as if our thoughts exist independently of our expressions of them, but I have shown that the phenomena can be answered in ways that do not threaten the importance of natural language. Perhaps the 'tip-of-the-tongue' experience occurs as the result of a momentary failure of memory. Or perhaps it occurs because of a failure in forming a thought. The point I have tried to make is that such phenomena do not need to be answered by appealing to translation difficulties between mentalese and natural language.

Einstein's claim that his thinking involved image manipulation has been used to show that thought is independent of natural language. However, Einstein's testimony does not show that thinking must therefore be carried out in mentalese. Nor does it show that thinking does not involve natural language. At most Einstein's thoughts show that it is possible for image manipulation to somehow be involved in thinking. I have therefore allowed that thinking may sometimes involve visual imagery, but I maintain that natural language is necessary for us to reason about visual

images. Furthermore, I have claimed that abstract concepts such as 'might', 'cause' or 'equal', cannot be thought about solely in visual terms and must therefore be thought about in sentences of mentalese or natural language. It is reasonable to believe that human thinking involves a combination of words and images, and Einstein's thoughts must have involved such a combination.

Cases of aphasia have been used to show that thinking can occur without language, but I have shown that the evidence does not give us sufficient reason to accept that conclusion. In the case of Brother John, all we can be certain of is the fact that he appeared to be unable to use or understand language, while at the same time he retained his ability to cope with the environment. Even Brother John's own testimony does not give us reason to believe that thinking can occur without language. We can account for Brother John's behaviour and his introspective reports by appealing to Chomsky's model of the language system. Chomsky's model shows us the possibility that linguistic structures can continue to work in the brain at an unconscious level even during an aphasic episode.

I have shown that the arguments against a cognitive role for natural language can be answered. The work in this chapter, coupled with the evidence presented in chapter 3, give us good reason to accept that natural language is the language of thought.

5

A Direction for Future Thought

1: The Review

The goal of this project has been to establish that human thinking should be understood as being constituted by processes involving sentences of natural language. This chapter will consist in a review of the preceding chapters and some speculation as to where further research might take us.

This thesis has taken the form of a debate between two types of language of thought theories. On one side of the debate I looked at the traditional language of thought theories of Fodor and Pinker, which require the existence of an innate and hidden processing language known as mentalese. On the other side of the debate I suggested that for reasons of economy, it is easier to suppose that the natural language we speak is our language of thought. Despite their differences, these two views agree on one thing -- that is, human thinking involves the manipulation of sentential structures and this activity is specified by a finite collection of rules. In chapter 2, I outlined the language of thought hypothesis. One of the main points of chapter 2 was to show that the language of thought hypothesis can be used to explain the productive and systematic nature of human thought. I then went on to show that there are some problems with supposing that the language of thought is innate and is hidden. These problems were not unanswerable by the mentalese theorist, but they provided sufficient reason to take a look at the simpler possibility that natural language might be the language of thought.

In chapter 3, I started out by describing the two roles that language plays. These are the *communicative* role and the *cognitive* role. When we speak of the *communicative* role, we are referring to language's obvious role in facilitating the transmission of thoughts to other individuals. Whether thoughts are constructed from mentalese or natural language, they are communicated via sentences of natural language. On the other hand, saying that language has a *cognitive* role is to say that language plays an important role in thinking. The position that I have been defending is that human thoughts are constructed from sentences of natural language and so I accept that language plays a cognitive role. My reasons for holding this belief were considered in chapter 3. First, I looked at the evidence of introspection, which leads us to believe that our thinking is carried out in natural language. When we introspectively examine our thought processes, we notice that our thinking seems to be constituted by sentences of natural language. We have a constant chatter of inner speech occurring in our heads, and most of the time this chatter is not intended to be communicated to other individuals. It seems to serve the primary purpose of determining and guiding our day to day activities. Mentalese supporters have a number of possible responses to my claims, but I have shown that these responses lead to a cluttered view of the mind. One of my driving forces is to keep explanations of the mind as simple as possible and not postulate the existence of a hidden language of thought if natural language can do the work itself.

To further empathize the role natural language plays in human thinking, I considered the case of Helen Keller, who did not acquire language in the usual fashion. Because she became blind and deaf within the first two years of her life, Keller did not learn a language until she was about 7 years old. It was at this time that

she learned to communicate through a unique language of touch. An examination of Keller's writings reveals a critical point that has remarkable implications for a discussion on the role of natural language in thought. Keller states that before she acquired language, she did not exist as a unified thinking thing. She says that her thoughts and behaviour were instinctual and contained no creativity or intelligent will. Furthermore, Keller tells us that after learning language she began to experience a phenomenon similar to the inner speech that we experience. Interestingly, Keller's inner speech involved the imaginary sensations of words being spelled out into her hand, and this became the way in which she entertained private thoughts. Keller's testimony thus provides us with further reason to suppose that thinking is done in natural language.

In the next section of chapter 3, I considered the productive and systematic nature of human thought. My claim was that human thoughts do not become productive or systematic until a certain stage of development is reached. This stage, which was categorized by Jean Piaget, is known as the object permanence stage and marks the time at which human infants acquire the ability to think symbolically. This is important because, as I explained in chapter 2, symbol manipulation is required for thoughts to be productive and systematic.

Chapter 4 contained a number of arguments that have been put forward to refute the claim that human thinking is carried out in natural language. These arguments pointed to features of human experience that seem to be incompatible with the idea that thinking involves sentences of natural language. Phenomena such as the tip-of-the-tongue experience, for example, were used to show that thoughts exist in mentalese and are independent of their natural language interpretations. Furthermore,

the very fact that we can learn new words and coin new words is supposed to show that our thoughts must be different from the words of natural language. I disagreed with these arguments and attempted to explain the phenomena in terms of natural language.

In chapter 4, I also looked at the mentalese supporters' claim that Einstein's thought processes involved the manipulation of visual imagery. I suggested that if Einstein's thoughts were entirely visual, then the role of mentalese would be called into question just as much as the role of natural language. Despite this point, I accepted that Einstein's thoughts involved visual imagery but I showed that the visual images must have been accompanied by sentences of natural language. The idea behind this claim was that certain ideas cannot be expressed entirely visually. Concepts such as 'tomorrow', 'all', and 'perhaps' are examples of ideas that cannot be described or understood in visual imagery. This is because such concepts have no appearance and visual imagery can (obviously) only represent the appearance of things.

The last section of chapter 4 examined the evidence from aphasia. Aphasia is a neurological disorder that has the effect of disabling certain elements of a person's language facility. There are differing degrees of effect, and in some extreme cases a person's entire linguistic system can appear to be disabled. Despite the loss of language, however, many people who suffer from aphasia continue to behave in intelligent, distinctly human ways. Theorists such as Pinker conclude from this evidence that thought and natural language must be distinct, and therefore there is a language of thought that is not natural language. Contrary to Pinker's belief, I think that there are other ways to explain the presence of intelligent behaviour in the

absence of natural language. To support this idea, I suggested that in cases of extreme aphasia, the sufferer has lost the ability to produce and understand sentences of natural language, but that person's behaviour is still being influenced by natural language structures at an unconscious level. My reason for suggesting this came primarily from the thought that during a lifetime, a person's mental structures are shaped and influenced by exposure to natural language. It seems reasonable to suppose that those mental structures may continue to exist despite the loss of overt language ability. In light of this possibility, I concluded that aphasia does not provide good evidence for the existence of a language of thought that is distinct from natural language.

2: The Conclusion

The work done in this thesis has served the dual purpose of: **1)** establishing that there is a language of thought, while **2)** calling into question which language the language of thought actually is. The debate between mentalese and natural language is not an easy one to resolve. As this paper has shown, any attempts to provide evidence that natural language is the language of thought can be answered by the mentalese theorist in a very simple way. All they have to do is reiterate their claim that mentalese does its work behind the scenes and provides a foundation upon which natural language can operate. Mentalese supporters may state that because we are only aware of natural language, it is natural for us to believe that natural language is the language of thought. But they would also remind us that this does not show that mentalese does not exist. I think the mentalese supporter would be correct in making this claim, but remember also that they cannot show us that mentalese *does* exist. Any attempt by the mentalese theorist to provide reasons for which we must believe that thinking is

done in mentalese rather than natural language can be answered, thus keeping alive the possibility that natural language is the language of thought. So, the tension arises because we have two candidates for the language of thought, both of which offer significant explanatory power. I have decided that the solution to the tension between these rivaling theories involves an appeal to economy -- namely, go for the simplest solution. Now, mentalese is hidden and its origins are extremely difficult to discover. Natural language, on the other hand, is observable and its origins can be discovered. If we decide that mentalese is the language of thought, our picture of the mind is one which involves the mind expending resources on supporting and translating between two languages -- mentalese and natural language. But both of these languages can explain human thinking. For reasons of simplicity I have been lead to the conclusion that natural language is the best candidate for the language of thought. I have not shown that there is no mentalese, but I have provided good reasons for considering the possibility that natural language is the language of thought. The conclusion I draw from the above thesis is that we should put aside the mentalese theories and concentrate on furthering our knowledge of natural language and its influence on the structure of the mind.

3: Expanding the mind beyond the confines of the biological brain

Accepting that natural language is constitutively involved in human thinking opens up a number of avenues for future research. Such research, which has been pointed to by Andy Clark (1997) and Merlin Donald (1991), might explore our ability to 'off-load' thoughts into other minds, or onto pieces of paper, or into electronic storage/retrieval systems. This ability is directly related to our possession of language and results in

the expansion of our minds beyond the confines of our biological brains. Consider the simple case of what happens when a shop assistant is cashing up at the end of the day. He/she may count the cash first before adding up the eftpos receipts and cheques. Rather than keeping the cash total stored in the brain, he/she will often write the number down on a piece of paper and then add up the eftpos receipts and cheques. This number is also written down and is added to the first number yielding a result that can be written down for the accountant. Now, this process could have been carried out entirely in the brain, but it was easier to use an external tool. This external tool had the effect of serving as a short term memory enhancement and became an integral part of the mind in solving an everyday problem. We do this sort of thing all the time and it has resulted in a huge expansion of our intellectual capacities. But is it plausible to consider a computer, or piece of paper, to truly be a temporary expansion of the mind? Or is the mind always a self contained system that simply knows how to use a few tools? Is there any real difference between storing data (or thoughts) in neural patterns, or storing them in external devices and enabling a data link via the stream of photons entering the eyes? And consider, if our minds are expanded into the non-neural environment, then where do our minds end? In some ways the answers to these questions may seems trivial. Indeed, asking these questions may seem absurd. However, a little reflection shows that they are valid questions that apply only to humans because of our language ability. These questions are important because they relate directly to the most curious of entities, 'the Self'. I think one of the first steps that further research should take would involve an analysis of the extent to which the Self is dependent upon language. Preliminary questions might determine whether or not conscious experience is possible in the absence of language, and these

questions could provoke further questions about the nature of the Self. Then further thought about the possible expansion of the mind into the environment could include questions about how such an expansion might affect the Self and consciousness. Does the self expand as the mind expands into the environment? Are human conscious experiences altered as the mind expands?

The issues discussed in this, and subsequent projects may have profound implications for the way we think about what it is to be human. For they show us that we are far more than mere biological devices. We have been enhanced through the installation of linguistic technology. This technology has restructured our minds and augmented our innate capacities by providing a mechanism through which we can assimilate the thoughts of others, and by allowing us to integrate socially constructed software into our biologically constructed programming. With language we can represent abstractions, think about time, and engage in deductive reasoning. Such abilities escape the other animals and sets us apart from them; for unlike the non-linguistic creatures, it is with language that we think.

References

Armstrong, D.M. (1973). *Belief, Truth and Knowledge.* Cambridge University Press, 1973.

Baillargeon, Renee. (1991). "Reasoning about the height and location of a hidden object" in *Cognition 38.* Pages 13-42.

Blackburn, Simon. (1984). *Spreading the Word.* Oxford University Press, 1984.

Braddon-Mitchell, David., and Jackson, Frank. (1996). *Philosophy of Mind and Cognition.* Blackwell Publishers, 1996.

Calvin, William H. (1996). *The Cerebral Code.* A Bradford Book, The MIT Press, Cambridge, Massachusetts, 1996.

Calvin, William H. (1996). *How Brains Think: Evolving Intelligence, Then and Now.* New York: Basic Books, A Division of HarperCollins Publishers, Inc. 1996.

Calvin, William H. (1998). "The Emergence of Intelligence" in *Scientific American, Vol 9, Number 4, 1998.* Pages 44-51.

Carruthers, Peter. (1996). *Language, Thought and Consciousness.* Cambridge University Press, 1996.

Chomsky, Noam. (1995). "Language and Nature" in *Mind, Vol 104.* Oxford University Press, 1995. Pages 1-61

Clark, Andy. (1990). "Connectionist Minds" in *Connectionism, Debates on Psychological Explanation.* Edited by Cynthia Macdonald and Graham Macdonald, Blackwell Publishers Ltd, 1995. Pages 339-356.

Clark, Andy. (1993). *Associative Engines.* Cambridge, MA: MIT Press, 1993.

Clark, Andy. (1997). *Being There: Putting Brain, Body, and the World Together Again.* A Bradford book, MIT Press, 1997.

Clark, Andy. (1998). "Magic Words: How Language Augments Human Computation" in *Language and Thought: Interdisciplinary themes.* Edited by Peter Carruthers and Jill Boucher, UK: Cambridge University Press, 1998.

Cole, David. (1997). "Hearing Yourself Think: natural language, inner speech and thought", http://www.d.umn.edu/~dcole/hearthot.htm. Downloaded 1 September 1999.

Cole, David. (1999). "I Don't Think So: Pinker on the Mentalese Monopoly" in *Philosophical Psychology, Vol 12, Number 3, September 1999.* Pages 283-296.

Corballis, Michael C. (1999). "The Gestural Origins of Language" in *American Scientist, March - April 1999.* Pages 138-145.

Cosmides, Leda. (1989). "The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason?" in *Cognition, Volume 31, 1989.* Pages 187-276.

Cowie, Fiona. (1999). *What's Within? Nativism Reconsidered.* Oxford University Press, 1999.

Crane, Tim. (1995). *The Mechanical Mind.* Penguin Books, 1995.

Curtiss, Susan (1977). *Genie: A Psycholinguistic Study of a Modern-day "Wild Child".* New York: Academic Press.

Davidson, Donald. (1975). "Thought and Talk" in his *Inquiries into Truth and Interpretation.* Oxford University Press, 1984. Pages 155-171.

Dawkins, Richard. (1976). *The Selfish Gene.* Oxford University Press, 1976.

Dennett, Daniel C. (1978). "A Cure for the Common Code" in his *Brainstorms.* The Harvester Press Limited, 1979. Pages 90-108.

Dennett, Daniel C. (1991a). *Consciousness Explained.* Penguin Books, 1993.

Dennett, Daniel C. (1991b). "Mother Nature Versus The Walking Encyclopedia" in *Philosophy and Connectionist Theory.* Edited by Ramsey, W., Stich, S., and Rumelhart, D. Laurence Erlbaum Associates, Inc., Publishers, 1991.

Dennett, Daniel C. (1995). *Darwin's Dangerous Idea.* Penguin Books, 1996.

Dennett, Daniel C. (1996). *Kinds of Minds.* Weidenfeld & Nicolson, 1996.

Dennett, Daniel C. (1998). "Reflections on Language and Mind" in *Language and Thought: Interdisciplinary themes.* Edited by Peter Carruthers and Jill Boucher, UK: Cambridge University Press, 1998.

Devitt, Michael & Sterelny, Kim. (1987). *Language and Reality.* Basil Blackwel Ltd, 1987.

Donald, Merlin. (1991). *Origins of the Modern Mind.* Harvard University Press, 1991.

Dretske, Fred. (1981). *Knowledge and the Flow of Information.* Cambridge, Mass.: MIT Press, 1981.

Fodor, Jerry. (1975). *The Language of Thought.* Thomas Y. Crowell Company, 1975.

Fodor, Jerry (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind.* Bradford/MIT, 1987.

Fodor, Jerry. (1990a). "Fodor's Guide to Mental Representation" in his *A Theory of Content and Other Essays.* Cambridge, Mass.: MIT Press, 1990. Chapter 1.

Fodor, Jerry. (1990b). "A Theory of Content, 1" in his *A Theory of Content and Other Essays.* Cambridge, Mass.: MIT Press, 1990. Chapter 3 (Pages 51-87).

Freud, Sigmund. (1940). "Some Elementary Lessons" in *The Complete Psychological Works of Sigmund Freud, Volume XXIII (1937-1939).* Pages 281-286.

Garfield, Jay. (1997). "Mentalese Not Spoken Here" in *Philosophical Psychology 10:4, 1997.* Pages 413-435.

Garfield, Jay. (1999). "Thought as Language: A Metaphor Too Far". Paper presented as part of the Erskine lecture series at the Department of Philosophy, University of Canterbury, 18th May 1999.

Gauker, Christopher. (1992). "The Lockean Theory of Communication" in *Nous, Volume XXVI, Number 3, September 1992.* Blackwell Publishers, Cambridge MA & Oxford UK, Pages 303-324.

Godfrey-Smith, Peter. (1996). *Complexity and the Function of Mind in Nature.* Cambridge University Press, 1996.

Harman, Gilbert. (1973). *Thought.* Princeton, N.J.: Princeton University Press, 1973.

Jaynes, J. (1976). *The Origin of Consciousness in the Breakdown of the Bicameral Mind.* Boston: Houghton Mifflin.

Kaye, Lawrence. (1998). "Another Linguistic Turn? Review of Language, Thought and Consciousness", http://psyche.cs.monash.edu.au/v4/psyche-4-02-kaye.html. Downloaded 17 July 1999.

Kaye, Lawrence. (1999). "The Language of Thought", http://www.uniroma3.it/kant/field/lot.html. Downloaded 17 July 1999.

Keller, Helen. (1909). *The World I Live In.* London: Hodder and Stoughton.

Laurence, Stephen and Eric Margolis (1997). "Regress Arguments Against the Language of Thought," in *Analysis, Vol. 57, No. 1.*

Lecours, Andre.R., and Yves, Joanette. (1980). "Linguistic and Other Aspects of Paroxysmal Aphasia" in *Brain and Language, 10, 1980.* Pages 1-23.

Malson, Lucien. (1964). *Wolf Children.* Translated by Edmund Fawcett, Peter Ayrton and Joan White. London: NLB, 1972.

Piaget, Jean. (1932). *The Language and Thought of the Child.* New York: Harcourt Brace.

Piaget, Jean. (1953). *The Origin of Intelligence in the Child.* Translated by Margaret Cook. London: Routledge & Kegan Paul, 1953.

Piaget, Jean. (1983). "Piaget's Theory" in *Handbook of Child Psychology, 1983.* Edited by Paul H. Mussen, Volume editor: Willian Kessen. Pages 103-128.

Pinker, Steven. (1994). *The Language Instinct.* Penguin Books, 1995.

Plotkin, Henry. (1997). *Evolution in Mind.* Penguin Books, 1998.

Sapir, E. (1921). *Language.* Harcourt, Brace and Company, Inc, 1921.

Schaller, Susan. (1991). *A Man Without Words.* Forward by Oliver Sacks. USA: University of California Press, 1991.

Searle, John. (1992). *The Rediscovery of the Mind.* Cambridge, Mass.: MIT Press, 1992.

Siegler, Robert S. (1986). *Childrens Thinking.* Second edition, USA: Prentice-Hall Inc, 1991.

Sterelny, Kim. (1990). *The Representational Theory of Mind.* Cambridge, Mass. Basil Blackwell, Inc, 1991.

Sterelny, Kim. (1999). "Dawkin's Bulldog" in *Philosophy and Phenomenological Research, Vol. LIX, No. 1, March 1999.*

Summerfield, Donna M. (1996). "Fitting versus Tracking: Wittgenstein on Representation" in *The Cambridge Companion to Wittgenstein.* Edited by Hans Sluga and David G. Stern, USA: Press Syndicate of the University of Cambridge, 1996.

Tye, Michael. (1991). *The Imagery Debate.* MIT Press, A Bradford Book, 1991.

Wade, Carole., and Tavris, Carol. (1993). *Psychology.* Third edition, USA: HarperCollins College Publishers, 1993.

Whitmeyer, Joseph M. (1998). "On the Relationship Between Memes and Genes: A Critique of Dennett" in *Biology and Philosophy 13, 1998.* Pages 187-204.

Wittgenstein, Ludwig. (1953). *Philosophical Investigations.* Edited by G.E.M. Anscombe and R. Rhees, Trans. G.E.M. Anscombe. Oxford: Blackwell.

Wittgenstein, Ludwig. (1958). *The Blue and Brown Books.* Oxford: Blackwell, 2nd edition, 1960.

Acknowledgments

First and foremost I wish to thank Derek Browne for supervising my efforts in constructing this thesis. He kept me on track and provided valuable criticism and encouragement. I would also like to thank the staff and graduate students in the Department of Philosophy at the University of Canterbury. They have all (directly and indirectly) influenced my thoughts -- especially Vanessa Scholes and Alex Wynne, with whom I shared an office and had many interesting conversations. My partner, Pamella Jenks, deserves my deep gratitude. She has been patient, kind, and supportive of my efforts.

Finally, I wish to offer my sincere thanks to my parents, Margaret and Winston Silby. They have offered me emotional and financial support, and have encouraged me through every stage of my life.

Brent Silby 3rd February 2000.