

Contextualist Solutions to Cartesian Scepticism

This Thesis is submitted in part requirement for the degree of
MPhil in Philosophy

Miguel Angel Fernandez

University College London

April 2001




ProQuest Number: 10016107

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10016107

Published by ProQuest LLC(2016). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

In this thesis I argue that the contextualist solution to scepticism about knowledge of the external world fails to accomplish its own goals.

In chapter I I identify the sceptical argument that contextualists focus on and describe the problem posed by that argument that contextualists want to solve. I point out that the goals of the contextualist solution to the sceptical problem are (1) dissolve the paradox of intuitions generated by the sceptical argument previously identified: we have intuitions for the premises and against the conclusion of that (valid) argument; and (2) justify the claim that relative to the epistemic standards operative in everyday contexts people can have knowledge of the world.

In section 1 of chapter II I expound and discuss the semantic theory of knowledge ascriptions within which contextualists pursue the two goals of their solution to the sceptical problem. According to that semantic theory the truth conditions of knowledge ascriptions are sensitive to facts about the context in which the ascriptions are made. In section 2 of chapter II I explain how this semantic theory is used to dissolve the sceptical paradox.

In chapter III I argue that there's evidence for thinking that, in effect, the semantic content of knowledge ascriptions is sensitive to contextual facts. I defend this idea against two objections. But then I argue that the particular manner in which contextualists construe that context sensitivity doesn't give them what they expected from it, *viz.* the resources to accomplish goals (1) and (2) of their solution to the sceptical problem.

Table of Contents

Contextualist Solutions to Cartesian Scepticism

Acknowledgements.....	5
Introduction.....	6
Chapter I. Posing the Sceptical Problem.....	9
1. Scepticism by hypotheses.....	9
2. Detachment and the “objective” conditions for knowledge.....	19
3. Why the sceptical condition cannot be fulfilled?.....	24
a. Stroud’s argument.....	24
b. Nozick’s argument.....	28
4. The sceptical problem posed.....	30
Chapter II. The Contextualist Solution to Scepticism.....	32
1. Contextualist semantics for knowledge ascriptions.....	32
a. The context sensitivity of knowledge ascriptions.....	32
b. How to understand <i>strength in epistemic position</i>	36
i. Strength in terms of ruling alternatives out: Cohen and Lewis.....	37
ii. Strength in terms of truth-tracking: DeRose.....	40
c. Contextual facts that affect the content of knowledge ascriptions: Subject facts vs ascriber facts.....	42
2. Solving the sceptical paradox.....	46
Chapter III. Problems for the Contextualist solution to scepticism.....	53
1. Is the contextualist solution of the paradox internally coherent?.....	53
2. Semantic or pragmatic phenomenon?.....	60
a. The pragmatic objection.....	61
b. Unger’s hypothesis of semantic relativity.....	66
c. It is a semantic phenomenon.....	71
3. Elusive knowledge and impossible theory.....	74
a. Elusive knowledge?.....	74
b. Impossible theory.....	78

Conclusion.....85
Bibliography.....87

Acknowledgements

During my MPhil studies I enjoyed the scholarship 136126 granted by the Consejo Nacional de Ciencia y Tecnologia (CONACYT) of Mexico. The writing of this thesis was possible thanks to this financial support.

I also want to express here my gratitude to my thesis supervisor at UCL Michael G. F. Martin.

Miguel Angel Fernandez
London, April 2001

Introduction

Scepticism about knowledge of the external world has proved to be a particularly intractable problem. Many apparently powerful attempts to refute the sceptic, or otherwise dissolve the problem he seems to pose, have been shown to be defective in different ways.¹ In this thesis I expound and examine a relatively new attempt to solve the problem of scepticism, the so called contextualist solution. I try to show that this solution to scepticism piles up with the other failed attempts to silence the sceptic's doubts. In what follows I'll describe what I do in this thesis in order to reach this conclusion.

The first problem in confronting the sceptic is to determine what is exactly the problem that he is posing, or if he is posing several problems, identify which one we will focus on. The first chapter is devoted to clearly identify the sceptical problem that contextualists tackle. Contextualists are interested in a form of scepticism, that I call Cartesian scepticism, according to which nobody knows anything about the external world. I explain how Cartesian scepticism can be obtained via the exploitation of undetectable cognitively disabling states such as *I'm a brain in a vat* or *there's an Evil Demon that always makes me believe false things*. I call such hypothetical states "sceptical hypotheses". I explain how there are two different kinds of sceptical hypotheses and argue that they can be exploited in two different arguments by the sceptic. Then I isolate one of those sceptical arguments as the one on which contextualists focus their attention. I explain how the sceptic can defend the premises of this argument, the purpose of such explanation is to show how powerful intuitions arising from the concept of knowledge militate in favour of the sceptic's premises. Despite the existence of those intuitions for the premises of the argument its conclusion seems to us totally unacceptable, we also have entrenched intuitions connected with the concept of knowledge that lead us to reject the conclusion of the sceptical argument. This conflict of intuitions sets up the problem contextualists want to tackle: it is a paradox of intuitions, some of them pull us to accept the premises and others pull us to reject the conclusion of a perfectly valid argument.

Contextualists develop their solution of the sceptical paradox within the

¹For a record of failed attempts to solve the sceptical problem cf. Stroud, 1984 and Wright, 1985, section II, pp. 439-449.

framework of a semantic theory about knowledge ascriptions. The first part of chapter II is devoted to the exposition of this semantic theory. The core idea of the theory is that the truth conditions or semantic content of knowledge ascriptions are sensitive to certain facts about the context in which the ascriptions are made. Most contextualists underestimate the task of giving a detailed account of the kind of context sensitivity allegedly present in knowledge ascriptions, I stress how important this point is for the success of the contextualist solution and try to discuss some aspects of the problems involved in providing the required detailed account.

For contextualists the truth conditions of a knowledge ascription set up how strong a subject's epistemic position with respect to a proposition p must be to count as a knower of p . Thus what in their view changes contextually are the standards for strength in epistemic position. After exploring different interpretations of epistemic strength I go on to present the contextualist view about what are the contextual facts that determine the standards for how strong an epistemic position must be. Their view is that linguistic facts of *mentioning* possibilities of error or *saying* that someone knows something are enough to produce contextual shifts that modify the epistemic standards operative in a given conversational context. Later on I call this view the "simply say it" view.

In the second part of chapter II I explain how contextualists utilise the semantic theory about knowledge ascriptions expounded in the first half of the chapter to solve the sceptical paradox. The essence of the solution can be stated briefly. They say that our conflicting intuitions concerning the premises and the conclusion of the sceptical argument arise only because the conversational context shifts throughout the exposition of the sceptical argument in virtue of the sceptic's mentioning certain possibilities of error. The contextual shift so produced brings about changes in the operative standards for how strong an epistemic position must be. We evaluate the premises and the conclusion of the sceptical argument relative to different standards for knowledge and that's why we end up having conflicting intuitions about the truth value of those propositions. Our paradoxical intuitions would disappear if the conversational context could be kept fixed throughout the sceptical argument.

The contextualist solution of the sceptical paradox can be regarded as an attempt to show that our concept of knowledge is not incoherent, that it doesn't

produce paradox. I draw attention to the fact that in addition to this claim about the non-paradoxical or coherent character of the concept of knowledge, contextualists also claim that this concept has true applications, albeit only relative to what they call “everyday contexts”; that is to say, they also claim that in everyday context people can have knowledge of the world. This means that contextualists expect that their semantic theory about the context sensitivity of knowledge ascriptions will give them the resources to solve the sceptical paradox and to justify their claim that in everyday contexts knowledge of the world is possible.

In chapter III I assess two aspects of contextualism. First, I inquire what can be the evidence for the context sensitivity of knowledge ascriptions being a semantic matter that affects their truth conditions. In the first two sections of chapter III I respond to two objections to contextualism that in different ways try to undermine the idea that there is a context sensitivity of knowledge ascriptions which is a semantic phenomenon. In responding to such objections I suggest that there’s evidence for granting to the contextualist the semantic context sensitivity of knowledge ascriptions. After granting this to contextualists I explore the question whether their particular view about what contextual facts determine the semantic content of knowledge ascriptions (what I call the “simply say it” view) gives them the two results they expect from it: solving the sceptical paradox and allowing them to justify their claim that relative to the epistemic standards of everyday contexts people have knowledge of the world. I argue that the “simply say it” view doesn’t give them these two results and for this reason I conclude that the contextualist solution to scepticism fails to accomplish its own goals.

Chapter I Posing the sceptical problem

In this work I shall call ‘Cartesian scepticism’ the following epistemological view:

Nobody knows anything about the external or physical or material world.

This is a negative global verdict about the possibility of human knowledge of the world. I call it ‘Cartesian’ because there is agreement that the basic considerations that lead to such a view are to be found in Descartes’ *First Meditation*.¹ There is agreement also in that something must be wrong in an argument that leads to that unacceptable view. However, there is large disagreement about *what are* the crucial doctrines and ideas involved in the argument and about *which of them are wrong* and why. The two issues are closely related, different representations of the sceptical reasoning can motivate different diagnoses of it, i.e. different views about what is wrong with that reasoning. If two representations of the argument highlight different principles or doctrines as involved in arriving at the conclusion, then those principles or doctrines will become the target of different diagnoses of the argument.

In the next two chapters I’ll expound and examine one diagnosis of scepticism called “contextualism”. In the present chapter I’m going to pose the sceptical reasoning in a way that the issues addressed by the contextualist response to it can be easily identified. I will not discuss the question whether there are other issues involved in the reasoning that should be examined in a thoroughgoing diagnosis of scepticism, if there are such other issues it may be that they don’t get highlighted in the present description of the sceptical problem.

1. Scepticism by hypotheses.

The Cartesian sceptic aims to reach his global verdict about knowledge of the world by somehow employing possibilities or hypotheses such that if they were the case that would preclude everybody from knowing anything about the world. In

¹ Cf. Descartes, 1641, pp. 13-15.

posing the sceptical problem we have to clarify what the nature of those hypotheses is and how the sceptic utilises them in reaching his conclusion.

The hypotheses the sceptic uses may preclude or be incompatible with knowledge in two different ways.² Since knowledge requires true belief plus something else (on the nature of which there's much dispute) hypotheses may be incompatible with knowledge by being incompatible with true belief or by being incompatible with the third-component in knowledge.³ I'll call hypotheses of the first type 'Hypotheses-T' (or H_T) and of the second type 'Hypotheses-K' (or H_K). There are many (perhaps infinitely many) hypotheses-T and K that are incompatible with each of the many *particular* things we think we know, but most of those hypotheses are proposition-specific in the sense that they are incompatible only with one proposition or with knowledge of that proposition. Given that the sceptic aims at a *global* conclusion about knowledge of the world he needs hypotheses that are not proposition-specific but incompatible with *any* proposition about the world or with *any* knowledge about the world.

For simplicity I shall concentrate the discussion on two hypotheses the sceptic can use for his purposes:

H_T : There is an Evil Demon that always makes me believe false things.

H_K : I'm a brain in a vat being stimulated by a super-computer to have all the experiences I have.⁴

It is pretty clear that H_T is incompatible with the truth of any belief one may have about the world, and in this way it is incompatible with any knowledge of the world. But it is not equally clear how H_K is incompatible with the third-component of any knowledge of the world, therefore I shall explain why it is incompatible with that component.

² Hereafter 'knowledge' means, unless otherwise indicated, 'knowledge of the external world'.

³ Philosophers that write on scepticism rarely stress this difference between sceptical hypotheses. Among the few exceptions are Wright, 1985, pp. 430-31; Stroud, 1984, pp. 28-29; Williams, 1991, pp. 332-333.

⁴ It must be noticed that the brain-in-vat hypothesis is also incompatible with the *truth* of many propositions I believe (e.g. that I have hands, etc.) but *not* with the truth of *all* the propositions I believe (not with e.g. there are trees in London). Most philosophers that take the brain-in-vat hypothesis as the paradigmatic example of sceptical hypothesis interpret it as a hypothesis-T. But so interpreted it *cannot* serve the *global* sceptical purpose, in order to serve such purpose it must be interpreted as a hypothesis-K: the brain-in-vat hypothesis is incompatible with the third-component of *any* knowledge of the world.

There has been a great debate about what's needed to turn a true belief into knowledge. We cannot look at the different views about that here, but we can confidently say that at least one thing in addition to true belief is required for knowledge, namely, an appropriate, knowledge-yielding connection between our beliefs and the obtaining of the worldly facts those beliefs are about.⁵ Intuitive reflection shows the necessity of such a connection for knowledge. Suppose I'm actually a brain in a vat. In that case my believing, say, that there are trees in London would be completely disconnected from the fact that happens to make my belief true; I would believe what I do not because things are the way I believe them to be, but simply because I happen to be artificially stimulated to believe them to be that way. This means that if I'm a brain in a vat what the world is like becomes irrelevant to what beliefs about the world I form. A subject of whom this is true cannot count as a knower of the world precisely because his beliefs about the world have lost some important dependence-connection to the world.

It's not easy to spell out the kind of general connection that there must exist between our beliefs and the facts we believe to be the case in order for such beliefs to count as knowledge. Some philosophers have spelt it out in terms of causal relations,⁶ others in terms of a dependence specifiable by subjunctive conditionals. Robert Nozick,⁷ for example, explicates the appropriate connection between beliefs and facts in terms of the following pair of conditionals:⁸

- If p were false, S wouldn't believe that p .
- If p were true, S would believe that p .

Hypothesis H_K is incompatible with the truth of these two conditionals. If I were being stimulated to believe that p is true I would believe that it is true even if p were actually false. And if I were being stimulated to believe that p is false I wouldn't

⁵ In saying that this is *at least one* of the things that are required for knowledge in addition to true belief, I want to leave open the question whether something else (perhaps an internalist condition) is needed for knowledge. Taking position with respect to this question would require a great deal of argument that would lead us away from the main topic of this chapter.

⁶ Cf. Goldman, 1967

⁷ Nozick, 1981, p. 176.

⁸ I'm ignoring here some complication in Nozick's account concerning methods of belief formation. But this omission is harmless for the point I making at this stage. The issue about methods becomes important below (section 3. b) when I discuss Nozick's argument that sceptical hypotheses cannot be known to be false.

believe it's true even if it were actually true. Nozick uses his account of knowledge in terms of those conditionals as the basis of a response to Cartesian scepticism. Nozick's response to scepticism is a complicated issue that I won't discuss here. I mention Nozick's account only as an example of the kind of connection between beliefs and facts that is required for knowledge and with which H_K is incompatible.

It is an open question whether Nozick's account of the knowledge-yielding connection between beliefs and facts is correct, but independently of such a general account we can gather intuitively compelling reasons that H_K is incompatible with our gaining knowledge of the world through any of the sources by which we think we can obtain it: perception, testimony, inductive inference; precisely because in each case H_K is incompatible with a certain connection between beliefs and facts. For a perceptual belief to count as knowledge it has to have a certain aetiology that cannot exist if we are brains in vats.⁹ A belief formed on the basis of the testimony of others counts as knowledge only if we *really* are getting the testimony of someone that is appropriately related to the facts he is giving testimony of, but if we are merely being stimulated to have the experience as if we were getting such a testimony, then we are not really getting it and so we are not fulfilling that necessary condition to obtain testimony-based knowledge of the world. When we form inductively the belief that all As are Cs, we think such a belief amounts to knowledge only if all, or most of, the previously *observed* As have been Cs. But if we have been merely stimulated to have the experiences as if all As we have encountered had been Cs, then we lack what would turn our belief into knowledge, namely, *observations* that all As examined have been Cs.¹⁰

⁹ Cf. Wright, 1985, p. 431 and 1991, p. 91. Wright prefers the hypothesis of dreaming, which is another hypothesis-K. For the purpose of explaining what's the belief-fact connection with which a hypothesis-K is incompatible, it doesn't matter which hypothesis-K we choose: all of them are incompatible with that connection in the same way.

¹⁰ Some of the knowledge we think we have about the world is based on our memory, so it might be said that memory is missing in the list of sources of knowledge I'm discussing. However, memory is not a *source* of knowledge but a capacity to *retain* and *recall* knowledge we have obtained in some other way. Anyway it is an interesting question whether the hypothesis-K is incompatible with remembering, that is, with retaining and recalling a piece of knowledge. The question gets complicated because different forms of memory seem to involve different conditions that must be fulfilled for genuinely remember: for example, the conditions for autobiographical memory seem to be different from the conditions to remember something that I didn't personally witness. However, it has been argued (cf. Martin & Deutscher, 1966) that in general a condition for S to remember that *p* is that the way in which S acquired the information that *p* somehow must be causally operative in producing S's present belief that *p*. Now, if S is being stimulated to have the state as if he were remembering that *p* then his present belief that *p* is entirely the causal product of the *present* brain stimulation, the way in which S acquired the information that *p* (even if at some time in the past S wasn't a brain in a vat and

Now that we have seen how the hypothesis-K is incompatible with any knowledge of the world, let's turn to the question how the sceptic uses his hypotheses to obtain his global conclusion.

Since sceptical hypotheses are *incompatible* (albeit in two different ways) with any knowledge of the world, the sceptic thinks that a necessary condition for *any* such knowledge is to eliminate the sceptical hypotheses. Let's represent any sceptical hypotheses as 'H', any proposition about the world as '*p*' and the knowledge operator as 'K'. We can say that the sceptic thinks that $K\sim H$ is a necessary condition for Kp , or in other words that Kp entails $K\sim H$. This way of looking at the sceptic's reasoning has led lots of philosophers (among which most of the contextualists are included) to think that the Cartesian sceptic is committed to the Closure Principle according to which if one knows one thing and knows that some other thing is a logical consequence of the first thing then one knows this second thing:

If Kp and $K[p$ entails $q]$, then Kq

The relevant instance of Closure that the sceptic is thought to be using is

If Kp and $K[p$ entails $\sim H]$, then $K\sim H$.

There is something powerfully compelling in the Closure Principle, something like it must be correct if knowledge by deduction is to be possible. However, the principle needs some refinements, for example, if 'K' means 'S knows...' the principle seems liable to easy counterexamples. Suppose that S knows that *p* and that *p* entails *q* but in a particular occasion S fails to put the two pieces of knowledge together, in that case this person might fail to infer *q* and then fail to actually know that *q*. This type of counterexample might perhaps be avoided if 'K' is interpreted in the weaker form 'S is in a position to know...'. Some philosophers have tried to construct other types of counterexamples to Closure in an attempt to rebut the very first step leading to

appropriately acquired the information that *p*) plays no causal role in S's present belief that *p*. For this reason the hypothesis-K is incompatible with a necessary condition for remembering something.

Cartesian scepticism.¹¹ But for many people giving up Closure looks as too high a price to pay for blocking the sceptical argument.¹² Contextualists eschew this conflict by not giving up Closure. They are friends of Closure, they want to show how we can retain Closure, and the powerful intuitions supporting it, and still provide a response to the sceptic.¹³

Once the relevant instance of the Closure Principle is in place the sceptic argues that nobody can fulfil the condition $K\sim H$ and then concludes that for that reason nobody know anything about the world:

1. If Kp and $K[p$ entails $\sim H]$, then $K\sim H$.
2. $K[p$ entails $\sim H]$
3. $\sim K\sim H$
4. $\sim Kp$ (From 1, 3 by Modus Tollens)¹⁴

We must notice that this argument can *only* exploit hypotheses-T to obtain the global sceptical conclusion. The reason is that the argument requires the sceptical hypothesis to be incompatible with *the truth of p*; in other words, it requires that p entails $\sim H$, as the second conjunct of the antecedent of premise (1) says. Hypothesis-T is incompatible with the truth of *any* proposition we believe about the world, *any* such proposition entails $\sim H_T$, and then the above argument can be used to exploit hypothesis-T with global sceptical results. But the same is not true of hypothesis-K. This hypothesis is incompatible with the truth of some but *not all* propositions about the world. For example, H_K is incompatible with the proposition that I have hands, but not with the proposition that there are trees in London. H_K is compatible with many propositions about the world, and so none of those propositions entail $\sim H_K$. Consequently the argument above cannot exploit hypothesis-K with global sceptical results. To get the global conclusion he wants through that argument the Cartesian sceptic has to restrict the range of hypotheses to only hypotheses-T, I shall explicitly mark this restriction replacing 'H' by ' H_T ' in the argument:

¹¹ Cf Dretske, 1970 and Nozick, 1981 for classic anti-Closure arguments.

¹² Cf. Vogel, 1990 for a defence of Closure against different putative counter examples.

¹³ Cf. below Chapter II.2.

¹⁴ The number of philosophers that think that *this* is the argument for Cartesian scepticism is huge, some examples are: Dretske, 1970, p. 1011; Nozick, 1981, p. 204; Cohen, 1987, p.18 and 1999, p. 62; Klein, 1995, p. 213; with minor variations De Rose, 1995, p. 1; Schiffer, 1996, p. 317; Sosa, 1999, p.143. They largely overlook the fact that *different arguments* are needed to exploit the two different types of sceptical hypotheses; see below.

CP

CP1. If Kp and $K[p \text{ entails } \sim H_T]$, then $K\sim H_T$

CP2. $K[p \text{ entails } \sim H_T]$

CP3. $\sim K\sim H_T$

CP4. $\sim Kp$ (From CP1, CP3 by Modus Tollens)

What would be an argument for Cartesian scepticism that exploits hypothesis-K? Since H_K is incompatible with knowledge in general not by being incompatible with the truth of the propositions we believe but with a *knowledge*-yielding connection between beliefs and facts, it may be suggested that we can couch the first premise of the new argument after the model of (CP1) just by replacing the idea that p entails the negation of hypothesis-T with the idea that *knowledge* that p entails the negation of hypothesis-K. The first premise of the new argument would then be:

E1: If Kp and $K[Kp \text{ entails } \sim H_K]$, then $K\sim H_K$

Despite its similarity this is *not* an instance of the Closure Principle, unlike premise (CP1) in the argument CP which *is* an instance of Closure. E1 is an instance of a principle that doesn't say, as Closure does, that if p is known the 'K' operator penetrates to all known logical consequences of p ; but that if p is known the 'K' operator penetrates to all known consequences of *knowing* that p . And the sets defined by 'known consequences of p ' and 'known consequences of knowing that p ' are different sets of propositions.

The full argument for Cartesian scepticism exploiting hypothesis-K would be this:

E

E1. If Kp and $K[Kp \text{ entails } \sim H_K]$, then $K\sim H_K$.

E2. $K[Kp \text{ entails } \sim H_K]$

E3. $\sim K\sim H_K$

/ $\sim Kp$ (From E1, E3 by Modus Tollens).

Argument E, which exploits hypothesis-K, doesn't employ an instance of Closure as its first premise, and this distinguishes it from argument CP. It must be noted at this point that part of the compelling character of argument CP comes from the fact that its first premise is justified as an instance of the intuitively compelling Closure Principle; since premise E1 cannot be justified in the same way argument E lacks the persuasiveness that CP borrows from Closure.¹⁵ But perhaps an argument for the exploitation of hypothesis-K with the global results the sceptic wants can be constructed using *an instance of the Closure Principle*, provided the substitutions for variable 'p' in the principle are restricted in a certain way and a certain auxiliary premise is annexed to the simple argument CP. The idea would be the following. The instance of Closure used in CP is:

If Kp and $K[p$ entails $\sim H_T$], then $K\sim H_T$

Hypothesis-K seemed unsuitable to take the place of 'H_T' because not all propositions *about the world* entail the negation of hypothesis-K. We saw that what entails the negation of hypothesis-K is any proposition asserting *knowledge* of the world. This suggests that if we stipulate that only proposition of the form ' Kp ' can take the place of the propositional variable 'p' in the above instance of the Closure Principle, then we could use it to exploit Hypothesis-K. The resulting instance of Closure would be:

If $K(Kp)$ and $K[(Kp)$ entails $\sim H_K$], then $K\sim H_K$,

and the argument for the exploitation of hypothesis-K would be:

1. If $K(Kp)$ and $K[(Kp)$ entails $\sim H_K$], then $K\sim H_K$.
2. $K[(Kp)$ entails $\sim H_K$]
3. $\sim K\sim H_K$
4. $\sim K(Kp)$ (From 1, 3 by Modus Tollens)

¹⁵One might think that E1 *is* compelling despite not being an instance of Closure. But I think that E1 may seem compelling only because it is easily confused with a genuine instance of Closure, viz. premise (CP1*) of argument CP* below. If this is so the apparent compelling character of E1 would be borrowed from its being confused with an instance of Closure.

This argument uses an instance of Closure to exploit hypothesis-K but falls short of the conclusion the sceptic wants. If successful the argument only establishes that one doesn't know that one knows p , but the sceptic wants the first level conclusion that one doesn't know that p . In order to obtain this latter conclusion an additional premise should have to be added in that argument, namely:

Kp entails KKp

This premise embodies the so called KK-principle according to which if one knows something then one knows that one knows it. Using the KK-principle, premise (4) in the argument above and Modus Tollens the sceptic would obtain the first level conclusion he wants: that one doesn't know that p . The full argument using an instance of Closure to exploit Hypothesis-K would then have to appeal to the KK-principle:

CP*

CP1*. If $K(Kp)$ and $K[(Kp)$ entails $\sim H_K$], then $K\sim H_K$.

CP2*. $K[Kp$ entails $\sim H_K$]

CP3*. $\sim K\sim H_K$

CP4*. $\sim K(Kp)$ (From CP1*, CP3* by Modus Tollens)

CP5*. If Kp , then KKp

CP6*. $\sim Kp$ (From CP4*, CP5* by Modus Tollens).

We have two sceptical arguments, CP and CP*, that use instances of Closure to exploit hypothesis-T and hypothesis-K respectively; although CP* requires an auxiliary premise (i.e. the KK-principle), that CP doesn't, in order to obtain the sceptical conclusion.¹⁶

In what follows I'll not complicate matters trying to take into account these two sceptical arguments. I'll concentrate on the argument of the simple form CP. The main reason for doing this is that contextualists, whose response to Cartesian scepticism will be the topic of the next two chapters, typically focus their attention on

¹⁶ As far as I know Crispin Wright, 1985, pp. 430-438, is the only philosopher that has emphasised that the two types of sceptical hypotheses are exploited in different sceptical arguments, although the arguments in which according to him the hypotheses get exploited don't have the structure of the arguments I've described here. As regards this I must repeat that I'm representing the sceptical argument(s) in such a way that the issues addressed by *contextualists* can easily be highlighted.

that simple argument CP. In order to present and assess their views it will be useful then to circumscribe the discussion to argument CP.¹⁷

I have argued that the simple argument CP can only exploit hypothesis-T with the global consequences the sceptic wants. Surprisingly contextualists usually choose a hypothesis-K, as we'll see in the next chapters, to fill in the argument (their favourite one is brain-in-vathood). They choose a suitable proposition about the world (like, *I have hands*) so that its truth entails $\sim H_K$ and the second conjunct of the antecedent of (CP1) comes out true. They don't seem to realise that, as we have seen above, a hypothesis-K cannot be exploited in the simple argument CP *with the global results the sceptic wants*: simply because H_K is not incompatible with *all* propositions about the world. This means that contextualists choose the wrong combination of hypothesis and sceptical argument. But this can be easily corrected by replacing the H_K they choose by a suitable H_T exploitable with the desired sceptical global results through the simple argument CP, which is the argument that contextualists want to focus on. Nevertheless, in order to avoid confusions I'll preserve the examples contextualists use. I just wanted to point out that strictly speaking the hypothesis that must figure in the argument contextualists address must be a hypothesis-T not a hypothesis-K .

After having examined the nature of sceptical hypotheses and the arguments in which they are exploited, it's time to see how the sceptic argues for the crucial stages of the argument we are going to focus on from now onwards:

CP

CP1. If Kp and $K[p$ entails $\sim H_T]$, then $K\sim H_T$

CP2. $K[p$ entails $\sim H_T]$

CP3. $\sim K\sim H_T$

CP4. $\sim Kp$ (From CP1, CP3 by Modus Tollens)

If it is true that a necessary condition of knowing anything whatsoever about the world is that we know hypothesis-T not to obtain, as (CP1) entails; and if we can

¹⁷ Consequently I'm going to ignore how contextualists would exactly respond to argument CP*. I think that the basic contextualist's ideas deployed in the response to CP can, without many modifications, be extrapolated into a contextualist treatment of CP*. However, word-limit prevents me from going into this issue.

never fulfil that condition, as (CP3) says, then Cartesian scepticism would be unavoidable. In the following sections I'll examine what the sceptic can say for premises (CP1) and (CP3) of the argument CP.

2. Detachment and the "objective" conditions for knowledge.

How can the sceptic defend the idea that eliminating hypothesis-T is a condition for knowledge? The sceptic might argue that in imposing that condition on knowledge he is just applying a certain pattern used in everyday life to establish necessary conditions for knowledge. The pattern in question is this: Suppose I'm a botanist and found a cactus with a feature F that makes me classify it as X. I know, and you know, there is other variety Y of cactus with feature F. So you point out to me: "If it is a Y, then you don't know it is a X. You should eliminate the possibility that it is a Y if you are going to know it's a X." We recognise that the possibility that the cactus is a Y is incompatible with me knowing that it is an X, and that recognition seems enough to convince us that I must know it is not a Y if I am going to know that it is a X. It seems that we undoubtedly follow this pattern in everyday life: if we recognise that a certain possibility or hypothesis is incompatible with our knowing that p , we conclude that a necessary condition for knowing that p is to eliminate the possibility in question. The sceptic could argue that in introducing the elimination of his hypothesis-T as a necessary condition for any knowledge of the world he just makes a particular application of that pattern.

But do we really follow *that* pattern in everyday life? Is it enough to make the elimination of a certain hypothesis a necessary condition for knowledge that we recognise it is incompatible with our knowledge? It seems it isn't, and precisely the sceptical hypothesis exemplifies it. We recognise hypothesis-T as incompatible with any knowledge of the world, but that is not enough in itself to make us accept that we must know it doesn't obtain to know anything about the world. Something else beyond the mere recognition of the hypothesis as incompatible with knowledge is needed to turn the elimination of that possibility into a necessary condition for knowledge. If we look at the everyday example of the cacti we find that if the elimination of the hypothesis that the cactus is a Y is going to be a necessary condition for knowing that it is an X, there must be some specific reason in those

circumstances for thinking that the hypothesis might obtain. For example, we must at least believe that cacti Y can grow there, that it wouldn't be terribly difficult to bring a Y to this environment, etc. Without these reasons we wouldn't think the elimination of that hypothesis is a necessary condition for knowledge, we would consider it too remote to deserve being taken into account. The plausibility of this pattern of everyday practice suggests that what makes the elimination of a certain hypothesis a necessary condition for knowledge is not just that we see it is incompatible with knowledge, but also that we have some specific reasons for thinking that the hypothesis might really obtain.¹⁸

The sceptical condition doesn't fit into that pattern: we never have specific reasons for thinking hypothesis-T might be actual. And even if we can perversely imagine situations in which such reasons exist it would still be true that in the overwhelming majority of cases we lack such reasons, and therefore, from the point of view of the pattern mentioned above, the elimination of the sceptical hypothesis cannot count as a *general* condition for knowledge. The fact that the sceptic cannot argue that in imposing his condition on knowledge he is just extending further a pattern we follow in everyday life, leaves us with an important question: how would the sceptic defend his condition if he cannot support it in everyday practice, and moreover, given that it is in flagrant conflict with that practice? This question raises complex issues about what the relation should be between *any* philosophical view about the conditions for knowledge and the conditions we actually impose in everyday epistemic practices. If philosophers (sceptics and non-sceptics) want to find out which are the epistemic conditions for knowing something, is reflection on ordinary practice the only source of information for that inquiry? The sceptic would have to answer 'no', since reflection upon ordinary practice speaks against his condition for knowledge; but what else can we reflect upon in order to decide which are the epistemic conditions for knowledge?¹⁹ I cannot discuss these questions here. What I'm going to do is sketch a possible interpretation of the pattern(s) we find in everyday epistemic practices that makes such practices

¹⁸ Cf. Austin, 1946, for a defence of the idea that to make the elimination of a possibility a condition for knowing something, we must have a special reason for thinking the possibility might be actual.

¹⁹ Cf. Kaplan, 2000 for a discussion of what should one reflect upon to settle the question of what are the epistemic conditions for knowledge, and especially of what's the place of reflection upon ordinary practice in settling that question.

consistent with the sceptic's condition for knowledge. Such interpretation of everyday practice doesn't show the sceptic to be correct, it just shows that what goes on in ordinary practice does not straightforwardly exhibit the sceptic to be wrong.

The sceptic accepts that in everyday life we never think of his condition for knowledge. But he never intended to be faithful to what we do in everyday life, he wanted to step back from everything we accept as true in ordinary practice and ask, from a detached point of view, if we "really" know any of those things about the world we ordinarily accept as true.²⁰ He, of course, comes up with a negative answer. But what matters here is that he finds feasible the idea of adopting a detached point of view from which the "objective" or "real" conditions for knowledge can be discovered.

The sceptic thinks that whether someone knows something of the world is an objective matter of fact that holds or not quite independently of anybody thinking, believing or saying that it obtains or not. We can be wrong in ascribing knowledge to someone, as we can be wrong about any other objective matter of fact. This conception of knowledge opens up the way to the idea that the procedures we follow in everyday life in settling the question whether someone knows something, may not embody "what, with detachment, we can be brought to acknowledge [as] the full conditions for knowledge".²¹ The sceptic might say that this is exactly what happens with our everyday epistemic practices. He might say that our epistemic practices of ascribing knowledge to people are not guided by what, strictly speaking and from a point of view detached of the contexts where those practices operate, are the objective and full conditions that must be fulfilled to know something of the world. The conditions he claims to be objectively necessary for knowledge are, from his point of view, clouded by the practical concerns and interests in everyday life. Barry Stroud has explained this way in which the sceptic conceives of the relation between our epistemic practices and the reflective or detached stance from which he assesses our knowledge by means of an analogy he adapts from an example of Thompson Clarke's.²²

²⁰ Cf. Descartes, 1641, p. 12.

²¹ Stroud, 1984, p. 71.

²² Cf. Stroud, 1984, pp. 67-74. Stroud's story is inspired on Clarke, 1972, pp. 759 ff.

Think of a group of plane spotters trained by their military manuals to perform the task of distinguishing planes of type E from planes of type F. The manuals say that any plane with characteristics xyw is an E, and any plane with characteristics xyz is an F. However, it turns out that there is a third type of plane, the Gs, that also have features xyz and so are indistinguishable from the Fs for the spotters. We can imagine several reasons why the manuals do not mention planes of type G. Perhaps the people in charge of writing the manuals do not know of the existence of the Gs, or perhaps they know of them but it is not useful for the war effort that the spotters be able to distinguish Fs from Gs, maybe because both types of planes play the same function for the enemy.

The practice of distinguishing Es from Fs in accordance with the manuals may be very useful from a practical point of view and even help crucially in winning the war, nevertheless abstracting that practice from the practical contexts in which it is embedded we can see that it is defective from a purely epistemological point of view. Whenever a spotter says: "Yes, I know the last plane was an F", he speaks falsely even if the plane was actually an F and he followed scrupulously the instructions in his manual. The spotter doesn't know the plane was an F because he classifies it as an F in virtue of its having features xyz , but if it's for that the plane *might* have been a G. We, observers of the position the spotters are objectively in, realise that someone cannot know that a plane with xyz is an F unless he is able to eliminate the possibility that it is a G. We see that this is, objectively speaking, a necessary condition for the spotters' knowledge. Since they cannot fulfil that condition they can't know that a given plane is an F.

The conditions for classifying a plane as an F stated in the training manuals do not supply conditions for the *truth* of claims of the form: 'I know that's an F', this is shown by the fact that someone may fulfil the conditions stated in the manuals and still don't know that the plane is an F. However, nobody would question the usefulness and practical appropriateness of the spotters' classificatory practice, in spite of being epistemically defective in the manner described. We may say then that what the manuals really set up are conditions for the *appropriate* and useful assertion of claims like: 'I know that's an F', but not the objective and real *truth* conditions for such claims.

From the point of view of the spotters they speak the truth when they say ‘I know that’s an F’ after identifying the plane in accordance with their manual, but from a detached position we can see that they are wrong, they are confusing the conditions for *appropriate* assertion of ‘I know that’s an F’ (which they fulfil), with the *truth* conditions of that assertion (which they don’t fulfil).

The example of the plane spotters illustrates two points that the sceptic can use to construct a certain conception of everyday epistemic practices. First, the example stresses the fact that the conditions for *appropriate assertion* of a claim and the *truth conditions* for that claim are two different things; and secondly, that conditions for appropriate assertion may be confused with truth conditions.

With these ideas in hand the sceptic may argue that our everyday epistemic practices only embody the conditions for appropriate assertion of ascriptions of knowledge like: ‘S knows that *p*’, but not the objective and full truth conditions of such ascriptions. We have learnt to ascribe knowledge to people when certain conditions obtain, for example, when people can eliminate hypotheses or possibilities for which there are specific reasons to think they might be actual (as the spotters learned to classify a plane as an F when it has the features mentioned in the manual). To make those ascriptions in the conditions we make them is a very useful and appropriate thing to do given the interests and purposes we have when we interact socially with our fellows every day (as the classificatory practice of the spotters is very useful and appropriate given the interest of winning the war). However, the conditions for the appropriateness of a knowledge ascription are something different from its truth conditions. Thus the sceptic may say that the conditions for knowledge displayed in everyday practice can be construed as conditions for appropriateness of knowledge claims, without necessarily being conditions for the truth of such claims.

This construal of everyday epistemic practices provides the sceptic with a defence of his condition for knowledge against the charge that it is not a condition for the truth of knowledge claims *because* it deviates from the patterns we follow in everyday life. Undoubtedly eliminating hypothesis-T is not among the conditions for *appropriately claiming* that we know something in a given circumstance, that is because the only hypotheses whose elimination is among such conditions are those for which there are specific reasons for thinking they may be actual. This is what ordinary practice shows. But that is consistent with the elimination of hypothesis-T

being a condition for the *truth* of knowledge claims. Our reluctance to accept in everyday life that the elimination of hypothesis-T is a necessary condition for the truth of knowledge claims may be a consequence of our confusing the conditions of appropriate assertability with truth conditions of knowledge claims.²³ Given this construal of everyday epistemic practices, the sceptic's position that eliminating hypothesis-T is a condition for the truth of any knowledge claim about the world cannot be shown to be mistaken simply for deviating from any conditions displayed in ordinary practice as conditions for knowing something.

The conception of the relation between our everyday epistemic practices and the sceptic's condition for knowledge sketched above does not lead by itself to the negative verdict that nobody knows anything about the world. The crucial stage in the way to this conclusion is that the necessary condition for knowledge that according to the sceptic is overlooked in our everyday epistemic practices, turns out to be such that *cannot* be fulfilled. In the next section I'll examine two prominent arguments for the view that the sceptical condition cannot be met indeed.

3. Why the sceptical condition for knowledge cannot be fulfilled?

There are two powerful arguments in the literature for the view that the sceptical condition on knowledge of the world (*viz.* knowing $\sim H_T$) cannot be met, one is due to Barry Stroud and the other to Robert Nozick. I'll expound them in this order.

a. Stroud's argument.

The sceptic claims that all the evidence available to us at any time doesn't enable us to know $\sim H_T$, i.e. that there's no Evil Demon that always makes me believe the false.²⁴ Since it is quite implausible to say that we can know $\sim H_T$ by a

²³ This kind of defence of the sceptical condition for knowledge provides the basis for a criticism to the contextualist's solution to scepticism. I expound and examine that criticism in Chapter III. 2.

²⁴ Stroud's argument is couched in terms of the hypothesis of dreaming (*cf.* Stroud, 1984, p. 21). I make the suitable (mainly verbal) adjustments to the argument required to couch it in terms of the hypothesis of the Evil Demon.

single strike of a priori intuition or by a single strike of perceptual experience, it seems that in order to know that there's no Evil Demon there should be a procedure the execution of which would determine that there's no Evil Demon. But no such procedure can possibly exist. For us to use a procedure to establish and so come to know that there's no Evil Demon we would have to know that we are really executing a procedure with such results and not merely being deceived by an Evil Demon into thinking that we are executing it. But this means that the procedure would be of use for knowing that there's no Evil Demon only if we *antecedently* knew that we are not merely being deceived by an Evil Demon into thinking that we are executing a procedure that establishes that there's no Evil Demon. However, that there's no Evil Demon is precisely what the execution of the procedure was meant to establish. Therefore the procedure would be successful only if what it is meant to show can be established on different grounds.

One could think that what this argument shows is that we require a second order procedure for knowing that the execution of the first order procedure really establishes that there's no Evil Demon. But the same problem would arise for the second order procedure: for it to give us knowledge that the execution of the first order procedure really establishes that there's no Evil Demon, we would need a third procedure for telling that the execution of the second order procedure really establishes that the first order procedure is successful and that we are not merely being deceived by an Evil Demon into thinking that the execution of such second order procedure establishes that the first order procedure is successful. For a similar reason we would need a fourth procedure to establish that the execution of the third order procedure really establishes that the second order procedure is successful, and so on. This leads to an infinite regress of procedures, we would always need a procedure for establishing that the execution of the previous procedure really has successful results and that we're not merely being deceived by an Evil Demon into thinking that such a procedure has those successful results, and so we would never be in a position to know that there's no Evil Demon.²⁵

It may be thought that what generates this infinite regress of procedures is the especially problematic nature of what the procedures are meant to establish: that

²⁵ Cf. Stroud, 1984, pp. 21-23.

there's no Evil Demon. It is because with respect to each procedure, if it is to have any probative force, we need to know that we are *not merely being deceived by an Evil Demon* into thinking that its execution has the results we want from it, that we always need a higher order procedure. However, it seems that the same type of regress can be generated without the Evil Demon hypothesis being involved. What we need to generate a regress of that type is just the following pair of principles:

- i. If a fact of the matter X cannot be known to obtain by a single strike of sensory experience or a priori intuition, then in order to know that X obtains we need to carry out a procedure.
- ii. Executing a procedure cannot give us knowledge that a matter of fact obtains if we don't know independently that we have properly executed the procedure.²⁶

The qualification *know independently* in (ii) makes it explicit that the execution of a procedure doesn't itself give us knowledge that the procedure has been properly executed, we need a distinct procedure for knowing that the first procedure has been properly executed. This idea seems compelling. Suppose that you are measuring the distance between two objects with an electronic device, if you don't know whether you are using the device properly, intuitively, you don't know, on the basis of using it, what the distance between the objects is. To know the distance you need to know you are using the device properly, but you don't know this latter thing just by using the device as you do; to know that you are using the device properly you need a distinct, further procedure.

Principles (i) and (ii) generate a sceptical infinite regress like the one we find in Stroud's argument in the following way. Let X_0 be a fact of the matter that can't be known to obtain by a single strike of sensory experience or a priori intuition such that $K[X_0]$ ²⁷ is a necessary condition for knowledge of the world. Principles (i) and (ii) entail that nobody is ever in a position to satisfy any $K[X_0]$ necessary condition for knowledge. To satisfy such a condition we need to carry out a procedure P_0 that gives us knowledge that X_0 (as (i) demands). However, for the execution of P_0 to give us

²⁶Principle (ii) is a version of what Crispin Wright calls the "Proper Execution Principle" (Wright, 1991, pp. 99-100). Wright formulates the Proper Execution Principle in terms of procedures to earn *warrant* for believing a proposition, whereas principle (ii) is formulated in terms of procedures to *know* a proposition. Despite this difference Wright's principle can also be used to generate a sceptical regress like the one I describe next in the text.

²⁷ That is: 'Knowing that $[X_0]$ obtains'.

knowledge that X_0 we need to know independently that we have carried out P_0 properly (as (ii) demands). This means that we need a distinct procedure P_1 to know that the fact of the matter X_1 : “we carried out P_0 properly” obtains. However, for the execution of P_1 to give us knowledge that the fact of the matter X_1 obtains we need to know independently that we carried out P_1 properly (as (ii) dictates). This means that we need a distinct procedure P_2 to know that the fact of the matter X_2 : “we carried out P_1 properly” obtains. However, for the execution of P_2 to give us knowledge that the fact of the matter X_2 obtains again we need to know independently that we carried out P_2 properly. This means that we need a distinct procedure P_3 to know that the fact of the matter X_3 : “we carried out P_2 properly” obtains. But for the execution of P_3 to give us knowledge that the fact of the matter X_3 obtains we need to know independently that we carried out P_3 properly. This means that we need a further procedure P_4 to know that the fact of the matter X_4 : “we carried out P_3 properly” obtains.....It’s obvious we have here an infinite regress: in trying to satisfy a necessary condition for knowledge of the world $K[X_j]$ we will always need to know through a certain procedure that a certain matter of fact obtains, and for that procedure to give us knowledge of that matter of fact we will have to know via other procedure that the previous procedure was properly executed, and we will need a further procedure to know that this latter procedure was properly executed and so on. As a consequence we will never be in a position to meet any condition $K[X_j]$ for knowledge of the world.

We have arrived at this strong conclusion solely on the basis of principles (i) and (ii), independently of the Evil Demon hypothesis. This suggests that we can generate a regress like the one we find in Stroud’s argument for the conclusion that the sceptical condition for knowledge cannot be met (i.e. that we don’t know $\sim H_T$) without appealing to anything peculiar or specific about *the sceptical condition for knowledge*, but only appealing to the general principles (i) and (ii).

I will present now other argument for the view that the sceptical condition cannot be satisfied that unlike Stroud’s gives to the specificity of the Cartesian condition a central role to play in arriving at that conclusion.

b. Nozick's argument

Robert Nozick has argued that knowing that p requires that one's belief that p *tracks the truth* of p . He defines the relation of truth tracking in terms of the following pair of counterfactual conditionals:

- a. If p were false and S were to use method M to arrive at a believe whether (or not) p , then S wouldn't believe, via M , that p .
- b. If p were true, and S were to use method M to arrive at a believe whether (or not) p , then S would believe, via M , that p .²⁸

Nozick maintains that one cannot know that hypothesis-T is false because one's belief that $\sim H_T$ doesn't track truth, it doesn't satisfy condition (a) of the truth tracking relation. I'll explain why.

In the first place we have to ask ourselves what method could someone use to form a belief as to whether $\sim H_T$.²⁹ We must notice that the method must be available both when $\sim H_T$ is false (as (a) implies) and when $\sim H_T$ is true (as (b) implies). In other words, the method must be available both when there is an Evil Demon deceiving me and when there's no Evil Demon deceiving me. What methods for forming beliefs about the world are available in these two cases? The usual methods we employ to form beliefs about the world: perception, testimony, inductive inference, are not available in both cases. If there's an Evil Demon *deceiving* me into thinking that I'm perceiving, receiving testimony or observing exceptionless correlations between A s and B s then I'm not perceiving, receiving testimony or making a warranted inductive inference. Any method the operation of which requires a certain connection with the facts my beliefs are about wouldn't be available in the case where there's an Evil Demon; the reason is that in that case the Demon would be deceiving me into thinking that the appropriate connection required for the operation of the methods obtains when really it *doesn't*. Once this is noted it appears that the only method for forming beliefs about the world that would be available both when there is and when there's not an Evil Demon is *how things (phenomenally) seem to me*, since things appearing to me a certain way doesn't require the existence of *any connection with*

²⁸ Nozick, 1981, p. 179.

²⁹ Nozick presents his argument using the brain in vat hypothesis, I shall make the suitable adjustments to present the argument using the Evil Demon hypothesis.

the facts that appear to me that way. Things can seem to me the *same way* if I'm perceiving them to be that way, or if I'm dreaming them to be that way or if a Demon is deceiving me by making things appear to me that way. This is why the method of how things seem to me is available both when there is and where there isn't an Evil Demon.

The problem is that the method of how things seem to me does not enable me to know $\sim H_T$ (i.e. that there's no Evil Demon) because it doesn't enable me to satisfy condition (a) of truth tracking. The relevant instance of that condition would say:

- a. If $\sim H_T$ were false and S were to use how things seem to him to arrive at a belief whether (or not) $\sim H_T$, then S wouldn't believe, via how things seem to him, that $\sim H_T$.

It is false that if there were an Evil Demon (i.e. if $\sim H_T$ were false) and I were to use the method of how things seem to me to arrive at a belief whether $\sim H_T$, I wouldn't believe via that method that there's no Evil Demon. The reason is that if there were an Evil Demon deceiving me by making it seem to me as if there were no Evil Demon, things would seem to me the same as they would if there really were not an Evil Demon. This means that the fact that there is or the fact that there isn't an Evil Demon makes no difference to how things seem to me, and for this reason makes no difference to what I come to believe on the basis of how things seem to me. Either if there is or if there is not an Evil Demon I would form the same beliefs (all of which would be false in the case where there's Demon); in particular, in both cases I would believe that there's *no* Evil Demon. Then, if there were an Evil Demon I would still believe, via how things seem to me, that there's no Evil Demon, which means that my belief that there isn't an Evil Demon doesn't satisfy the truth tracking condition (a) and for this reason doesn't count as knowledge.

One may think that Nozick explanation of why we don't know $\sim H_T$ depends too heavily on his theoretically loaded analysis of knowledge in terms of counterfactual conditionals (a) and (b). To the extent that such analysis is not free from difficulties Nozick's explanation of why $\sim H_T$ is unknowable would also seem inconclusive or problematic. However, I think Nozick's explanation rests on a quite intuitive and conclusive basis, which is the following: If the only resource I have to tell whether there is or there isn't an Evil Demon is the way things seem to me, and

things would seem the *same* to me if there were and if there weren't an Evil Demon, then I don't have the resources to *differentiate* the two cases. If I cannot differentiate one state from the other, I'm not in a position to know that I'm in one and not in the other, it's precisely because of this that on the basis of how things seem to me I cannot know that there's no Evil Demon:

[If the hypothesis of the Evil Demon were true], everything would seem the same to us. There is no way we can know it is not happening for there is no way we could tell if it were happening; and if it were happening we would believe exactly what we do now – in particular, we still would believe that it was not. For this reason, we feel, and correctly, that we don't know –how could we?- that it is not happening to us.³⁰

4. The sceptical problem posed

What is then the sceptical problem? We have seen that different arguments, *viz.* CP and CP*, can be used to reach Cartesian scepticism. Even after focusing on CP we have found different arguments for the crucial premise that we don't know there's no Evil Demon, those arguments also trace different routes to the sceptical conclusion. The diversity of arguments that can be employed by the sceptic suggests that there's no single problem about Cartesian scepticism. It looks more like a cloud of different problems posed by different lines of argument.

However, we can identify at least two *general* problems posed by the different lines of sceptical reasoning. The most direct sceptical challenge is how to *demonstrate* that the sceptical conclusion is false and we actually have knowledge of the world. But there is another challenge. Consider argument CP, this is a perfectly *valid argument* for whose premises, as we have seen, powerful intuitions can be brought in, and nevertheless we have also deep intuitions against its conclusion. We have intuitions attached to our concept of knowledge that pull us to accept the premises of CP and other intuitions attached to that concept that pull us to accept the negation of the conclusion of CP. This means that different intuitions arising from our concept of knowledge pull us to accept a set of mutually inconsistent propositions; *viz.* the premises of CP and the negation of its conclusion. Our concept

³⁰ Nozick, 1981, p. 201.

of knowledge seems paradoxical, incoherent. Here we have a different sceptical challenge, the challenge of showing that our concept of knowledge is not incoherent.

The official view of contextualists is that their solution to Cartesian scepticism is an attempt to meet the second, not the first, of the two challenges distinguished above. They say that they want to show that our concept of knowledge doesn't sink us into paradox, but that they don't try to demonstrate to the sceptic that his conclusion is false. Here is a clear statement of this view:

We need to be very clear about the nature of the enterprise in which we are engaged... we are faced with a paradox. We are inclined to accept each member of a set of propositions we know to be inconsistent. What we seek is a way out of the paradox –a resolution of our inconsistent inclinations. It is not a constraint on the acceptability of a resolution that it appeal to the sceptic....To resolve the paradox is not to demonstrate to the sceptic that we know. Rather it is to demonstrate to ourselves that we can claim to know without paradox.³¹

The official contextualist enterprise is to show that our concept of knowledge is not paradoxical. However, as we will see, they not just want to say that our concept of knowledge is coherent and paradox-free, they also want to say that in some contexts that concept has true applications, that in some contexts we know things. But in order to justify this latter claim it seems that contextualists must face the most direct sceptical challenge: show that the sceptical conclusion is false. This sets up our itinerary for the next two chapters. How do contextualists attempt to solve the paradox created by the sceptical argument CP? Do they succeed in this enterprise? Even if they succeed in this enterprise, can they justify their further claim that relative to some contexts we know things? In order to go into all these questions we first have to expound the theoretical machinery contextualists devise to attempt to solve the sceptical paradox. That is the task of the first part of next chapter.

³¹ Cohen, 1988, p. 113.

Chapter II

The Contextualist Solution to Scepticism

1. The contextualist semantics for knowledge ascriptions.

Contemporary contextualism is a theory about the truth conditions or semantic content of knowledge sentences, it says that “the truth conditions of sentences of the form “S knows that p ” or “S does not know that p ” vary in certain ways according to the context in which the sentences are uttered”.¹ According to contextualists when we say that S knows that p we are saying that S’s epistemic position is strong enough to make him qualify as a knower of p , thus the truth conditions or content of a knowledge sentence define how good S’s epistemic position must be to qualify as a knower of p .² The claim that the truth conditions of a knowledge sentence vary according to the context in which it is used is then equivalent to the claim that how good S’s epistemic position must be for him to know that p varies contextually.

There are at least two components in this theory that must be explained and defended. First, we need an account of the context sensitivity of knowledge sentences, and it is reasonable to expect that such account will make it clear which factors in a context of utterance affect the truth conditions of knowledge ascriptions and how exactly those factors operate in affecting the truth conditions.³ Secondly, we are told that the truth conditions of a knowledge ascription consist in how strong S’s epistemic position must be for him to know, so we need an explanation of what “strength of epistemic position” is. I will expound these components of the contextualist theory in that order.

a. The context-sensitivity of knowledge ascriptions.

There are different ways in which the semantic content of the token of an expression may be sensitive to certain features of the context in which the token is

¹ De Rose, 1992, p. 914. Cf. also Cohen, 1988, pp. 98-9 and 1999, p. 60.

² Cf. De Rose, 1995, pp. 4, 29; 1996, p. 195.

³ By ‘knowledge ascription’ I understand a particular *utterance* of a knowledge sentence like ‘S knows that p ’ or ‘S doesn’t know that p ’. Knowledge sentences are not themselves knowledge ascriptions, but they can be used to make knowledge ascriptions. Utterances of the *same* knowledge sentence in suitably different contexts constitute *different* knowledge ascriptions.

used. Context may help to figure out what language is being spoken and in that way help to determine what is being said in using an expression. Some expressions belonging to a determinate language can be used with different meanings in different contexts, they are ambiguous expressions like 'bank' and exemplify a second way in which context helps to fix content. Indexical expressions like 'I' or 'tomorrow' can also be used to say different things in different contexts, but not because they are ambiguous, the mode in which their content depends on context is different. Descriptive expressions such as 'round' can also be used to make different claims in different contexts,⁴ but not because it's not clear what language they belong to, or because they are ambiguous or because they are indexical; the way in which their content depends on context is different from all these. In each kind of context-dependence the mode in which the context operates to fix content is different and so each kind of context-dependence needs a different treatment.⁵

Given the varieties of context-dependence one might expect that contextualists should provide a thorough explanation of which type of context dependence is present in the case of knowledge ascriptions. However, they have either eschewed or underestimated this question. Stewart Cohen, for example, says that "these semantic issues, as near as I can tell, are irrelevant to the epistemological issues. As long as we allow for contextually determined standards, it doesn't matter how formally we construe the context-sensitivity".⁶ I think Cohen is wrong. We will see later that if one construes the "semantic issues" concerning knowledge ascriptions in a certain way then the contextualist theory becomes incapable of handling the "epistemological issue" of solving the sceptical paradox.⁷ The "semantic issues" are not just relevant but crucial for the contextualist's treatment of the sceptical problem as they pose it.

As far as I know the only contextualist that elaborates a bit on the issue of what kind of context-dependence we must think as present in knowledge ascriptions is Keith DeRose. He thinks that what we have in knowledge ascriptions is a case of indexicality.⁸ Stewart Cohen also expresses the belief that the context-dependence in

⁴ Cf. Travis, 1996, pp. 454-456.

⁵ Cf. Perry, 1997, pp. 593-94.

⁶ Cohen, 1999, p. 61.

⁷ Cf. below my discussion of Schiffer's objection (Chapter III.1) and of the "pragmatic objection" (Chapter III.2) to the contextualist solution to scepticism.

⁸ De Rose, 1992, p. 920.

question may be construed as indexicality,⁹ but doesn't say anything about how we should account for it. David Lewis is very cautious in not classifying the context-dependence of knowledge ascriptions as indexicality, he just says that "ascriptions of knowledge are subtly context-dependent".¹⁰ Given the silence of Cohen and Lewis concerning how we should understand the context sensitivity of knowledge ascriptions, I shall take DeRose's effort of providing at least a minimal account of this issue as the official contextualist account.

The distinctive feature of the context sensitivity called 'indexicality' is that the conventional or linguistic meaning of an indexical expression provides a *rule* or *function* that takes us from a token of the expression and a context where the token is used to the designation of the expression:

...[what is common to indexicals] is that the referent is dependent on the context of use and that the meaning of the word provides a rule which determines the referent in terms of certain aspects of the context.¹¹

In this way of understanding indexicality context plays the semantic role of fixing the content of an indexical expression via the meaning of the expression, which is embodied in a rule. For example, the meaning of 'I' is embodied in the rule:

'I' refers to the speaker or writer of 'I'.¹²

For any token of 'I' and any context where it is used, this rule selects an individual as the designation of that token. Given an utterance of 'I' its meaning automatically, so to speak, determines its designation, no supplementary actions or intentions on the part of the speaker are needed.¹³

There is more than one account of the semantics of indexicals,¹⁴ but for the purposes of conceiving 'know' as an indexical expression DeRose prefers David Kaplan's account for reasons I explain now.¹⁵ In Kaplan's account an indexical has

⁹ Cohen, 1987, p. 15; 1988, p. 97; 1999, p. 61.

¹⁰ Lewis, 1996, p. 550.

¹¹ Kaplan, 1989, p. 490. Cf. Heal, 1997, p. 622; Perry, 1997, p. 594 for similar statements of that view.

¹² Kaplan, 1989, p. 505.

¹³ I'm ignoring here the existence of "impure indexicals" or demonstratives.

¹⁴ Cf. Perry, 1997, pp. 590-607 and Kaplan, 1989, pp. 489-507, for two different accounts.

¹⁵ Cf. De Rose, 1992, p. 921.

two semantic levels, one he calls *character*: it is the meaning embodied in the rules governing the use of indexicals, it is the meaning that is “set by linguistic conventions”, “what is known by the competent language user”.¹⁶ The second level he calls *content*: it is the semantic contribution that the indexical makes in a given context to the propositional content of a sentence containing it.¹⁷ The character of indexicals remains constant across contexts, but it determines different contents in different contexts. The sentence ‘I was there yesterday’ has always the same linguistic meaning, but it expresses different propositions in different contexts because the content of ‘I’, ‘there’ and ‘yesterday’ shifts from context to context. The account of indexicals in terms of character and content explains nicely our intuition that in some sense ‘I was there yesterday’ always means the same, but in some other sense it means different things in the mouth of different people. Contextualists want to say something similar about sentences of the form ‘S knows that *p*’, they think that in some sense a sentence of this form means different things when uttered in different contexts, but they want to respect our intuition that in some other sense such a sentence always means the same. ‘Know’ is always used with the same “character”, it always means that the subject of the knowledge ascription is in a strong enough epistemic position to qualify him as a knower of a certain proposition. But ‘know’ changes its “content” across contexts, how strong is strong enough to qualify as a knower is determined by facts of the contexts where the ascription (or denial) of knowledge is made.¹⁸ Kaplan’s account gives contextualists what they need to state this basic theory of knowledge ascriptions, nothing more and nothing less, I think this is why DeRose chooses Kaplan’s view.

It seems that thinking of the context sensitivity of ‘know’ as indexicality suits pretty well the purposes of contextualists, however not everything runs smoothly. We have seen that the way in which context affects the content of indexicals is via a rule that embodies the character of the indexical, this rule “direct us” or “leads us” from an utterance and its context to a designation. If we grasp the character of an indexical *n*, then we are able to tell in any giving context what an utterance of *n* designates, and this is so because in grasping the character of *n* we become aware of the contextual

¹⁶ Kaplan, 1989, p. 505.

¹⁷ Cf. *ibid.*, pp. 500-1.

¹⁸ Cf. DeRose, 1992, p. 922.

facts that are relevant for fixing the content of *n*. For example, if we grasp the character of 'I' and 'yesterday', we see that the relevant contextual facts are, for 'I' who is the speaker and for 'yesterday' which is the day of utterance. Does anything parallel apply to 'know'? Perhaps we can grant that what a competent language user grasps in understanding the meaning of 'know' is that this word is used to describe someone as being in a good enough epistemic position with respect to a certain proposition. But it is not clear that *that* understanding can "direct" or "lead" the speaker to the contextual facts he has to look at as the relevant ones for fixing the content of 'know' (i.e. for fixing how good is good enough) in a particular occasion in which this word is used. Clearly not everything that happens in a context is relevant to the question whether someone knows something in that context, only some facts about that context are relevant for that question. If 'know' were really an indexical we would expect that its conventional meaning would tell us which are the relevant facts, as the meaning of 'yesterday' or 'I' tells us which are the relevant facts for determining which day or person you are talking about when you say *yesterday* or *I*. But the meaning of 'know' doesn't seem to tell us that much and this suggests that 'know' cannot be conceived as an indexical. I will propose a way of overcoming this difficulty after I have revised the kinds of contextual factors to which contextualists think knowledge ascriptions are sensitive. But before revising such factors I'll examine what is supposed to be sensitive to them, *viz.* the content of knowledge ascriptions. Contextualists say that the content of knowledge ascriptions defines how strong S's epistemic position must be in the context of ascription for him to know, but what is meant by 'strength in epistemic position'?

b. How to understand *strength in epistemic position*.

For contextualists what is context sensitive about knowledge ascriptions is how strong someone's epistemic position must be. There are two ways in which they have explained epistemic strength: in terms of the range of alternatives someone's evidence has to eliminate and in terms of the range of counterfactual situations within which someone's belief has to truth-tracking. These are two aspects of knowledge ascriptions that are sensitive to context, but there seem to be others. Here are some examples: the degree of confidence we demand on someone's part before ascribing

knowledge to him varies contextually;¹⁹ how reliable the method through which someone formed a belief in order for it to count as knowledge also varies contextually; whether or not someone must be able to recall a piece of knowledge before we ascribe to him that knowledge is also a context sensitive matter. Contextualists largely ignore the variety of aspects of knowledge ascriptions that are context sensitive, and it is not clear how they should treat them. Since I'm going to present the actual contextualist theory and how it is used against the sceptic, I'll also ignore the variety of aspects of knowledge ascriptions that are context sensitive and focus on the one contextualists do: strength in epistemic position.

i. Strength in terms of elimination of alternatives: Cohen and Lewis

Probably the most discussed version of contextualism is the so called Relevant Alternatives theory.²⁰ The theory says that a subject S knows that p only if his evidence rules out all relevant alternatives to p , i.e. all relevant situations in which something incompatible with p obtains. The contextualism championed by D. Lewis and S. Cohen is of this type. In this type of contextualism the standard for how strong S's epistemic position must be to know is understood in terms of the alternatives that are relevant for S in a given situation. How strong the epistemic position of S must be varies in relation with the range of $\sim p$ possibilities he has to eliminate in order to know. In this way of understanding epistemic strength the standards for how strong S's epistemic position must be are context sensitive because the range of relevant alternatives is context sensitive, it expands or contracts according to changes in the context.

Lewis and Cohen have different views about how the context determines which the relevant alternatives are for a given proposition claimed to be known, but they both share the conception of strength in epistemic position in terms of elimination of relevant alternatives.

¹⁹ Cf. DeRose, 1992, p. 922, fn. 18.

²⁰ DeRose, 1992, pp. 918-921, has emphasised that not all relevant alternative (RA) theories are contextualist. Contextualist versions of RA theories allow that facts about the *context of ascription* of knowledge determine which alternatives are relevant. Non-contextualist versions of RA theories hold that what alternatives are relevant is entirely determined by facts about the putative *knower*. The versions of RA theories that I'll be considering are contextualist.

There are problems with understanding strength in terms of elimination of relevant alternatives, some of them concern the notion of “elimination” or “ruling out”. What does it mean to eliminate or rule out an alternative? Cohen thinks that to eliminate an alternative q is to *know* that $\sim q$.²¹ But this is problematic given that Cohen, as well as other recent contextualists like DeRose and Lewis, want to preserve some form of the Closure Principle.²² Contextualists want to say that in some contexts one can know that p despite being unable to rule a sceptical hypothesis H_T out, so long as the H_T is an irrelevant alternative. Now, if we retain Closure, in those contexts in which one knows that p , by Closure, one also *knows* $\sim H_T$, despite being unable to rule it out. Therefore, if we retain Closure ruling out H_T has to be something *different* from knowing $\sim H_T$.²³

David Lewis gives a different account of “ruling out”. For Lewis the experiences and memories possessed by a subject S rule out or eliminate an alternative q in the sense that *the existence* of those experiences and memories, with the content they have, is incompatible with q being the case:

.....it is the existence of the experience that conflicts with W : W is a possibility in which the subject is not having experience E . Let E have propositional content P . Suppose even.....that E is, in some sense, fully characterized by P . Then I say that E eliminates W iff W is a possibility in which the subject’s experience or memory has content different from P .²⁴

An actual experience of S with content p rules out a certain possibility W just in case it is true that if W were the case then S ’s experience would have a different content than p . It is clear why this is so: if W is a possibility in which S ’s experience would have a different

²¹ Cf. Cohen, 1987, p. 17.

²² Cf. below section 2 of this chapter for how they preserve Closure.

²³ A possible reply to this objection is to hold that one may know the negation of an alternative in two different ways: on the basis of having evidence that rules the alternative out, or non-evidentially on the basis of the “intrinsic rationality” of believing the denial of the alternative (this is Cohen’s suggestion, 1988, pp. 111-113). Then we can consistently say that in contexts in which one knows that p , by Closure one also counts as knowing (*non-evidentially*) that $\sim H_T$, despite our inability to rule out (i.e. know *evidentially*) $\sim H_T$. I’ll not pursue this reply any further since ultimately I’ll prefer the alternative account of strength in epistemic position in terms of truth tracking that doesn’t have the complications of distinguishing evidential from non-evidential knowledge, this is in itself a reason to prefer it.

²⁴ Lewis, 1996, p. 553.

content that p , the fact that S 's experience *actually* has content p entails that W is not the case, W is ruled out by the existence of S 's experience with content p .

By not identifying 'ruling out' with 'knowing $\sim q$ ' Lewis's view avoids the problem of Cohen's, he can consistently say that in the contexts in which we know that p by Closure we *know* that $\sim q$ too even if we still can't rule q out. He can consistently say that because in his account "ruling out" q does *not* amount to knowing $\sim q$.²⁵ However Lewis's interpretation of 'ruling out' has its own problems. How would Lewis explain the possibility of some sort of inductive knowledge?²⁶ Sometimes we claim to *know* that the next a will be G on the basis of memories and observations that all a 's we have encountered have been G s, but it is not true that the existence of such memories and observations rules out, in Lewis's sense, all relevant alternatives to the next a being G . A relevant possibility in which the next a turns out to be not- G is one in which still my memories and perceptual experiences concerning a 's being G are as they actually are. In other words, the existence of my memories and experiences, with the contents they actually have, is *compatible* with relevant possibilities in which the next a isn't G . Lewis's account of 'ruling out' seems to imply that we *can't know* that the next a will be G on the basis of memories and experiences of all encountered a 's being G , simply because the *existence* of those memories and experiences is always compatible with relevant possibilities to what we inductively believe: that the next a will be G . But that the existence of a certain body of memories and observations is compatible with alternative hypotheses to what we believe on the basis of it, seems to be the very nature of inductive *knowledge* (!).

These problems with the account of strength in epistemic position in terms of ruling relevant alternatives out make it difficult to see how it can be developed satisfactorily. I'll present an alternative account of strength in terms of the concept of truth tracking. I'll generally prefer this account when I illustrate contextualist claims later on.

²⁵ A further issue Lewis would have to explain is what's the nature of the knowledge we have of $\sim q$ when we can't rule q out, i.e. when the existence of our evidence is *consistent* with q . But this is a further issue, the point I'm making now is just that Lewis's explanation of 'ruling out' is consistent with retaining Closure, whereas Cohen's is not, unless we are willing to accept the complications described in footnote 23.

²⁶ This problem has been pointed out by Vogel, 1999, p. 168-172.

ii. Strength in terms of truth tracking: DeRose

DeRose has proposed to understand strength in epistemic position in terms of the notion of truth-tracking as defined by Robert Nozick.²⁷ There is a difference between the way Nozick uses the notion of truth tracking and the way DeRose uses it. Nozick builds it into the concept of knowledge itself, that S's belief that *p* tracks the truth is for him a necessary condition of S knowing that *p*. But for DeRose truth tracking is not necessary for knowledge, as we'll see he thinks one can know that sceptical hypotheses are false even though our beliefs that those hypotheses are false do not track truth. Instead of building truth tracking into the concept of knowledge itself DeRose builds it into the concept of *knowledge ascription*. When the knowledge that *p* is ascribed to S in a given context, that S's belief that *p* tracks the truth within a contextually determined domain of possible worlds is a condition on the ascription being true; but so long as the ascription is not made S may well know that *p* without being a condition of his knowledge that his belief that *p* tracks the truth in the domain of possible worlds within which it would have to if the ascription had been made.

In a conversational context in which we ascribe knowledge to people, our ascriptions define how strong people's epistemic position must be for those ascriptions to be true. In DeRose's view what our ascriptions define in defining a required strength is the range of possible worlds within which people's beliefs must track truth:

An important component of being in a strong epistemic position with respect to *p* is to have one's belief as to whether *p* is true match the fact of the matter as to whether *p* is true, not only in the actual world, but also at the worlds sufficiently close to the actual world.....The further away one can get from the actual world, while still having it be the case that one's belief matches the fact at worlds that far away and closer, the stronger a position one is in with respect to *p*.²⁸

²⁷ Cf., Chapter I. 3. b.

²⁸ De Rose, 1995, p. 34.

According to this view, S's strength of epistemic position with respect to p is proportional to how far from actuality S's belief that p tracks the truth of p , the further from actuality truth-tracking reaches, the stronger the epistemic position is. But when we ascribe to S the knowledge that p what determines the domain of worlds within which S's belief that p must track truth for our attribution to be true? DeRose believes that the domain of worlds is determined by facts about the context where the ascription is made:

Context, I've said, determines how strong an epistemic position one must be in to count as knowing. Picture this requirement as a contextually determined sphere of possible worlds, centred on the actual world, within which a subject's belief as to whether p is true must match the fact of the matter in order for the subject to count as knowing....Call this sphere the sphere of epistemically relevant worlds. As the standards for knowledge go up, the sphere of epistemically relevant worlds becomes larger –the truth tracking of one's belief must extend further from actuality for one to count as knowing.....*When it is asserted that S knows (or doesn't know) that p , then....the sphere of epistemically relevant worlds [gets enlarged] so that at least includes the closest worlds in which p is false.*²⁹

One of the facts about the context where the ascription 'S knows that p ' is made is precisely the fact that this ascription was made, according to DeRose *this* contextual fact determines that the sphere of relevant worlds within which S's belief that p must track truth has in its outer boundary, so to speak, the closest worlds in which p is false. In this way the context of the knowledge ascription determines how strong S's epistemic position must be (i.e. how far from actuality S's belief that p must track truth) for the ascription to be true.

I have presented two accounts of the notion of strength in epistemic position, I want to draw attention to the fact that both of them rest on the same idea of relevance of possibilities. The first account defines S's strength in terms of the range of alternatives to p S has to eliminate for him to know that p , the set of such alternatives is called the set of relevant alternatives to p . The second account defines S's strength in terms of the $\sim p$ worlds up to which S's belief that p must track truth for the ascription 'S knows that p ' to be true, the set of possible worlds having in its

²⁹ *Ibid.*, pp. 36-7, emphasis added.

outer boundary those $\sim p$ worlds is called the sphere of epistemically relevant worlds. The idea of relevance is used in different ways, in one case the subject must be able to rule out the possibility that he *is* in any of the $\sim p$ relevant situations, in the other it must be true that if the subject *were* in any relevant $\sim p$ situation, then he wouldn't believe that *p*. However, what defines strength in both cases is a certain relation the subject has to have with the relevant $\sim p$ situations, in both accounts the crucial concept is that of *relevant $\sim p$ situation*.

We saw that the truth-conditional content of a knowledge ascription represents how strong the epistemic position of the subject of the ascription must be, we have seen that strength in epistemic position is defined in terms of a certain relation (evidential or counterfactual) the subject of the ascription has to have with a set of relevant possible situations. What makes the content of knowledge ascriptions shift contextually is that the set of relevant possible situations changes contextually. Next I'll expound the contextualist's account of how exactly a context operates in bringing about those shifts in the domains of epistemically relevant possible worlds.

c. Contextual facts that affect the content of knowledge attributions: subject facts vs attributor facts

Some contextualists underestimate the importance of explaining which are the contextual facts that determine which worlds are to be included in the domain of epistemically relevant worlds. Perhaps they have this attitude because they think their theory doesn't lose its anti-sceptical force even if such contextual factors cannot be fully and precisely specified.³⁰ Other contextualists attempt to identify *kinds* of contextual facts as the determinants of the content of knowledge ascriptions. DeRose, for example, distinguishes the facts about the context of the *ascriber* of knowledge from the facts about the *subject* of the ascription.³¹ He doesn't provide a full list or anything like that of each kind of facts, but thinks that *only* ascriber facts affect the truth-conditional content of knowledge ascriptions. Cohen also makes of the claim

³⁰ An expression of this thought can be found in Cohen, 1988, pp. 94-5, 115.

³¹ De Rose, 1992, pp. 918-23 and 1999, pp. 190-91

that only ascriber facts affect the content of knowledge ascriptions a defining characteristic of contextualism.³²

It is not easy to say in general which contextual facts are on each side of the divide subject facts/ascriber facts. Alvin Goldman was the first contextualist to emphasise this distinction, he thinks of the subject's context as "a complete specification [of the subject's] situation".³³ This specification does not only include what actually exists in the subject's physical and social environment, but also facts about the likelihood in that environment of certain possibilities and about the similarity of some possibilities with the actual environment.³⁴ In relation to the ascriber facts Goldman just says that they include "the speaker's own linguistic and psychological context".³⁵ DeRose doesn't say anything more than Goldman about what lies on each side of the divide.³⁶ There are easy scenarios where the divide is easily appreciated, but there are cases where it tends to be blurred. One type of cases in which that happens are those where the subject and the ascriber are in the same physical environment and linguistic context; other type of cases are those of first-person present tense knowledge claims where the subject of the ascription and the ascriber are the same person. De Rose thinks that even in *these* cases it's still legitimate to speak of the contextual facts that attach to the individual *qua* putative knower and those that attach to him *qua* ascriber of knowledge.³⁷

I want to put aside the problem of how to classify contextual facts as pertaining to the putative knower of a knowledge ascription or to the context of the ascriber of knowledge and emphasise that the contextualist theory used to solve the sceptical paradox is a view according to which the content of knowledge ascriptions, the standards for how strong an epistemic position must be, is determined by facts about *the context of the knowledge ascription*. Consequently this is the contextualist account of what are the contextual factors that determine the content of knowledge ascriptions that I'll be discussing in the rest of this thesis.

However, it is worth noting that for some contextualists how strong an epistemic position must be, *viz.* what's the domain of relevant possible worlds for the

³² Cf. Cohen, 1998, pp. 289, 294.

³³ Goldman, 1976, p. 775.

³⁴ *Ibid.*, p. 776.

³⁵ *Ibidem.*

³⁶ De Rose, 1992, pp. 918-9.

³⁷ *Ibid.*, p. 919.

evaluation of a given knowledge ascription, may be determined both by facts about the context of the knowledge ascription and by facts about the putative knower. David Lewis is an example of this type of contextualist. In Lewis's theory irrelevant $\sim p$ worlds are worlds we may properly ignore and relevant $\sim p$ worlds are worlds we may not properly ignore.³⁸ Lewis offers six rules that define which $\sim p$ worlds may and which may not be properly ignored in a given context. Four of Lewis's rules define which $\sim p$ worlds may *not* be properly ignored, i.e. which are *relevant*, he calls them the Rules of Actuality, Belief, Resemblance and Attention. The two remaining rules define which $\sim p$ worlds *may be* properly ignored, i.e. which are *irrelevant*, he calls them the Rule of Reliability and the Rule of Method.³⁹ Some of Lewis's rules depend on facts about the context of the *ascription* of knowledge, the Rule of Attention is a clear example. This rule says that if we, deciding whether someone knows or not, verbally raise, or in some other way attend to, a certain $\sim p$ possibility then that possibility becomes relevant.⁴⁰ The fact that an ascriber of knowledge considers a $\sim p$ possibility is, *a fortiori*, a fact about the context in which the knowledge ascription occurs, and it is this fact on which the Rule of Attention depends. But other of Lewis's rules depend on facts about the *putative knower*, a clear example is the Rule of Belief, this rule depends on facts concerning what beliefs the putative knower has, it says that if the subject believes, or ought to believe, that a certain $\sim p$ world obtains, then that world is relevant.⁴¹ Despite the diversity of his rules, in handling the sceptical paradox Lewis relies almost exclusively on the Rule of Attention.⁴² This aligns him with Cohen and DeRose, all of them solve the sceptical paradox appealing to a theory according to which the standards for how strong an epistemic position must be are determined by facts about the context of the knowledge ascription. This is why I'm going to take this "context-of-ascription account" of what are the factors that determine the content of knowledge ascriptions as the defining feature of the contextualist solution of the sceptical problem.

³⁸ Lewis, 1996, p. 554.

³⁹ *Ibid.*, pp. 555-60.

⁴⁰ *Ibid.*, p. 559.

⁴¹ *Ibidem.*

⁴² Cf. Lewis, 1996, pp. 560-567

When contextualists talk about “the context of the ascription”, what do they have in mind? It is important to realise that they don’t mean the whole physical and social setting in which the ascription takes place, they mean only the *linguistic or conversational context* within which the ascription occurs. How should we conceive such a conversational context? Contextualists are not very explicit about this, but since they hold that *mentioning* a certain possibility of error or *saying* that someone knows something is enough to change the conversational context, and with it the content of the knowledge ascriptions produced in it, it seems plausible to suggest that they must conceive the conversational context of ascription as a body of information shared, or capable of being shared, by the participants in a conversation. That information is the resulting sum of the bits of information contributed by the speech acts produced by the speakers, and is the background against which subsequent speech acts get interpreted. A conversational context in this sense is simultaneously the object speech acts act upon and the source of information relative to which speech acts are interpreted.⁴³ This notion of conversational context makes good sense of the contextualist’s claim that merely mentioning a certain $\sim p$ possibility or saying that someone knows something changes the context. In effect those acts of mentioning and saying modify the body of information shared, or capable of being shared, by the participants in the conversation in such a way that subsequent ascriptions of knowledge are interpreted relative to this new body of information. The question of whether mentioning possibilities of error and saying that someone knows something produce exactly the contextual shifts that contextualists require (i.e. raising epistemic standards), is a further issue I’ll discuss in the next chapter. Here I’m just pointing out how contextualists should conceive a conversational context to *make sense* of the idea that mentioning or saying something can change the context.

I want to finish this section going back to the question concerning the plausibility of conceiving the context sensitivity of knowledge ascriptions as indexicality. I pointed out that one problem with this view is that ‘know’ does not seem to fit quite well the standard model of other indexicals because its linguistic meaning does not seem to provide us with a rule that “directs us” to the contextual

⁴³ For a detailed development of this notion of conversational context see Stalnaker, 1970 and 1998.

facts that are pertinent for fixing its content. Given our preceding discussions I think that the fairest thing to say, if we want to stick to the idea that the sensitivity of 'know' is indexicality, is that this word is a special kind of indexical that differs from other indexical expressions in that which contextual facts it is sensitive to is itself a context sensitive matter.⁴⁴ 'Know' is not like 'I' or 'yesterday' which are always sensitive to the same contextual facts, 'I' to the person who utters it and 'yesterday' to the day of utterance.⁴⁵ Sometimes the facts that fix the content of a knowledge ascription may concern the conversational context in which the ascription is made, sometimes they may concern facts about the subject of the ascription, but not always the same type of facts about the subject may be operative in determining the content of the ascription, in different contexts different facts may be operative.

In the remainder of the present chapter I shall explain how contextualist use the context sensitivity of knowledge ascriptions expounded so far to provide a solution to the sceptical paradox we identified in Chapter I.

2. Solving the sceptical paradox

As I explained at the end of chapter I for contextualists solving the sceptical problem consists in solving the paradox of intuition generated by the argument CP we identified in chapter I:

- CP1. If Kp and $K[p$ entails $\sim HT]$, then $K\sim HT$.
- CP2. $K[p$ entails $\sim HT]$
- CP3. $\sim K\sim HT$
- CP4. $\sim Kp$ (From CP1, CP3 by Modus Tollens)

This valid argument generates a paradox because we have intuitions for CP1-CP3 and against CP4. The resolution of this paradox consists in explaining how

⁴⁴ I think this is the view of P. Yourgrau when he points out that "different contexts might determine different criteria of relevance" (1983, p. 184).

⁴⁵ We might say that 'know' differs from other indexicals in that its context dependence is "free" from rather than "controlled" by its linguistic meaning. I borrow this terminology from Recanati, 1989. He argues that the context dependence of many expressions is "free" and not "controlled" by their linguistic meaning in the sense that their meaning "constrains its possible semantic values but does not consist in a 'rule' or 'function' taking us from context to semantic value" (Recanati, 1989, p. 298). Recanati doesn't suggest that some *indexicals* may be conceived in this way, I'm just speculating about the possibility of treating some expressions as indexicals whose context sensitivity is "free" from its linguistic meaning. If this is a possibility then 'know' could be counted among them.

acceptance of the premises of the argument and rejection of the conclusion are not really inconsistent attitudes, and only appear to be so because the sceptic manipulates our intuitions for the premises and against the conclusion exploiting the context sensitive features of knowledge ascriptions that we examined in the first half of this chapter. Let's see the details of this resolution of the paradox.

Contextualists think that (CP1) is true so long as the conversational context⁴⁶ relative to which the embedded knowledge sentences in (CP1) get their truth conditions is kept fixed.⁴⁷ Then they explain our conflicting intuitions regarding (CP3) and (CP4) as arising from the fact that we unwittingly evaluate these premises in different conversational contexts where different epistemic standards are operative and where those premises get different truth values. Contextualists say that if those premises were evaluated in the same context they would get the same truth value, no conflicting intuitions would arise and there would be no paradox.

DeRose provides the most detailed resolution of the paradox along these lines. He proposes a rule that controls which are the standards for knowledge operative in a context, he calls it the "Rule of Sensitivity", it says that

When it is asserted that some subject S knows (or does not know) some proposition P, the standards for knowledge (the standards for how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require S's belief in that particular P to be sensitive for it to count as knowledge.⁴⁸

[A belief's being sensitive means that it tracks truth up to the closest $\sim p$ worlds]

According to early contextualists, such as Dretske,⁴⁹ the Closure Principle, and consequently premise (CP1), fails; rejecting Closure was his way of rebutting scepticism. More recent contextualists, such as DeRose, Lewis and Cohen want to retain Closure and explain the apparent failures of Closure as mere illusions due to

⁴⁶ It is crucial to notice that the notion of context in play throughout the resolution of the paradox is that of *the conversational context of the ascriber of knowledge*, as I explained it earlier in this chapter.

⁴⁷ Cf. Cohen, 1988, p. 111; De Rose, 1995, pp. 29-33; Lewis, 1996, p. 564.

⁴⁸ De Rose, 1995, p. 36.

⁴⁹ Cf. Dretske, 1970.

shifts in contexts.⁵⁰ DeRose's Rule of Sensitivity gives us that explanation.⁵¹ When we assert or just consider the antecedent of (CP1) the Rule of Sensitivity dictates that our belief that p must be sensitive for it to be knowledge, since normally our beliefs about worldly facts are sensitive, we feel that the antecedent of (CP1) is true. But as we go on to assert or just consider the consequent of (CP1) the same rule dictates that our belief that $\sim H_T$ must be sensitive for it to be knowledge, since this belief is *never* sensitive⁵² we feel that the consequent of (CP1) is false. It seems to us that we can know that p without knowing that $\sim H_T$, so it seems to us that (CP1) fails. But in reaching this verdict we don't notice, according to contextualists, that the conversational context, and with it the standards for knowledge, have changed midway through our assertion of or reflection upon (CP1). We start with an everyday conversational context which defines an epistemic standard consisting of a sensitivity condition that our beliefs do meet, but our mentioning or considering the possibility $\sim H_T$ changes the conversational context in such a way that it now defines a tougher standard consisting of a sensitivity condition that our belief that $\sim H_T$ do not meet. However, if we held fixed the conversational context, and with it the operative epistemic standard, (CP1) will always come out true:

In everyday contexts. On the one hand, contextualists say that relative to the standards operative in everyday conversational contexts we know both that p and that $\sim H_T$. Notice that for us to know that $\sim H_T$ relative to some standard this standard should not mandate that our belief that $\sim H_T$ be sensitive (because it is never sensitive). But the Rule of Sensitivity installs a standard that mandates precisely that sensitivity requirement whenever $\sim H_T$ is asserted or considered. Therefore, if there is such a thing as everyday standards relative to which we know $\sim H_T$, those standards have to be in place *before* the Rule of Sensitivity operates, that is, before the possibility H_T is mentioned or considered. This means that everyday conversational contexts relative to which we know both p and $\sim H_T$ are contexts where sceptical hypothesis are not mentioned.⁵³

⁵⁰ Cf. the references of footnote 47.

⁵¹ Lewis's and Cohen's explanation of why (CP1) sometimes seems to be false, and in general their dissolution of the paradox, is not substantially different from DeRose's. The difference between their explanations is that Lewis's appeals to his Rule of Attention (1996, p. 564) and Cohen to a Rule of Saliency (1988, pp. 106-110). But the operation of these rules brings about the same results that the operation of DeRose's Rule of Sensitivity; so, for the sake of brevity, I take DeRose's as the representative example.

⁵² Cf. Chapter I. 3. b.

⁵³ In chapter I. 3 I expounded two arguments that purport to show that *nobody knows $\sim H_T$* , one due to Stroud and the other to Nozick. Given that contextualists hold that *we know $\sim H_T$* relative to everyday contexts, one would expect that they can provide an answer to those arguments, can they? In everyday contexts the Rule of Sensitivity is, so to speak, suspended with respect to $\sim H_T$, and this explains why

In sceptical contexts. A sceptical context is one in which sceptical hypotheses are mentioned. Once the hypothesis H_T is mentioned the Rule of Sensitivity drives the standards for knowledge up, and relative to these new tougher standards we fail to know both $\sim H_T$ and p because those standards now require our beliefs that p and that $\sim H_T$ to be sensitive in sceptical possible worlds, i.e. worlds in which H_T obtains, and neither our belief that p or our belief that $\sim H_T$ are sensitive in those worlds.

In both an everyday context and a sceptical context it is true that if we know p then we know $\sim H_T$; in the former context we know p and also $\sim H_T$, and in the latter we don't know $\sim H_T$ nor p . (CP1) is "true regardless of what epistemic standards it's evaluated at",⁵⁴ so long as the epistemic standards are held constant all the way through.⁵⁵

Our conflicting intuitions about endorsing (CP3) and rejecting (CP4), as the apparent failures of (CP1), are explained too as resulting from overlooking contextual shifts brought about by the Rule of Sensitivity. DeRose thinks that the truth conditions and the truth value of (CP3) and \sim (CP4) vary with context, in

in these contexts we can know $\sim H_T$ despite that our belief that $\sim H_T$ is not sensitive. This means that the suspension of the Rule of Sensitivity helps contextualists to explain how we can know $\sim H_T$ in everyday contexts *despite what Nozick's argument shows*, viz. that our belief that $\sim H_T$ is not sensitive (cf. Chapter I. 3. b). This is the contextualist answer to Nozick's argument, but how would they respond to Stroud's argument? What Stroud's argument purports to show is *not* that our belief that $\sim H_T$ is not sensitive, but rather that the use of any procedure to establish $\sim H_T$ leads to an infinite regress of procedures (cf. Chapter I. 3. a). Consequently, the suspension of the Rule of Sensitivity with respect to $\sim H_T$ in everyday contexts can't help contextualist to explain how we can know $\sim H_T$ in everyday contexts *despite what Stroud's argument shows*. Even when the Rule of Sensitivity is suspended with respect to $\sim H_T$ in everyday contexts it may still be true that the use of any procedure to establish $\sim H_T$ leads to an infinite regress. Stroud's argument is not answered by the contextualist's response to Nozick's argument, contextualists need to provide a different treatment for Stroud's argument but I haven't found it in the contextualist literature I've studied. This difficulty stresses a point I made in Chapter I: there's not one single line of argument in route to Cartesian scepticism, there are several.

⁵⁴ DeRose, 1995, p. 39

⁵⁵ The fact that contextualists preserve Closure in the way explained above (i.e. within fixed contexts) raises the worry of how would they accommodate Closure with the failure of transmission of warrant from p to $\sim H_T$ (for a discussion of failure of transmission of warrant cf. Wright, 1985 and Davis, 2000). Contextualists accept that in everyday contexts warrant, justification or evidence fails to be transmitted from p to $\sim H_T$. Lewis and Cohen think that in everyday contexts we possess evidence for p but no evidence that rules H_T out, nevertheless in those contexts, according to them, we know both p and $\sim H_T$. In everyday contexts we have closure of knowledge without transmission of warrant (cf. Davis, 2000, p. 401). How do things stand in sceptical contexts? Here is a suggestion consistent with contextualism. For the reasons exposed in the text, in sceptical contexts we don't know p nor $\sim H_T$. As regards evidential warrant matters are as follows. For any of our evidence for p to count as such we must presuppose that $\sim H_T$, since for anyone who doubted that there's no Evil Demon *always deceiving us* the usual evidence that establishes p would be no evidence for p at all, it would lose its pro-vative force. But in sceptical contexts where the doubt whether $\sim H_T$ is raised we are not entitled to presuppose that $\sim H_T$ and then nothing of what we customarily take as evidence for p counts as such. Hence, in sceptical contexts we have closure of knowledge without transmission of warrant, not because there were evidence for p and no evidence for $\sim H_T$, but because in such contexts there's no evidence for *any* of them.



particular (CP3) is true and \sim (CP4) false in the sceptical context in which (CP3) is asserted, but (CP3) is false and \sim (CP4) true in everyday conversational contexts:

[We endorse (CP3)] as true, not at all standards, but only at the unusually inflated standards conducive to scepticism. Thus, on our solution, we do know, for instance, that we're not BIVs, according to ordinary low standards for knowledge. But, though [(CP3)] is false when evaluated according to those ordinary low standards, we're able to explain its plausibility, as we have seen, by means of the fact that the high standards at which [CP3] is true are precisely the standards that an assertion or denial of it put into play. Since attempts to assert [CP3] are bound to result in truth, and attempts to deny it are destined to produce falsehood, it's no surprise that we find it so plausible.⁵⁶

On the present solution, claims to know ordinary propositions are true according to ordinary low standards but false according to the highly inflated standards that, by the Rule of Sensitivity, are put in place by the assertion of [CP3].⁵⁷

The paradox arises because (CP3) and (CP4) are not evaluated in the same conversational context in the course of considering the sceptical argument. When the sceptic *asserts* (CP3), he creates a conversational context that, by the Rule of Sensitivity, determines that for S to know $\sim H_T$ his belief that $\sim H_T$ must be sensitive. Since the belief that $\sim H_T$ is never sensitive (CP3) seems to us and actually is true in the sceptical context. This explains the origin of our intuitions for (CP3). Once the sceptic has asserted (CP3) and thereby secured its truth, the sceptic points out that by logic alone (CP4) follows from (CP1) and (CP3). We feel this consequence as intolerable because when *we* think of (CP4) we evaluate it against the background of the everyday conversational contexts in which *we* ascribe knowledge to people. In these contexts, our beliefs about worldly facts meet the requirements of how strong one's epistemic position must be to count as knowing. This is so because, as we have seen above, in those ordinary conversational contexts the question whether $\sim H_T$ is not considered and so the Rule of Sensitivity does not raise the truth tracking or sensitivity requirements for knowledge up to the point where they cannot be met. That's why we know all sort of things relative to the standards of everyday conversational contexts, and that's why (CP4) seems to us and actually is false in

⁵⁶ DeRose, 1995, pp. 39-40.

⁵⁷ *Ibid.*, p. 40.

those contexts. This explains the origin of our intuitions against (CP4). But if we held constant the conversational context the sceptic uses to evaluate (CP3) in our evaluation of (CP4), we wouldn't be shocked by the consequence that in *that context* (CP4) is true. The standard operative in that context determines for 'S knows that *p*' truth conditions that can't be satisfied. Keeping the sceptical context fixed we should recognise peacefully that, given the truth conditions for knowledge ascriptions determined by that context, we indeed don't know anything about the world. The paradoxical conflict of intuitions arises only because the conversational context shifts in going from (CP3) to (CP4), i.e. when (CP3) is evaluated against the context installed by the sceptic and (CP4) against our everyday conversational contexts.

It is crucial for the argument that I'll develop in chapter III.3 to notice that in solving the sceptical paradox contextualists not only claim that there's no paradox, that the concept of knowledge is coherent; they also think that the concept has true applications relative to the epistemic standards operative in everyday contexts, that is to say, that people know things about the world in everyday contexts (see for example the last quotation from DeRose above). We will see how problematic this claim becomes within the contextualist theory itself.

There are a couple of issues concerning the resolution of the sceptical paradox that seem crucial for its success. First, we have the issue of the alleged context sensitivity of knowledge ascriptions as something that affects the semantic content or truth conditions of such ascriptions. One may wonder what's the evidence for that context sensitivity, and especially for its alleged semantic character. The question becomes pressing when we notice that the solution of the paradox seems to require that speakers don't notice which contextual shifts have an impact on the content of knowledge ascriptions.⁵⁸ Normally, with respect to unquestionable context sensitive expressions, such as 'I' or 'tomorrow', speakers know which contextual changes affect their semantic value. What can be the evidence for the semantic context sensitivity of knowledge ascriptions given that the solution of the paradox requires such sensitivity to be, so to speak, imperceptible to speakers? In different ways the two first objections to contextualism I'll examine in the next chapter press

⁵⁸ A full elaboration of this issue is next in Chapter III.1

on this question. I'll argue that both objections can be met and that we can grant to contextualists the semantic context sensitivity of knowledge ascriptions.

The second issue that seems crucial to contextualism is this: assuming that the context sensitivity of knowledge ascriptions can be established, there would still be a question about whether the consequences of knowledge ascriptions being sensitive in that way are precisely the consequences that contextualists want. The third and final objection to contextualism I'm going to develop in the next chapter presses on this question. I'll argue first that it is doubtful that the context sensitivity of knowledge ascriptions, as contextualists construe it, enables them to solve the sceptical paradox, i.e. to show that the concept of knowledge is coherent or paradox-free; secondly, I'll argue that their construal of that sensitivity has catastrophic consequences for the prospects of justifying their claims about the possibility of knowledge relative to the epistemic standards operative in everyday contexts.

Chapter III

Problems for the Contextualist Solution to Scepticism

1. Is the contextualist resolution of the paradox internally coherent?

Recall the argument that generates the sceptical paradox:

- CP1. If K_p and $K[p \text{ entails } \sim HT]$, then $K\sim HT$
- CP2. $K[p \text{ entails } \sim HT]$
- CP3. $\sim K\sim HT$
- CP4. $\sim K_p$ (From (CP1), (CP3) by Modus Tollens)

The paradox is that we have intuitions for (CP1)-(CP3) but against (CP4), or what is the same, for \sim (CP4). That is to say, we have intuitions for mutually inconsistent propositions: (CP1)-(CP3) and \sim (CP4). Contextualists dissolve the paradox by explaining how our conflicting intuitions arise thanks to the indexicality of knowledge ascriptions. In going from (CP1)-(CP3) to (CP4) we shift conversational context and with it we *unnoticeably* shift the epistemic standards that determine the content or truth conditions of knowledge ascriptions. We end up having conflicting intuitions about (CP1)-(CP3) and (CP4) because we *don't notice* the contextual shift that occurred, had it not occurred and had we perceived that it didn't, we wouldn't have got the conflicting intuitions that we got and so no paradox would have arisen in the first place. The first objection to contextualism I want to examine concerns the coherence of dissolving the paradox in *this way*, i.e. by invoking *unnoticed* changes in the content of knowledge ascriptions whose possibility is in turn explained by reference to certain *indexical* aspects of knowledge ascriptions.

Stephen Schiffer has argued that the contextualist's resolution of the paradox is not coherent. According to him the resolution of the paradox requires an "error theory" that depicts putative knowers as misunderstanding what they say or think. It represents competent speakers of English as confounding the proposition expressed by a given knowledge sentence in one conversational context (e.g. the proposition expressed by (CP4) in the sceptical context) with other proposition that would have been expressed by uttering the same sentence in a different conversational context (e.g. the proposition expressed by (CP4) in an everyday context). Speakers are

mistaken about what they are really saying (or thinking) when they say (or think) that someone knows or doesn't know something, and this is because they don't appreciate when a relevant contextual shift, affecting the content of what they say, has occurred. Schiffer thinks that this error theory "refutes"¹ the initial contextualist claim about the indexicality of knowledge sentences, that is, the claim that the same knowledge sentence can express different propositions in different contexts. If this is so the contextualists are in real trouble. The contextualist's resolution cannot succeed without either the error theory or the indexical semantics, he needs them both. The error theory is needed to explain why we get the conflicting intuitions that (CP1)-(CP3) are true and (CP4) false, and the semantics is needed to allow for the fact that (CP4) can express different propositions in different contexts, something that is in turn presupposed by the error theory. To apply the error theory that says that we confuse the proposition expressed by (CP4) in one context (i.e. the sceptical one) with a different proposition it would express in other context (i.e. an everyday context), first we have to presuppose that (CP4) *can* express different propositions in different contexts, and this means to presuppose the indexical semantics for knowledge sentences. If the error theory "refutes" the semantics then it would destroy a condition of its own application.

Schiffer briefly examines the prospects of combining the error theory with two semantic theories that would account for the context sensitivity of knowledge ascriptions, one sees it as indexicality and the other as vagueness.² I'll leave aside vagueness and discuss only his criticisms of the idea of combining the error theory with an account of knowledge ascriptions in terms of indexicality. I'll proceed in this way because the contextualists I'm discussing in this thesis use the model of indexicality, not vagueness, to account for the contextual sensitivity of knowledge ascriptions. If I'm correct that Schiffer's criticisms to the combination of the error theory with indexicality can be answered, then he fails to show that there's an incoherence in the way *the contextualists I discussed in chapter II* resolve the sceptical paradox, even if it turns out that his criticisms to the combination of the error theory and vagueness are successful.

¹ Schiffer, 1996, p. 325.

² *Ibid.*, pp. 326-328.

How is the error theory supposed to refute the semantics? Schiffer thinks that if speakers were confused about which proposition a knowledge sentence expresses in a given conversational context, then the indexical account of such sentences couldn't be correct, which is to say that if the indexical account were correct then it could not be the case that speakers are confounded in the way the error theory says they are. The reason according to Schiffer is that "speakers would know what they were saying if knowledge sentences were indexical in the way the Contextualist requires".³ I think this expresses Schiffer's assumption that the content of indexical sentences is transparent to speakers that utter them. It is because he makes this assumption that he thinks that if knowledge sentences were indexical then speakers could not be confounded about what they say in uttering those sentences. How can that assumption about transparency be defended or, in any case, motivated?

Schiffer starts off by trying to show how implausible it would be to suppose that the transparency assumption is false. He says that to suppose that a speaker is systematically confounded about the content of certain knowledge ascriptions would be as bad as supposing that in uttering the sentence 'It's raining' in London a speaker, "who knows where she is",⁴ mistakenly thought that she was asserting the proposition that it's raining in Oxford. This remark is strange because, of course, *if the speaker knows where she is* it would be totally implausible to suppose that she might commit the mistake of thinking that when she asserts 'It's raining' she is asserting the proposition that it is raining in a place different from the place where she knows she is. In the same way, it would be implausible to suppose that a speaker might commit the mistake of thinking that when he asserts 'S knows that p' he is asserting the proposition that S meets standard N of strength in epistemic position, *if he knows that he is asserting that in a context that determines M as the standard to be met.*

If we describe situations in which the speaker *knows* which context he is in, or knows the contextual parameters that define what he is saying in uttering a sentence with indexical components, then obviously the content of what he says is going to be transparent to him. But our ability to describe situations in which the transparency assumption holds doesn't show that the assumption is true *in general.*

³ *Ibid.*, p. 328.

⁴ Schiffer, 1996, p. 326.

To establish this Schiffer needs to show that the capacity to use indexical sentences *guarantees* knowledge of the contextual parameters that define the content of those sentences, only if this is shown it will be true that “speakers would know what they were saying if knowledge sentences were indexical”.

I don't see how that strong connection between the use of indexical sentences and knowledge of their content can be established, there seems to be plenty of evidence that such a connection does not exist. One way in which the connection is severed occurs when speakers have *false beliefs* about the contextual parameters that determine what they are saying in uttering a sentence with, either explicit or hidden, indexical components. In such cases the speakers might take themselves to be saying something they are not really saying. Here is an example: John is travelling from Paris to Oxford by bus stopping at London, he sleeps from Paris to London and when he wakes up in London, thinking that the bus has arrived at Oxford, he looks through the window and says: “It's raining again!”. John is confused about what he is saying, he thinks he is saying that its raining in Oxford but what he is actually saying is that it's raining in London. In appealing to the alleged indexicality of knowledge ascriptions to dissolve the sceptical paradox the contextualist has in mind a phenomenon like that in which John's confusion arises. We may have false beliefs about what standard of epistemic strength a subject has to meet to know in a given conversational context, and this may lead us into thinking that what we say in saying in that context that he knows is something we are not really saying. In such a case we would be confused about what we are saying when we say that S knows, as John is confused about what he says in saying that it's raining.

It might be objected that the contextualist cannot really model his case on something like John's case because John's case is one where the confusion of contexts and propositions expressed by his utterances is not *systematic* whereas the contextualist case is meant to exemplify a case of systematic confusion or misunderstanding. It's true that in the example of John, as I described it, he is not victim of systematic confusion. Other people taking the same trip are not confused about what they say in saying 'It's raining' in London, or perhaps some people in the bus are also confused but not in the same way that John: some might have different false beliefs from John's about where the bus has stopped. However, we can embellish the description of John's case in such a way that *everybody* taking that bus

will end up being confused *in the same* way about what they say in saying 'It's raining' in London. The point is that systematic confusion about what we say in uttering a sentence with indexical components is a perfectly possible phenomenon that (contrary to what Schiffer suggests) doesn't seem to "refute" a semantic theory about those sentences that treats them precisely as what they are: as indexical.

One way of motivating the transparency assumption against the possibility of *systematic* confusion of content of indexical sentences may appeal to some considerations concerning the conditions of successful communication involving that kind of sentences. Communication by means of indexical sentences is possible because speakers and hearers are typically able to keep track of the contextual parameters the content of the indexical sentences is sensitive to. If they were systematically unable to keep track of those parameters transmission of information would fail. But communication about knowledge or involving knowledge ascriptions seems to be typically successful, if knowledge ascriptions are indexical the straightforward explanation of that success would be that transparency holds: speakers and hearers typically are able to keep track of the contextual standards that define the content of knowledge ascriptions. But then the idea of systematic confusion evaporates. If the contextualist semantics were correct the error theory would be wrong.

I think this way of motivating the transparency assumption in order to show that there's incoherence in the contextualist resolution of the paradox overstates the systematicity of the confusion the contextualist needs. The contextualist doesn't say that most of the time in *most conversational contexts* we are confused about what we say in ascribing knowledge. What he says is that systematic confusion occurs in *some* unusual contexts, rarefied by discussions of scepticism. In these contexts communication about knowledge effectively is likely to fail because speakers and hearers are likely not to keep track of the operative epistemic standards.

But transparency of sentences we regard as indexical not only fails in some unusually strange conversational contexts, I've pointed out that it may fail in quite ordinary contexts in which speakers have false beliefs about the relevant contextual parameters. Other cases in which the transparency assumption is falsified occur when the speaker has *no beliefs at all* about those parameters. Suppose that instead of taking a trip John is kidnapped, when he is got into London he has no idea where he

is but he is aware that wherever he is it is raining there, so he whispers: 'It's raining'. If he was asked by one of the kidnappers 'What do you mean?', John would just repeat what he said because he is not in a position to grasp and articulate the content of what he said: that it is raining in London. Schiffer seems to ignore the possibility that the content of indexical sentences may not be transparent to speakers because they have no beliefs about the contextual parameters that fix the content of their utterances. That he ignores this possibility is shown by the fact that he contrasts 'It's raining' with 'S knows that p' on the score that speakers know "full well" what they say in uttering the former but are "totally ignorant of the sort of thing [they are] saying"⁵ in uttering the latter. He clearly thinks that this alleged contrast undermines the idea that knowledge sentences are indexical because if they really were the contrast in question shouldn't exist, the content of 'S knows that p' should be transparent to speakers, as the content of 'It's raining' allegedly is:

One who implicitly says that it is raining in London in uttering 'It's raining' knows full well what proposition she's asserting; if articulate, she can tell you that what she meant and was implicitly stating was that it was raining in London. But no ordinary person who utters 'I know that p', however articulate, would dream of telling you that what he meant and was implicitly stating was that he knew that p relative to such-and-such standard. If, for example, this ordinary guy says 'I know that Placido Domingo is scheduled to sing at the Met this season' and you ask him what exactly he said, he'll tell you that what exactly he said, and meant, was that he knew that Placido Domingo was scheduled to sing at the Met this season.⁶

But the alleged contrast between 'It's raining' and 'S knows that p' doesn't hold in general. Kidnapped John cannot articulate the content of his utterance 'It is raining' for a similar reason that the ordinary guy cannot articulate the content of his utterance 'I know Placido Domingo is scheduled to sing at the Met this season', they both are ignorant or unaware of what are the contextual parameters that define the content of their utterances. Notice, however, that both, the ordinary guy and John, might rectify their ignorance of what they are saying. If the ordinary guy carried out a more or less sophisticated reflection, he could come to appreciate that for him to know about Placido Domingo's concert he would have to rule out some alternatives to what he

⁵ Schiffer, 1996, p. 326.

⁶ *Ibid.*, pp. 326-327.

believes which are salient or relevant in that context. This means that he might get himself into a position to see how strong his epistemic position must be in that context in order for him to know, and in this way he could come to recognise what the content of his self-ascription of knowledge is: that he must eliminate such and such contextually relevant alternatives to what he believes. Similarly, by getting the information about where he is kidnapped John would get himself into a position to grasp the content of what he said in saying 'It's raining'.

There's no general contrast between the content of indexical sentences being transparent and the content of knowledge ascriptions not being transparent, both can be transparent or opaque. Therefore, the fact that the content of knowledge ascriptions isn't typically transparent to speakers doesn't entail that such ascriptions are not indexical.

The failure of transparency with respect to indexical sentences is the key to see why Schiffer fails to show that the contextualist's error theory "refutes" the contextualist's semantics for knowledge ascriptions. The error theory requires that speakers do not fully grasp the content of the knowledge sentences they utter, that's why they get confused in the sceptical context about what they are really saying. Schiffer points out that typically the portion of the content of their utterances of knowledge sentences that speakers grasp is actually so minimal that they are "totally ignorant of the sort of thing they are saying". Even if this is so it cannot "refute" the idea that knowledge sentences are indexical, simply because transparency of content, as I've argued, is not a mark of indexical sentences.⁷ Schiffer's charge that an error theory is inconsistent with an indexical semantics for knowledge ascriptions can be answered.

In the foregoing discussion of Schiffer's objection I've been assuming that contextualists are right in conceiving the context sensitivity of knowledge ascriptions as a *semantic* matter that affects the truth conditions of those ascriptions. But is this true? I want to consider an objection that questions, not the coherence of combining the contextualist semantics and the error theory but the very *semantic* character of the

⁷ I'm *not* arguing that the context dependence of knowledge ascriptions is indexicality, what I've argued is that it cannot be shown that they are not indexical simply for not conforming with Schiffer's transparency assumption. If knowledge ascriptions should not be conceived as indexical it must be for some other reason.

context sensitivity of knowledge ascriptions assumed by contextualists in their theory of knowledge ascriptions.

2. Semantic or pragmatic phenomenon?

The contextualist solution to scepticism is built upon what is meant to be a *semantic* theory of knowledge ascriptions, i.e. an account of what determines the content or truth conditions of those sentences. The core idea of the account is that the content of those ascriptions is sensitive to facts of the conversational context in which they are made; this context sensitivity is supposed to be a semantic phenomenon that affects truth conditions. There's an objection to contextualism which says that the context sensitivity of knowledge ascriptions does not really affect their *truth conditions* but only their conditions of *warranted assertability*. The problem for contextualists posed by this objection is what evidence they can find to show that the context sensitivity in question *is* a semantic matter that affects truth conditions.

We saw above that one aspect of Schiffer's complaints about the contextualist semantics for knowledge ascriptions is that speakers are unable to articulate the content or truth conditions that according to that semantics those ascriptions have. This exacerbates the problem of finding evidence for the contextualist semantics because, if Schiffer is right, it would seem that we couldn't appeal to evidence about the understanding speakers have of knowledge ascriptions, and about their capacity to articulate the content of such ascriptions, in order to support the contextualist semantics. If Schiffer is right *that evidence* wouldn't support this semantics. However, I responded to Schiffer's complaint about the incapacity of speakers to articulate the content that contextualists attribute to knowledge ascriptions by saying that such incapacity can be corrected and on reflection speakers can come to appreciate that the content of their knowledge ascriptions is that which the contextualist semantics says it is. A further development of this idea is going to be the key of my answer to the objection to the contextualist semantics I shall consider now.

a. The pragmatic objection

What is exactly the objection to the *semantic status* of the context sensitivity of knowledge ascriptions that contextualists advocate? In Chapter I.2 we saw the seeds of this objection in the sceptic's conception of the relation between his detached assessment of human knowledge and our everyday epistemic practices. The sceptic's conception of that relation parallels the attitude we have towards the plane spotters in the example we describe in Chapter I. As we saw the example of the spotters is meant to illustrate (a) that the conditions for warranted assertability and the truth conditions of knowledge ascriptions are two different things: a knowledge ascription may be warranted and false or unwarranted and true; and (b) that we can confuse one kind of conditions for the other.

The objection to contextualists I'll examine utilises the distinction between conditions for the warranted assertability and conditions for the truth of knowledge ascriptions. The objection says that contextualists correctly identify a certain context-sensitivity of knowledge ascriptions, but they misclassify what they see. They think that context-sensitivity bears upon the truth conditions of knowledge sentences, i.e. upon what are the real conditions for knowing something, but it really bears only upon the conditions of warranted assertability of knowledge ascriptions, i.e. upon the conditions under which it is appropriate, useful, relevant, helpful or otherwise conversationally correct to ascribe knowledge in different contexts: "[T]he contextualist", the objector says, "has confused a variance in the warranted assertability conditions of knowledge for a variance in their truth conditions".⁸

In what follows I'll say that this objection to contextualists amounts to the charge that the context sensitivity of knowledge ascriptions is not a *semantic* but a *pragmatic* phenomenon. Given the diversity of senses attached to the expressions 'semantic' and 'pragmatic' I'm going to make the following stipulation:⁹ 'Semantic phenomenon' is to mean linguistic phenomena that bears upon the truth conditions of a sentence, and 'pragmatic phenomenon' is to mean linguistic phenomena that doesn't bear upon truth conditions but upon the conditions that must be fulfilled to

⁸ DeRose, 1999, p. 196.

⁹ Cf. Bach, 1999, Appendix A pp. 81-82 for a list of the different ways in which 'pragmatics' is understood by philosophers and linguists.

achieve the purposes sentences are used for. Here is one example of the distinction semantic/pragmatic drawn in these terms.

...truth must be the central notion of *semantics*, and the boundary between what does and what does not bear logically on the truth of what is strictly said must be the boundary between the science of *semantics* and the science of the further effects obtaining in speech-exchange. There is a precedent for calling the latter *pragmatics*, but obviously this name is still little more than a catch-all for what does not bear on strict meaning.¹⁰

Before going on we must be very clear about what the pragmatic objection would accomplish if successful. It *wouldn't* establish that the truth conditions of knowledge ascriptions are the truth conditions the sceptic says, the objection would only establish that the truth conditions of knowledge ascriptions are not sensitive to context. This means that the objection would eliminate a "competitor" to the sceptic's view about what the truth conditions of knowledge sentences are, but to accept the sceptic's view we would still need positive reasons, different from the objection itself. However, what the objection would accomplish if correct would cause enough damage to the contextualist solution to scepticism. Contextualists say that Cartesian scepticism is true *only* relative to a conversational context in which sceptical hypotheses are considered.¹¹ They think this is so because the truth conditions of the sceptical conclusion are *tied to a sceptical context* and the *semantics* of knowledge ascriptions precludes the generalisation of such conclusion to other conversational contexts. If the pragmatic objection is correct, there wouldn't be any such semantics of knowledge ascriptions for the contextualists to appeal to in order to restrict the scope of the sceptical conclusion to the context in which it was reached. The semantic obstacle contextualists put to the unrestricted scope of the sceptical conclusion would have been removed.

The crucial question raised by the pragmatic objection is: how are we going to determine whether the context sensitivity of knowledge sentences is a semantic or a pragmatic phenomenon? Some contextualists introduce their theory that the context sensitivity in question is semantic as a hypothesis to explain the fact that a knowledge sentence like:

¹⁰ Wiggins, 1971, p. 21.

¹¹ Cf Chapter II.2 above and Chapter III.3.b below.

(*) S knows that p

uttered in one context seems true whereas the same sentence in other context seems false.¹² This is the context sensitivity that constitutes the data that the semantic theory of contextualists is an explanation of. The theory says that there is a context sensitive component built into the meaning of (*), that component is the domain of relevant possible worlds within which S's belief has to track truth.¹³ Different contexts fix different domains of relevant worlds and so fix different propositions to be expressed by (*). In an everyday context (*) *seems true* because the proposition it expresses there *is true*, and in a sceptical context (*) *seems false* because the different proposition it expresses there *is false*.

For the sceptic (*) has *always the same truth conditions* and (*) is *always false*, therefore he has to give a *non-semantic* account of the context sensitivity that constitutes the initial data of the contextualist: why (*) appears true in everyday contexts and false in sceptical contexts? The sceptic can give a pragmatic explanation of this sensitivity along the following lines.

There are maxims that govern our conversations in everyday contexts and determine what we are warranted in asserting and what not given the purposes and interests we seek to satisfy by saying what we do in different contexts. Most of the time we implicitly assume that the participants in a conversation implicitly follow those maxims, that is to say, we assume that they will make conversational contributions such as required by the purposes and interest that direct the conversation.¹⁴ To explain the context sensitivity of knowledge sentences we can posit a conversational maxim that dictates the following:

(M) If S's belief that p tracks truth within a domain of relevant worlds fixed by a given conversational context, then it is appropriate to say that S know that p.

¹² Cf. for example DeRose, 1992, pp. 913-914 and Cohen, 1999, pp. 58-59.

¹³ I'm using here DeRose's understanding of the standards for how strong S's epistemic position must be, but the points I'm making can be rephrased using Lewis's and Cohen's understanding of standards of epistemic strength in terms of S's evidence ruling out alternatives.

¹⁴ This is what H. P. Grice called the "Cooperative Principle". Cf. Grice, 1975, p. 26.

The contextual variance of the domain within which S's belief that *p* has to track truth is not built into the meaning of (*) but into the conditions under which it is warranted to assert (*). It's not true that the truth conditions of (*) change from an everyday conversational context to a sceptical context. The truth conditions of (*) are always the same and (*) is always false, not only false in a sceptical context. (*) *seems* to change truth value in going from an everyday context to a sceptical context for the following reason. In everyday contexts, when we are placed in what we take to be optimal epistemic circumstances, (*) *appears to be true* because asserting it, plus the assumption that we are following (M), generates the *true implicature*¹⁵ that S's belief that *p* tracks truth in the relevant domain of alternatives. The generation of this implicature warrants our assertion of (*). But the implicatures generated by what someone says are not part of the truth conditions of what is said. What someone says may be false and still produce true implicatures. This is what happens with (*), it is false in everyday context but in the contexts in which, in virtue of (M), it generates true implicatures that warrant our assertion of (*) we confuse this warranted assertability with truth.

The sceptic recognises that \sim (*), which he thinks is true in all contexts, seems false in everyday context. He explains this appearance of falsehood as he explained away the appearance of truth of (*). In everyday contexts \sim (*) *appears to be false* because asserting it, plus the assumption that we are following (M), generates the *false implicature* that S's belief that *p* doesn't track truth in the relevant domain of possible worlds. The generation of this implicature makes the assertion of \sim (*) inappropriate or unwarranted and we confuse this with falsehood.

This sketched view that the context sensitivity of knowledge sentences is just a pragmatic phenomenon might look a bit artificial, but the problem is how the contextualist is going to show that it is wrong and that the sensitivity in question is a semantic phenomenon.

Keith DeRose has provided a response to the pragmatic objection. The response does not consist in arguing directly that the context-sensitivity of knowledge sentences indeed affects their truth conditions. What DeRose does is to argue that the pragmatic objection doesn't meet certain criteria of success that an objection of that

¹⁵ In Grice's sense, cf. Grice, 1975, p. 27.

type to a theory about the truth conditions of sentence forms in natural language must meet. The criteria in question has two points:

- a. If the appearance of falsehood of an assertion like $\sim(*)$ in a context is going to be explained away in terms of the assertion not fulfilling its conditions of warranted assertability, then $(*)$ *must also appear to be false*. Why? DeRose says that because when in a context $\sim(*)$ appears false and $(*)$ true there's no pressure or motivation at all to explain away the appearance of falsehood of $\sim(*)$ pragmatically: the best thing to do is to take the appearances at face value and say that $\sim(*)$ is false and $(*)$ is true. But when both $\sim(*)$ and $(*)$ seem false, then we are under the pressure of explaining *some* appearance of falsehood away¹⁶ in order to avoid an inconsistency. We can put this point by saying that for DeRose it is a necessary condition for the success of the pragmatic view that there are apparent inconsistencies to explain away.
- b. The appearance of falsehood of $\sim(*)$ is relegated to the false implicatures generated by asserting $\sim(*)$ in some contexts. Those implicatures arise because it would be assumed in those contexts that a certain maxim of conversation is observed. DeRose says that the maxim invoked must be a *general* one, it must cover an appropriately wide range of linguistic phenomena. If it is not general but designed *ad hoc* to apply, say, only to the case of $(*)$, then we could invent conversational maxims at will to explain away as a mere pragmatic phenomenon any apparent counter example to any account about the truth conditions of any sentence form in natural language.¹⁷

I think these criteria for judging the success of a pragmatic objection of the sort invoked by the sceptic are quite inconclusive. The sceptic's pragmatic view doesn't meet (a) because while $\sim(*)$ seems *false* $(*)$ seems *true* in an everyday context, so there's no pressure to explain $\sim(*)$'s appearance of falsehood away pragmatically in order to avoid an inconsistency. This probably shows that the pragmatic explanation lacks some sort of motivation but that doesn't show that it is incorrect, even more,

¹⁶ DeRose, 1999, p. 198

¹⁷ *Ibidem*. For simplicity's sake I've couched the criteria as criteria for explaining away an appearance of *falsehood* pragmatically, but it must be clear that they are criteria for explaining away the appearance of having a truth value (either true or false) pragmatically.

that doesn't deprive it of all plausibility. It is less clear that the sceptic's pragmatic view doesn't meet (b). How broad is the range of linguistic phenomena covered by (M)? Is it broad enough to make (M) count as a *general* maxim of conversation as opposed to an *ad hoc* one? I don't know how to answer these questions, but it seems to me that this unclarity as to whether the sceptic's pragmatic view meets DeRose's criteria (b) should not lead us to judge that that view is incorrect.

At this point I would restate the question I asked before: How can we determine whether it is the semantic or the pragmatic view the correct account of the context-sensitivity of knowledge sentences? Some contextualists have struggled in the search for a satisfactory answer to this question. Stewart Cohen, in one of his recent papers gives up the search and says that neither side of the dispute can demonstrate the correctness of its view to the other side. However, he thinks the "strongest argument" against the pragmatic sceptical view is, precisely, that it is sceptical.¹⁸ The fact that the pragmatic view is designed to support Cartesian scepticism is enough to prefer the contextualist semantic view which has the nice consequence that in everyday life we know basically everything we pre-philosophically thought we knew.

Cohen's attitude suggests the thought that the discussion between the contextualist and the sceptic is doomed to reach a standoff in which the "strongest argument" the contextualist can give against the pragmatic sceptical view consists in reminding us that it is sceptical (!). But is that standoff really unavoidable? I'll present Peter Unger's arguments for the view that the standoff is unavoidable, then I'll argue that Unger's view is mistaken and the contextualists's despair in facing the pragmatic objection too hasty, since there seems to be a quite straightforward argument to show that the context sensitivity of knowledge ascriptions *is* a semantic phenomenon.

b. Unger's hypothesis of semantic relativity

Unger thinks the debate between the sceptic and the contextualist is an instance of a more general dispute between two *semantic* approaches he calls

¹⁸ Cohen, 1999, p. 83. Cf. also his 1990, p. 168

‘invariantism’ and ‘contextualism’. To be an invariantist with respect to the semantics of a particular expression is to hold that there is a single standard for its correct application in any context in which it may be employed. To be a contextualist with respect to that expression is to hold that the standard for its correct application changes from context to context. The invariantist is not necessarily a sceptic. If the invariant standards for the use of expression X are set *low*, so that they are usually met, the invariantist will not be a sceptic about the use of expression X. The invariantist becomes a sceptic only when he sets the invariant standards so *high* that they are never met. In the following discussion I’ll equate the invariantist with sceptic invariantist, as Unger himself does.¹⁹

The dispute between invariantism and contextualism doesn’t arise with respect to all expressions of a natural language. There are expressions (like ‘here’ or ‘yesterday’) whose semantic value indisputably can be assigned only relative to a given context. And there are expressions whose semantic value is entirely independent of the contexts in which they are used.²⁰ The dispute invariantism/contextualism arises only in connection with expressions with respect to which it is *plausible* to say that they may determine one semantic value when evaluated independently of any particular context in which they may be used and other semantic value when evaluated relative to a particular context of use. ‘Know’ is one of such expressions.

Unger’s view is that with respect to the range of expressions concerning which the dispute contextualism/invariantism arises, there’s no way of determining “objectively” whether it is the invariantist or the contextualist theory the correct account of the expressions’s meaning. He calls this view the hypothesis of “semantic relativity”.²¹ In the areas of language in which this hypothesis applies our choice of a semantic account “will be somewhat arbitrary, not determined by objective facts,

¹⁹ Cf. for instance Unger, 1984, p. 51. It’s also worth noting that the contextualist is not necessarily an anti-sceptic. If the contextually varying standards happen to be so tough that they can’t never be met we would have a sceptical contextualist. However, in the discussion I’ll equate the contextualist with the anti-sceptical contextualist, as Unger also does.

²⁰ Though the extent to which this may be true has been questioned by Charles Travis, cf. his 1996 and 1997.

²¹ Unger, 1984, p. 15.

including facts of logic and language”.²² Let’s flesh out this alleged objective indeterminacy with an example.

Consider the following interchange:

A. Is there any wine left?

B. No, the bottle is empty.

For the contextualist the proposition expressed by what A said is: “Is there any wine left *to make one more drink?*” And so, in replying as he did B said something true: the bottle is empty *relative to the purpose of making another drink*, the few drops remaining in the bottle just aren’t enough. For the contextualist contextual parameters enter into the meaning of the sentences uttered by A and B. For the invariantist, on the contrary, such parameters do not play that semantic role, for him the proposition expressed by what A said is: “Is there any wine left *absolutely?*” And so, in replying as he did B said something false: the bottle is *not* empty, there are some few drops remaining. Of course the invariantist agrees that in saying what he says A gets B to focus on the contextually relevant question *Is there any wine left to make another drink?*, but he thinks this phenomenon concerns the pragmatics of communication and need not be explained as a “complex *semantic* operation or understanding”.²³

From the invariantist’s point of view if B had answered: ‘Yes, there is some wine left’, his assertion would have been *true* but *inappropriate*, it would have generated the false implicature that there is enough wine to make another drink, and that would be because it had violated some maxim of communication (perhaps the maxim: Be Relevant).²⁴ On the other hand, by saying ‘the bottle is empty’, B says something *false* but succeeds in conveying the information (in the form of true implicatures) that the maxims of cooperative communication required from him in the context where his conversation with A takes place.

The invariantist thinks the contextual parameters exploited by speakers in getting their audience to focus on a certain contextually relevant piece of information enter into the conditions of appropriateness of their conversational contributions, but not into the determination of the truth-conditions or propositions expressed by the

²² *Ibidem*.

²³ Unger, 1984, p. 26.

²⁴ Cf. Grice, 1975, p. 27.

sentences they utter. Those parameters “enter our thought only at another place, one removed from semantic grasping and evaluation”.²⁵ In this way the invariantist explains away pragmatically what the contextualist takes to be a semantic phenomenon. The availability for the invariantist of this pragmatic move seems to deprive the contextualist of the evidence from communication, interpretation and understanding that he could otherwise have used to support his view. Such evidence from language use can be interpreted pragmatically, as the invariantist does, or semantically, as the contextualist does, and there are no “objective” facts that decide which interpretation is the right one.

One might think that Unger’s hypothesis of semantic relativity is just wrong because a certain divergence of invariantism from our linguistic practices decides the dispute in favour of contextualism. The divergence is that from the invariantist point of view many of the things we say in ordinary life have the opposite truth value from the one we commonly think they have. We think that when B says “The bottle is empty” he speaks the truth but for the invariantist he says something false. In this respect contextualism is aligned with ordinary thoughts about the truth values of what we say: we think B’s utterance is true and contextualism gives the same verdict. Unger says that this divergence of invariantism from and alignment of contextualism to ordinary practice doesn’t show that contextualism is the objectively correct account of *the proposition expressed by what B says*. The agreement of contextualism with linguistic practice only shows that we have entrenched contextualist habits of interpretation, understanding and attribution. The entrenchment of those habits makes our ordinary judgements about the truth value of what we say diverge from the verdicts ensued by invariantism:

...owing to our contextualist habits of attribution, the psychological posits that the invariantist requires are also strange and surprising, just as are the semantic posits of his total view. But none of this strangeness need indicate any greater objective truth for the more habitual mode of attribution or for the more common-sensical judgements encouraged thereby.²⁶

The divergence of invariantism from and the agreement of contextualism with

²⁵ Unger, 1984, pp. 25-26.

²⁶ *Ibid.*, p. 36.

habitual practice don't count as objective facts to decide the dispute in favour of contextualism because both, the divergence and the agreement, can also be interpreted pragmatically or semantically. The semantic interpretation favours the contextualist, but the pragmatic one favours the invariantist. On the pragmatic interpretation the contextualist, common-sensical assignments of truth values to the *proposition* expressed by what we ordinarily say are wrong (the correct assignments are the invariantist ones), but we think that they are correct because what we are really evaluating and assigning a truth value to is not the proposition strictly expressed by our words but the *information conveyed in a given context* by our use of those words. For the invariantist the truth value of the proposition expressed by B's utterance: "The bottle is empty" is *false*, but we think that it is *true* because what we evaluate is not the strict propositional content of that sentence but what it communicates in the context in which it is used: that there's not enough wine left to make another drink.

It seems again that any facts about linguistic practice that could be thought to favour contextualism over invariantism can be interpreted semantically or pragmatically. Under the semantic interpretation the facts favour contextualism, under the pragmatic interpretation they favour invariantism, and there doesn't seem to be anything to decide in turn which interpretation of the facts is the correct one. I think this is what Unger has in mind when he says that there are no "objective facts", "*including facts of logic and language*",²⁷ that settle the debate between contextualism and invariantism.

Unger's hypothesis about the objective indeterminacy infecting the debate between contextualism and invariantism may be seen as the explanation of why the discussion between the contextualist about knowledge ascriptions and the Cartesian sceptic ends up in the standoff we saw some contextualists think it does. The sceptic raises the objection against the contextualist that the context sensitivity of knowledge sentences is not a semantic but a pragmatic phenomenon, the contextualist finds himself unable to show that it is a semantic phenomenon. His inability isn't due to his missing some facts that would establish the semantic character of the

²⁷ Unger, 1984, p. 15.

phenomenon in question, it is rather that there aren't any such facts. The context sensitivity of knowledge ascriptions can be interpreted pragmatically, then it will be consistent with the sceptic's invariantist theory about the truth conditions of such ascriptions; or it can be interpreted semantically, then it will favour the contextualist theory about the truth conditions of such ascriptions. Any facts about linguistic behaviour that could be thought to be usable to decide which interpretation is correct will seem to be in turn susceptible of being construed pragmatically or semantically, favouring the sceptic under the former construal and the contextualist under the latter. The result is that there are no facts that determine whether it is the contextualist or the sceptic who is right.

c. It is a semantic phenomenon

I think Unger's account of what goes on in the discussion between the contextualist and the sceptic is ill grounded. There seems to be a very straightforward way of determining whether the context sensitivity of knowledge sentences is a semantic or a pragmatic phenomenon.²⁸

We can state the sceptic's position by saying that he considers that an utterance of the sentence

(*) S knows that p

expresses a complete proposition independently of the context in which it is uttered, for notational convenience let's represent that proposition as follows:

[E] $(\forall PW)$ (S's Bp tracks truth with respect to PW)²⁹

[E] says that S's belief that p tracks truth in every possible world, and *a fortiori*, in every sceptical world.³⁰ The contextualist says that an utterance of (*) does not express a complete proposition independently of the context in which it was uttered,

²⁸ In what follows I'll draw on ideas presented by Recanati in his 1989.

²⁹ Here again I'm using DeRose's understanding of standards for how strong S's epistemic position must be, but the point I'm making can be easily rephrased using Lewis's or Cohen's understanding of standards of epistemic strength in terms of S's evidence ruling out alternatives.

³⁰ Of course, since S's belief that p doesn't track truth in sceptical worlds the proposition that according to the sceptic is expressed by (*), viz. [E], is always false.

detached from context an utterance of (*) expresses a proposition that we can represent as follows:

[C] (\forall PW in D) (S's Bp tracks truth with respect to PW)

This is not a complete proposition because the variable D for the domain of relevant possible worlds within which S's belief must track truth hasn't been specified, and it can only be specified relative to a given context. Therefore, independently of its context of utterance the proposition expressed by an utterance of (*) cannot be assigned a truth value for the general reason that a truth value cannot be assigned to an incomplete proposition.

To see how we can determine which of the two accounts of (*) is the correct one we have to realise first that the disagreement between the two accounts concerns what is the proposition expressed by *an utterance of (*)*. We may wonder why shouldn't we suppose that the speaker who produced the utterance knows what is the proposition *he is expressing*? If we want to know what proposition *S expresses* in uttering (*) the best thing to do is to ask S what he means in uttering (*). After all what we are interested in is *what propositions people express* in making knowledge ascriptions. Compare with the following case. Suppose John says 'All students came to Paris', one could come up with the theory that the proposition John expressed was that all students that have ever existed on Earth came to Paris. But of course that's not the proposition John expressed, and we can corroborate this conjecture by asking John whether it was that proposition the one he expressed.

Similarly, if we ask to our subject B above: "Did you express the proposition *the bottle is empty absolutely* by saying 'the bottle is empty'?" He would certainly answer: "No, what I said, what I wanted you to take as the truth evaluable item I was expressing, was the proposition *the bottle is empty relative to the purpose of making another drink*". We would get a similar reaction with respect to (*). If we ask a subject: "Did you express the proposition *S's belief that p tracks truth in all possible worlds* by uttering (*)?" He would certainly answer: "No, what I said, what I wanted you to take as the truth evaluable item I was expressing, was the proposition *S's belief tracks truth in some possible worlds*". It may be difficult for a speaker to identify all the possible worlds picked out by his utterance of (*) in a given

conversational context; but with certainty he would at least reject that the proposition he expressed had the *form* [E] above, he would rather say that it has the *form* [C].

The intuitive appeal of these facts suggests that the proposition expressed by S in uttering (*), or at least the form of that proposition, is known, or may be known on reflection, by S. This recommends that “any tentative identification of what is said [by uttering (*)] has to be checked against the speaker’s intuitions”.³¹ As a result any identification of the proposition expressed by the speaker, or at least of the form of such proposition, inconsistent with the speaker’s beliefs about what that proposition is, or about what the form of that proposition is, must be rejected; unless some reason is given that defeats the assumption that the speaker knows what proposition he is expressing or what is the form of that proposition.

The foregoing considerations take us back to one aspect of Schiffer’s complaints to contextualism. We saw that one of Schiffer’s complaints is that the content that the contextualist semantics attributes to knowledge ascriptions typically is *not transparent* to speakers and this is reflected in their incapacity to articulate such content. I granted that but suggested that the lack of transparency may be corrected up to a certain point if speakers carry out a more or less sophisticated reflection on what the content of their knowledge ascriptions amounts to. In answering to the pragmatic objection I’ve developed this suggestion. I have said that if the speaker’s reflection about the content of their knowledge ascriptions is prompted by the question whether it is [E] or [C] that depicts the *form of the content* of their knowledge ascriptions, they can come to recognise that such content is represented by [C], as the contextualist semantics holds. This reflective recognition still falls short from grasping the *full content* that the contextualist semantics attributes to each particular knowledge ascription, that grasping would require knowledge of the different domains of possible worlds defined by each individual knowledge ascription, and such knowledge is too much to expect. However, my argument against the pragmatic objection only requires that on reflection speakers

³¹ Recanati, 1989, p. 313. This is what Recanati calls the “Availability Principle” which says that in making a decision concerning what of the information conveyed by an utterance is part of the proposition expressed by the speaker, we should always try to preserve the speaker’s intuitions on the matter (*cf.* pp. 310-15). Recanati doesn’t discuss knowledge ascriptions but I’m trying to apply his discussion to the debate between the sceptic and the contextualist.

recognise [C] as the correct representation of the *form* of the propositional content of their knowledge ascriptions; that's enough evidence for the contextualist semantics.

The previous discussion gives us a criterion to decide whether the context sensitivity of knowledge ascriptions is a semantic phenomenon or not. The fact that on reflection speakers would reject the invariantist interpretation [E] of the *form of the proposition* expressed by utterances of (*) and accept the contextualist interpretation [C], shows that the contextual determination of a domain of possible worlds is necessary to determine the proposition expressed by a speaker's utterance of (*). And this means that the context sensitivity of knowledge ascriptions is a *semantic* phenomenon.

I have argued that evidence can be gathered for the semantic status of the context sensitivity of knowledge ascriptions. So we can grant to contextualists that knowledge ascriptions are context sensitive and that this context sensitivity is semantic in character. Even granting this there is a further question which is decisive for the success of the contextualist solution of the sceptical paradox, it is the question whether the context sensitivity of knowledge ascriptions, as construed by contextualists, has the consequences they expect from it. We saw in Chapter II.1.c that contextualists have a particular view about what contextual facts determine the propositional content of knowledge ascriptions and how that determination comes about. The question I shall explore now is if the contextualists particular view about how the context sensitivity of knowledge ascriptions operates has the consequences they want it to have.

3. Elusive knowledge and impossible theory

a. Elusive knowledge?

We saw in chapter II that part of the contextualist semantics for knowledge ascriptions is the idea that there are conversational rules or mechanisms that control how strong S's epistemic position must be to know by controlling the size of a domain of epistemically relevant possible worlds. Those rules define which conversational events constitute changes in the conversational context that have an

impact on the content of knowledge ascriptions. DeRose's Rule of Sensitivity, recall, says that the event of *saying* 'S knows that *p*' makes it the case that for S to know that *p* S's belief that *p* must track the truth up to the closest $\sim p$ worlds, i.e. S's belief that *p* must be sensitive. The standard for knowing that *p*, i.e. the content of 'S know that *p*', is determined by the very act of saying that S knows that *p*. This is why when the belief that *p* is *insensitive* any *assertion* of 'S knows that *p*' is bound to result in falsehood.³²

Lewis also thinks that conversational events of asserting and mentioning affect the content of knowledge ascriptions. His Rule of Attention embodies this idea, it says that merely attending to a certain hypothesis incompatible with what we believe makes it a relevant one:

No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in *this* context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative.³³

In Lewis's view the psychological event of attending to a certain hypothesis incompatible with *p* is enough to change the truth conditions of 'S knows that *p*' in such a way as to require the elimination of the hypothesis considered. What we say in saying 'S knows that *p*' is different before and after we consider the hypothesis in question *just in virtue of* our having thought of that hypothesis. When the envisioned hypothesis cannot be eliminated, as it happens with sceptical hypothesis, then our having thought about it causes our assertions and thoughts about knowledge to have truth conditions that cannot be satisfied, and so those assertions and thoughts are bound to be false:

Simply *mentioning* any [sceptical hypothesis] aloud or even in silent thought, is a way to attend to the hitherto ignored possibility, and thereby render it no longer ignored, and thereby create a context in which it is no longer true to ascribe the knowledge in question to yourself or others.³⁴

³² DeRose, 1995, p. 39.

³³ Lewis, 1996, p. 559. Lewis' emphasis.

³⁴ *Ibid.*, p. 562. Lewis' emphasis.

The examples above illustrate the contextualist's view about *what* events or facts in a conversational context have an impact on the content of knowledge ascriptions and *how* that impact comes about. The view is that events of saying or thinking that S knows that *p*, or of mentioning or considering a hypothesis incompatible with *p*, make the content of knowledge ascriptions shift, and those events have this impact on knowledge ascriptions just in virtue of occurring. For the sake of brevity I'll call this view the "simply say it" view.

It's important to emphasise that the "simply say it" view is *essential* to the contextualist resolution of the sceptical paradox. Contextualists set up the sceptical problem as arising in a conversational context in which the sceptic asserts, or somehow makes us consider, the premises of his argument. The contextualist says that the *mere assertion* of the premises of the sceptical argument CP makes the content of knowledge ascriptions shift in such a way as to require unmeetable standards in strength of epistemic position. Given those standards most of our claims to know anything in that conversational context will be false. But later, when we come to consider the conclusion of the argument (CP4), the *mere assertion* of (CP4) switches the standards for knowledge back to what they were before we started considering the sceptical hypothesis. Those reinstated standards are usually meetable, we know a lot relative to them; consequently we end up having intuitions that (CP4) is false. These intuitions arise only because of the shifts in standards caused by our simply *saying* or *thinking* what we did, or what the sceptic did. For *this explanation* to work it must be very easy to shift epistemic standards, as easy as simply mentioning or considering H_T , or simply claiming or thinking that S knows or doesn't know H_T . The "simply say it" view embodies the easy mechanisms of standard shifting that the contextualist needs for his resolution of the sceptical paradox.

I want to object to the "simply say it" view that it is unfaithful to the ways we actually react to challenges to our knowledge claims and to some plain claims to know. It looks as if the "simply say it" view were correct then simply saying that one doesn't know that a sceptical hypothesis doesn't obtain, for example that one is not a brain in a vat, would prompt agreement among our audience, since according to the "simply say it" view merely saying or considering such a hypothesis is enough to raise the standards for knowledge in such a way that any assertion or thought to the

effect that someone knows that it doesn't obtain is bound to be false. But the reaction to saying that one doesn't know that the sceptical hypothesis doesn't obtain is not agreement, "the typical response...is to find preposterous the mere assertion that one doesn't know one is not a brain in a vat. People require more coaxing and more explanation, although at the end they might agree to it. They need reasons to think that we lack such knowledge".³⁵

The "simply say it" view also seems to imply that my saying 'I know I'm not a brain in a vat' would always bring on dissent since according to that view any such assertion is bound to be false. But this is not so either. Suppose we are playing one of those card-games where one has to answer question on general knowledge, after failing six straight times to provide the correct answer I desperately say: "At least I know this: I'm not a brain in a vat". This mere assertion will not make the standards of knowledge go up in such a way that my assertion will be false, as the "simply say it" view implies; that's why none of my playmates would disagree with my claim to know that I'm not a brain in a vat.

I think the unfaithfulness of the "simply say it" view to our reactions to knowledge claims has its roots in the fact that that view conflates the *epistemic relevance* of a possibility of error with *our attending* to it, and the *epistemic irrelevance* of a possibility with *our not attending* to it. And this is wrong: The psychological fact of attending to a certain possibility of error doesn't settle by itself alone the normative question whether that possibility deserves to be taken into account in setting standards for knowledge. It seems obvious that we can attend to a certain possibility of error and deem it epistemically irrelevant.³⁶

The conflation of psychological facts and normative questions in the "simply say it" view makes it too easy to raise epistemic standards, as the contextualist requires for his resolution of the sceptical paradox. But raising standards for knowledge is not as easy as that view says, if it were that easy we would give up knowledge ascriptions more readily than we do.³⁷

³⁵ Feldman, 1999, p. 100.

³⁶ On this point *cf.* Williams, 2000, p. 83.

³⁷ Lewis anticipates the criticism I've been making to the "simply say it" view and he replies as follows:

You will say that no amount of attention can, by itself turn [the far-fetched sceptical hypotheses] into relevant alternatives....If you say this, we have reached a standoff. I started

I've argued that the 'simply say it' view is implausible, however I've pointed out that it is essential to the contextualist resolution of the sceptical paradox. Now I will argue that the "simply say it" view puts contextualists in a quandary for justifying claims they make about the truth of knowledge ascriptions relative to everyday contexts. Given the indispensability of the "simply say it" view for the purposes of contextualists they are not in a position to avoid the quandary I'm going to describe by just dropping out the "simple say it" view.

b. Impossible theory

The "simply say it" view entails that merely attending or mentioning sceptical hypothesis H_T "destroys" our knowledge by raising epistemic standards up to an unmeetable level. Since every epistemology student has to come across sooner or later sceptical hypotheses, epistemology becomes the quite peculiar discipline that destroys its subject matter:

Do epistemology. Let your fantasies rip. Find uneliminated possibilities of error everywhere. Now that you are attending to them, just as I told you to, you are no longer ignoring them, properly or otherwise. So you have landed in a context with an enormously rich domain of potential counter-examples to ascriptions of knowledge. In such an extraordinary context, with such a rich domain, it never can happen....that an ascription of knowledge is true....That is how epistemology destroys knowledge.... The epistemology we've been doing, at any rate, soon became an investigation of the ignoring of possibilities. But to investigate the ignoring of them was *ipso facto* not to

with a puzzle: how can it be, when his conclusion is so silly, that the sceptic's argument is so irresistible? My Rule of Attention [is] ...built to explain how the sceptic manages to sway us...If you continue to find it eminently resistible in all contexts, you have no need of any such explanation. We just disagree about the explanandum phenomenon. (Lewis, 1996, p. 561)

This reply doesn't quite answer the objection I've presented. The objection, in effect, is that simply attending to a possibility cannot by itself make it epistemically relevant. Raising epistemic standards is not that easy, and I've pointed out that we can confirm this by looking at how we ordinarily react to knowledge claims. But to object in this way we need not disagree with the contextualist about what his "simply say it" view is meant to be an explanation of, namely, the compelling character of the sceptical argument. The objector may agree with the contextualist that the argument is very compelling, perhaps even irresistible, and still disagree in that the argument's persuasiveness is to be explained by appeal to the "simply say it" view. The source of the compelling character of the argument may be something different from the operation of the conversational rules posited by the contextualist's "simple say it" view.

ignore them...it will be inevitable that epistemology must destroy knowledge. That's how knowledge is elusive. Examine it, and straightway it vanishes.³⁸

Epistemology [is] an investigation that destroys its own subject matter.³⁹

In Chapter II we saw that contextualists not only want to show that the concept of knowledge is not incoherent or paradoxical, they also want to say that in some contexts that concept has true applications but in others don't, in other words, they want to say that in some contexts we know but in others we don't. Here is an explicit statement of this claim

On the present solution, claims to know ordinary propositions are true according to ordinary low standards but false according to the highly inflated standards [operative in sceptical contexts].⁴⁰

The discursive context of Epistemology is one in which sceptical possibilities are considered and therefore, by definition, a sceptical context in which Cartesian scepticism is true, and nobody knows anything about the world. But still, in everyday conversational contexts where sceptical possibilities are ignored, we know a lot. I want to enquire how the contextualist can justify this latter claim given that he endorses the "simply say it" view.

Notice that in arriving at the claim that in everyday contexts where sceptical possibilities are ignored people know things, the contextualist occupies the reflective position of an epistemologist. In doing Epistemology the contextualist reflects on, and so attends to, the previously ignored sceptical possibilities, and in attending to them he makes knowledge vanish as the "simply say it" view says. But then how can the contextualist, in the reflective context he occupies as an epistemologist, be justified in endorsing the claim that people know in everyday contexts? It looks as if he couldn't, for how could he see as true, *from his reflective position*, any knowledge ascription? No matter whether made before or after his reflective stance is reached he can look at a knowledge ascription only from his reflective position, and from this position any knowledge ascription he can entertain will appear to be *false*. It looks as if the contextualist could not be justified in endorsing the claim that people know in

³⁸ Lewis, 1996, pp. 559-560

³⁹ *Ibid.*, p. 550.

⁴⁰ DeRose, 1995, p. 40.

everyday contexts from the reflective position he has to occupy in order to formulate both this claim itself and the “simply say it” view. Lewis is perfectly aware of this quandary:

If the story I told was true, how have I managed to tell it?...Does not my story deconstruct itself? I said: S knows that P iff S's evidence eliminates every possibility in which not-P –Psst!– excepts for those possibilities that we are properly ignoring. That ‘psst’ marks an attempt to do the impossible –to mention that which remains unmentioned. I am sure you managed to make believe that I had succeeded. But I could not have done. And I said that when we do epistemology, and we attend to the proper ignoring of possibilities, we make knowledge vanish. First we do know, then we do not. But I had been doing epistemology when I said that. The uneliminated possibilities were *not* being ignored –not just then. So by what right did I say even that we used to know?⁴¹

So, the problem seems to be that from the reflective perspective the contextualist needs to adopt to formulate his theory about knowledge ascriptions, he cannot be justified to endorse the anti-sceptical part that theory which says that pre-reflectively, before considering sceptical hypotheses, we know many things.

We can look at this problem in the following alternative way. To be justified in endorsing, *in the reflective position he occupies in which sceptical hypotheses are mentioned*, the claim that in everyday contexts where sceptical hypotheses are not mentioned people know things, the contextualist must be justified in endorsing this conditional:

ST: If the sceptical hypothesis had not been mentioned we would have been able to know things about the world.

Can contextualists be justified in endorsing (ST)? I think they can't. I'll give my reasons in detail.

To be justified in believing (ST) is to be justified in believing that whenever its antecedent is true its consequent is true as well, i.e. that whenever sceptical hypotheses are not mentioned we can know things about the world. I'm going to

⁴¹ Lewis, 1996, p. 566. DeRose also points out that this is a problem for the contextualist, cf. his 1995, p. 40, footnote 36.

show that contextualists are not justified in believing that *whenever* sceptical hypotheses are not mentioned we can know things about the world.

Consider an everyday conversational context where sceptical hypotheses are not mentioned, in such a context the contextualist claims that people know many things, and his basis for claiming this is that such a conversational context determines epistemic standards that people usually meet. This already looks like a good justification for claiming that *whenever* sceptical hypotheses are not mentioned we can know things. However matters are more complicated. When the contextualist says that everyday contexts determine epistemic standards that are usually meetable, what he is saying is that those conversational contexts determine domains of possible worlds within which people's beliefs must track truth in order for them to know;⁴² and people usually meet those standards because their beliefs usually track truth within those relevant domains of possible worlds.

All contextualists must accept that the possibility that actually obtains, *viz.* the actual world, is *always* one of the relevant possible worlds within which people's beliefs must track truth in order to count as knowledge. Lewis's Rule of Actuality captures this important requirement which, as he points out, amounts to the idea that "only what is true is known".⁴³ This entails that if the contextualist's claim that people know relative to the standards of everyday contexts is going to be true, then people's beliefs must track truth in the actual world. But notice that the truth of the contextualist's claim also requires that the relevant domains of possible worlds within which people's beliefs must track truth defined by everyday contexts do not include any sceptical world, i.e. any world in which a sceptical hypothesis obtains. The reason for this is that people's beliefs do not track truth in sceptical worlds, and therefore, if sceptical worlds were in the relevant domains of possible worlds defined by everyday contexts it wouldn't be true that people know relative to the standards set by everyday contexts. Now, the contextualist's justification for thinking that sceptical worlds are excluded from the domains of relevant worlds defined by everyday contexts is that in those contexts sceptical hypothesis *are not mentioned*. But the not mentioning of a sceptical possibility does not exclude one crucial way in

⁴²Once again I'm using DeRose's understanding of standards for how strong an epistemic position must be, but what I say here can be rephrased using Lewis's or Cohen's understanding of standards of epistemic strength in terms of evidence ruling out alternatives.

⁴³ Lewis, 1996, p. 554.

which a sceptical world may be part of the domain of relevant worlds defined by everyday contexts, *viz.* by being the actual world itself (!) If the actual world were a sceptical world in which all the everyday conversational contexts take place, then the fact that in such conversational contexts the sceptical hypotheses are *not mentioned* wouldn't exclude the sceptical world from the domain of relevant worlds defined by everyday contexts, simply because the sceptical world would be the actual world and the actual world is *always* part of any domain of relevant worlds. This means that the contextualist's claim that relative to the epistemic standards of everyday contexts people know things requires the *assumption* that the actual world is not a sceptical world.

The foregoing argument may look a bit intricate, we can improve its clarity summarising it in the following way. The truth of the proposition

- (i) sceptical worlds are excluded from the domain of relevant worlds defined by everyday contexts

is a necessary condition for the truth of the proposition

- (ii) S knows things relative to the standards of everyday contexts

The not mentioning of sceptical hypotheses in everyday contexts does not preclude one way in which a sceptical world may be in a domain of relevant worlds defined by everyday contexts: by being the actual world. Therefore, endorsing (i) involves assuming that:

- (iii) the actual world is not a sceptical world.

Since contextualists believe (ii), they must endorse (i) and therefore assume (iii).

Now, how does all this bear on the question whether contextualists are justified in endorsing conditional (SC)? I said that if they are going to be justified in endorsing, from the reflective stance they occupy, the part of their theory which says that in everyday contexts, where sceptical hypotheses are not mentioned, people know things, they must be justified in endorsing the conditional

(SC) If sceptical hypotheses had not been mentioned we would have been able to know things about the world,

and being justified in endorsing (SC) is being justified in believing that *whenever* sceptical hypotheses are not mentioned we can know things of the world. Now we are in position to explain why contextualists can't justifiably believe this from the reflective stance they occupy.

Not mentioning sceptical hypotheses does not *secure* that the domains of possible worlds defined by everyday contexts exclude sceptical worlds; as we have seen there's one way in which a sceptical world may be part of the relevant domains of possible worlds which is not precluded by the fact that sceptical hypotheses are not mentioned: being the actual world itself. If the actual world, where the contexts in which sceptical hypotheses are *not mentioned* take place, is a sceptical world then *ipso facto* the domains of possible worlds defined by everyday contexts include one sceptical world: the actual world. This would be a case in which sceptical hypotheses are not mentioned and nevertheless people cannot know anything about the world, which is a counterexample to (SC). In order to secure that *whenever* sceptical hypotheses are not mentioned people can know things of the world, the contextualist must make assumption (iii) identified above: that the actual world is not a sceptical world. The problem is that *in the reflective stance he occupies* the contextualist is not entitled to make that assumption. In his reflective context making that assumption would be claiming to know something about the world *viz.* that it is not a sceptical world, but his own theory says that in that reflective context where sceptical hypotheses are considered nobody knows *anything* about the world. This is why in his reflective context the contextualist is not entitled to make the assumption about the actual world he would have to make in order to justifiably believe from that reflective perspective that if sceptical hypotheses had not been mentioned we would have been able to know things about the world.

To put the point another way. Contextualists can be justified in endorsing the part of their theory that says that in everyday contexts, where sceptical hypotheses are not mentioned, people know things, only by assuming that *they*, contextualists, know something about the world (*viz.* that it is not a sceptical world) in a context in which

their own theory says they, as epistemologists, don't know *anything* about the world. The part of the theory that grants Cartesian scepticism in reflective contexts, such as the one they occupy in formulating their own theory, makes impossible for contextualists to establish and be justified in endorsing the part of their theory that purports to rescue the truth of everyday knowledge ascriptions.

In this section I have examined the question whether the particular way in which contextualists construe the context sensitivity of knowledge ascriptions gives them the consequences they want from it. The results of my examination have been negative.

First, I argued that the contextualists' view about the contextual factors that determine the content of knowledge ascriptions, what I called the "simply say it" view, is implausible and seems to conflate psychological facts associated with mentioning and saying something, with normative questions about the epistemic relevance of possibilities of error. A consequence of this conflation is that it becomes too easy to shift epistemic standards, but as a matter of fact, changing epistemic standards is not that easy. However, as I pointed out, the "simply say it" view is essential to the contextualist solution of the sceptical paradox. The solution of the paradox says that our conflicting intuitions concerning the premises and the conclusion of the sceptical argument arise because of conversational contextual shifts that change the operative epistemic standards, but the contextual shifts in question are just those described in the "simply say it" view. Given the problems infecting this view we cannot accept the solution of the paradox that makes such an essential use of it. Secondly, I argued that the "simple say it" view makes it impossible for contextualists to justifiably endorse the part of their theory that says that relative to the epistemic standards of everyday contexts people know things about the world. The reason is that for them to justifiably believe this they would have to make an assumption about what the actual worlds is like, but they are not entitled to make such assumption in the reflective context they occupy because their "simply say it" view entails that in that context *they* don't know anything about the world.

Ironically, the contextualists' view about how the context sensitivity of knowledge ascriptions actually operates doesn't give them what they expected from it.

Conclusion

I began this thesis distinguishing different lines of sceptical argument that pose different sceptical problems. Then I isolated one sceptical argument on which contextualists focus their attention, *viz.* argument CP. I stressed the point that even with respect to CP at least two general sceptical problems can be identified. First, the most direct sceptical challenge of demonstrating that the sceptical conclusion is false and people possess knowledge of the world; and secondly the challenge of showing that the concept of knowledge doesn't yield paradox or incoherence. Contextualists officially focus on the second sceptical challenge but we saw that eventually they reveal that they don't want to claim just that the concept of knowledge is not incoherent but also that it actually has true applications, that people actually know in everyday contexts. What they seek is "a resolution that rescues common sense from the sceptical worries posed by the paradox"¹, and rescuing common sense means precisely showing that people are right in thinking that in everyday life they can know things about the world.

In order to solve the paradox and rescue everyday knowledge contextualists devise a semantic machinery I examined in detail in chapter II. This semantic theory depicts knowledge ascriptions as having truth conditions or semantic content sensitive to certain facts of the conversational context where those ascriptions occur. In chapter III I examined two crucial questions about that semantic theory. First the question of how can it be shown that knowledge ascriptions *are* context sensitive and that such sensitivity is a *semantic* matter that affects their truth conditions. In sections 1 and 2 of chapter III I presented two objections (Schiffer's objection and the "pragmatic objection") that in different ways try to undermine the idea that there is a semantic context sensitivity of knowledge ascriptions. I argued that such objections can be answered and that we can grant to the contextualist that there is such a thing as a semantic context sensitivity of knowledge ascriptions.

Then I went on to examine a second crucial question concerning the contextualist semantic theory of knowledge ascriptions, the question whether the contextualists' particular view about the contextual facts that determine the

¹ Cohen, 1988, p. 113.

content of knowledge ascriptions gives them the results they want from it, *viz.* solving the sceptical paradox and rescuing the truth of everyday knowledge ascriptions. I argued that disappointingly their view about what are the contextual facts that knowledge ascriptions are supposed to be sensitive to doesn't yield those results. Such view (which I called the "simply say it" view) makes their resolution of the sceptical paradox implausible and makes it impossible for them, as epistemologists, to defend their claim that relative to epistemic standards of everyday contexts knowledge ascriptions are true.

The conclusion of this work is then that, for reasons internal to their own account of the context sensitivity of knowledge ascriptions, contextualists fail to solve the sceptical paradox and fail to justify their claim that relative to everyday contexts people can know things about the world. I think that of these two failures the second one strikes us as the most disappointing. This is perhaps because when we are first confronted with the sceptical argument what we feel is not that it exhibits our concept of knowledge as incoherent or paradoxical, our first reaction is rather to regard the conclusion of the argument as *false*. We ordinarily think we know many things about the world and we, in effect, want a solution of the argument that "rescues" the truth of this common sense reaction. We expect a solution that shows how after all we know all those things we ordinarily think we know. This means that what we expect is a solution of the argument that tackles the most direct sceptical challenge: demonstrate that the sceptical conclusion is false. This is why the failure of contextualists to justify the claim that we have knowledge, at least relative to the epistemic standards of everyday contexts, looks so disappointing to us: it is really a failure to answer the most direct sceptical challenge. If I'm right in that the sources of that failure are internal to the contextualist construal of the context sensitivity of knowledge ascriptions, then contextualism not only leaves the most direct sceptical challenge with us but predicts that this challenge cannot be successfully met in the context of epistemological inquiry.

Bibliography

Austin, J. L. [1946] "Other Minds", reprinted in his *Philosophical Papers*, Urmson J. O. & Warnock G. J. [Eds.], Oxford: Oxford University Press, 1990.

Bach, K. [1999] "The Semantic-Pragmatic Distinction: What It Is and Why It Matters". In *The Semantics/Pragmatics Interface from Different Points of View*. Oxford: Elsevier Science, 1999.

Clarke, T. [1972] "The Legacy of Skepticism". *The Journal of Philosophy*. Vol. 69, no. 20, 1972.

Cohen, S. [1987] "Knowledge, Context and Social Standards". *Synthese*. Vol. 73, 1987.

_____ [1988] "How to be a Fallibilist". In Tomberlin, J., ed. *Philosophical Perspectives*, 2: Epistemology, 1988.

_____ [1990] "Scepticism and Everyday Knowledge Attributions". In Ross, G. & Roth, M. eds. [1990]

_____ [1998] "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier and the Lottery". *Australasian Journal of Philosophy*. Vol. 76, no.2, 1998.

_____ [1999] "Contextualism, Scepticism and the Structure of Reasons". In Tomberlin, J. ed., [1999].

Davis, M. [2000] "Externalism and Armchair Knowledge". In Boghossian, P. & Peacocke, C. eds. *New Essays on the A Priori*. Oxford: Clarendon Press. 2000.

DeRose, K. [1992] "Contextualism and Knowledge Attributions". *Philosophy and Phenomenological Research*. Vol. 52, no. 4, 1992.

_____ [1995] "Solving the Sceptical Problem". *The Philosophical Review*. Vol. 104, no. 1, 1995.

_____ [1996] "Relevant Alternatives and the Content of Knowledge Attributions". *Philosophy and Phenomenological Research*. Vol. 56, no. 1, 1996.

_____ [1999] "Contextualism: An Explanation and Defence". In Greco, J. & Sosa, E. [eds.] *The Blackwell Guide to Epistemology*. Oxford: Basil Blackwell, 1999.

Descartes, R. [1641] *Meditations on First Philosophy*. In J. Cottingham, R. Stoothoff, D. Murdoch [eds.] *The Philosophical Writings of Descartes. Volume II*. Cambridge: Cambridge University Press, 1984.

Dretske, F. [1970] "Epistemic Operators". *The Journal of Philosophy*. Vol. 67, no. 24, 1970.

_____ [1981] "The Pragmatic Dimension of Knowledge". *Philosophical Studies*. Vol. 40, no. 3, 1981.

Feldman, R. [1999] "Contextualism and Scepticism". In Tomberlin, J., ed., [1999].

Goldman, A. [1967] "A Causal Theory of Knowing". *The Journal of Philosophy*. Vol. 64, no. 12, 1967.

_____ [1976] "Discrimination and Perceptual Knowledge". *The Journal of Philosophy*. Vol. 73, no. 20, 1976.

Grice, H. P. [1967] "Logic and Conversation". Reprinted in his *Studies in the Way of Words*. Massachusetts: Harvard University Press, 1991.

Heller, M. [1999] "The Proper Role for Contextualism in an Anti-luck Epistemology". In Tomberlin, J., ed., [1999].

Kaplan, D. [1989] "Demonstratives". In Almog, J., Perry, J., Wettstein, H. [eds.] *Themes from Kaplan*. New York: Oxford University Press, 1989.

Kaplan, M. [2000] "To What Must an Epistemology Be True? *Philosophy and Phenomenological Research*. Vol. 61, no. 2, 2000.

Klein, P. [1995] "Skepticism and Closure: Why the Evil Genius Argument Fail". *Philosophical Topics*. Vol. 23, no. 1, 1995.

Lewis, D. [1996] "Elusive Knowledge". *Australasian Journal of Philosophy*. Vol. 74, no. 4, 1996.

Hale, B. & Wright C. eds. [1997] *A Companion to the Philosophy of Language*. Oxford: Blackwell, 1997.

Heal, J. [1997] "Indexical Predicates and Their Uses". *Mind*. Vol. 106, no. 424, 1997.

Martin, C. B. & Deutscher, M. [1966] "Remembering". *The Philosophical Review*. Vol. 75, 1966.

Perry, J. [1997] "Indexicals and Demonstratives". In Hale & Wright, eds. [1997]

Recanati, F. [1989] "The Pragmatics of What is Said". *Mind and Language*. Vol. 4, no. 4. 1989.

Ross, G. & Roth, M. eds. [1990] *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht: Reidel Publishing Company, 1990.

Nozick, R. [1981] *Philosophical Explanations*. Cambridge: Harvard University Press, 1981.

Schiffer, S. [1996] "Contextualist Solutions to Skepticism". *Proceedings of the Aristotelian Society*. Vol. 96, no. 3, 1996.

Sosa, E. [1999] "How to Defeat Opposition to Moore". In Tomberlin, J., ed., [1999].

Stalnaker, R. [1970] "Pragmatics", reprinted in his [1999].

_____ [1998] "On the Representation of Context", reprinted in his [1999].

_____ [1999] *Context and Content. Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press. 1999.

Stroud, B. [1984] *The Significance of Philosophical Scepticism*. Oxford: Clarendon Press, 1984.

Tomberlin, J. [ed.] [1999] *Philosophical Perspectives*. Volume 13: Epistemology.

Travis, C. [1996] "Meaning's Role in Truth". *Mind*. Vol. 105, no. 419. 1996.

_____ [1997] "Pragmatics", in Hale & Wright, eds. [1997]

Unger, P. [1984] *Philosophical Relativity*. Oxford: Basil Blackwell, 1984.

Vogel, J. [1990] "Are There Counterexamples to the Closure Principle". In Ross, G. & Roth, M. eds., [1990].

_____ [1999] "The New Relevant Alternatives Theory". In Tomberlin, J., ed., [1999].

Wiggins, D. [1971] "On sentence-sense, word-sense, and difference of word-sense". In *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics and Psychology*. Jacobovits, L.A. & Steinberg, D.D. [eds]. Cambridge: Cambridge University Press. 1971.

Williams, M. [1991] *Unnatural Doubts. Epistemological Realism and the Basis of Scepticism*. Princeton: Princeton University Press, 1991.

_____ [2000] "Is Contextualism Statable?" In Villanueva E. & Sosa E. [eds.] *Philosophical Issues* 10, 2000. Published as a supplement to *Nous*.

Wright, C. [1985] "Facts and Certainty". Henriette Hertz Philosophical Lecture for the British Academy, in *Proceedings of the British Academy*, 71, 1985.

_____ [1991] "Scepticism and Dreaming: Imploding the Demon". *Mind*. Vol. 100, no. 397, 1991.

Yourgrau, P. [1983] "Knowledge and relevant Alternatives". *Synthese*. Vol. 55. 1983.

