

**Logic and The Limits of Explanation:
The Justification of Deduction, Carrollian Regress, Logical
Validity, and Deductive Inferential Knowledge**

By

Dean Chapman CHPDEA001

A dissertation submitted in fulfilment of the requirement for the award of the

Degree of Master of Arts

Department of Philosophy

University of Cape Town

April 2006

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Declaration

This work has not been previously submitted in whole, or in part, for the award of any degree. It is my own work. Each significant contribution to, and quotation in, this dissertation from the work, or works, of other people has been attributed, and has been cited and referenced.

Signed by candidate

Signed

18/4/2006

Date

Requests for reprints can be sent to:

Dean Chapman
The University of Zululand
P. Bag X1001
Internal Box no. 279
Kwadlangezwa
3886

Acknowledgements

I sincerely acknowledge the great value of the input that my supervisor, Bernhard Weiss, has given to this work. He has provided me with patient guidance and mentorship while allowing the free development of my thoughts, and I thank him for that.

Other members of the Department of Philosophy at the University of Cape Town, and associates thereof, also deserve my grateful acknowledgement, for stimulating conversations, contributions, support, and criticisms. They include Jeremy Wanderer, Elisa Galgut, Robert Segall, Antony Holiday, David Benatar, Jeremy Allcock, Annemie Gildenhuys, Leo Townsend, Monique Whitaker, and Neil. I presented parts of this essay in a less developed form at the Rhodes University Spring Colloquium, 2005, benefiting from the insightful comments and suggestions of audience members there, especially David Maartens, Francis Williamson, and Luke Buckland.

I would like to thank Carmen Perez, for inestimable support on Thursdays; my great friend Jang Jae Phil; my family, for much love and persistent encouragement throughout; and Sarah, for inspiration and assistance with the development of my ideas, but mostly for keeping me smiling, and, quite often, laughing.

Abstract

This essay engages with the problems of the justification of deduction, Carrollian regress, and deductive inferential knowledge. Also, it is considered whether Lewis Carroll's tale of what the tortoise said to Achilles can be interpreted as suggesting an argument against the possibility of logically *valid* argument. Such an argument is presented and shown to be unsound.

Any justification of one of our basic rules of deductive inference, such as *modus ponens*, will inevitably make use of the very rule it means to justify. It will be a 'rule-circular' argument and invite charges that it begs the question and 'keeps bad company'. Following Paul Boghossian, the contention in this essay is that a thinker need not know that the *rule* according to which a given inference proceeds is sound in order to be entitled to carry out the inference. Thus, a rule-circular argument for the soundness of *modus ponens* does not beg the question. Also, by a conceptual role semantics which takes as its starting point that of Boghossian, and with insights gained from Robert Brandom's inferentialism, it is argued that a thinker who carries out an inference which is meaning-constituting of some concept for her is entitled to that act of inference, in part because she is *epistemically blameless* in it.

One of the ways to counter a Cartesian sceptic is to maintain that some of our beliefs are beliefs we are entitled to have no doubt about. To make that claim good, it is argued, one must hold two things: first, that some of our beliefs are such that we have *conclusive* evidence for them, evidence which *guarantees* their truth; and second, that for some of *these* beliefs, we *know* that we have conclusive evidence for them - there are infallibilist and internalist constraints on the possibility of us having knowledge that is *certain*. *Pace* Boghossian, the contention here is that anyone who carries out an inference which is meaning-constituting of some concept for her, in fact *knows* that inference to be valid.

It is argued that knowing of the validity of an inference is sufficient for being entitled to carry it out so that one can thereby come to have what is here called 'certain knowledge' of the truth of its conclusion, given the satisfaction of other broadly applicable constraints. Thus, it is held that a thinker who carries out a meaning-constituting inference can thereby come to have certain knowledge of the truth of its conclusion.

The central undertaking of this essay is therefore to face up to the problem of Carrollian regress, insofar as the main difficulty it raises for us has to do with the possibility of deductive inferential *knowledge*. By a conceptual role semantics it is argued that the error committed in allowing ourselves to be led into Carrollian regress is that of distinguishing too sharply between *practical* and *propositional* knowledge. It is a proper requirement that a reasoner must know that an inference is valid in order to be entitled to carry it out; but we need not think further that such knowledge must be deployed with the proposition of whose truth the thinker knows as a *premise* of the argument for her conclusion. A belief can serve as the basis of a subject's carrying out of an *act* - such as an assertion, or an *inference* - and an inferer's belief that an inference is valid being a case of knowledge is sufficient for her being entitled to that *act* of inference. The account is satisfyingly an account of both epistemic inferential *entitlement* as well as the *rationality* of inference.

Table of Contents

Chapter 1		1
Introduction		
Chapter 2		
The Justification of Deduction, Carrollian Regress, and Logical Validity		
I	The Justification of Deduction	6
II	Rule-Circularity	8
	II-i Bad Company	10
	II-ii Begging the Question	12
III	What the Tortoise said to Achilles	14
IV	What of What the Tortoise Said to Achilles?	16
	IV-i We Want What We Cannot Have	20
	IV-ii An Invalid Argument Against Valid Arguments	23
	IV-iii A Challenging Argument Challenging Deductive Inferential Knowledge	28
Chapter 3		
Descartes' Doubts, Infallibilism, and Internalism		
I	Introduction	32
II	Descartes' Doubts as Motivation for a Special Interest in <i>Deductive Inferential Knowledge</i>	34
III	Infallibilism and Internalism	43
	III-i Infallibilism	43
	III-ii Internalism	46
IV	Know it or Not	51

Chapter 4
The Sceptic Up Close

I	Introduction	53
II	The Sceptic's Argument, the Natural Reply, and the Problem of Regress	55
	II-i The Sceptic's Argument and the Most Natural Reply	56
	II-ii The Problem of Regress	60
III	The Form of the Requirement - Explicit vs. Implicit Knowledge	64
	Section III Preliminaries	64
	Section III Overview	67
	III-i Explicit Knowledge; Explicit Premise	68
	III-ii Explicitly Accepting that <i>P</i> entails <i>C</i> in the Course of Inferring	70
	III-iii Being Able to Know <i>Explicitly</i> that <i>P</i> entails <i>C</i>	73
	III-iv Implicit Knowledge	77
IV	Telling Myths to get into Regress?	78

Chapter 5
No Need to Know?

I	Introduction	86
II	My Beliefs are Not the Matter?	89
III	All that My Knowledge Entails	97
IV	Modus Ponens Premises; Deadlock; Progress	101

Chapter 6
Putting Theory into Practice

I	Introduction	108
II	Pointing the Way Forward	110
III	Boghossian's Conceptual Role Semantics	113
	III-i Simple Inferential Internalism	116
	III-ii Rational Insight	117
	III-iii Boghossian's Way Forward	119
	III-iv Putting Theory into Practice	122
IV	Brandom's Semantic Inferentialism and Logical Expressivism	127
V	Knowledge Needed, Though not as a Premise	134
	V-i Knowledge Needed	134
	V-ii Reason to Infer	138
VI	Entitlement to <i>Act</i>	142

Chapter 7
Objections, Implications, and Conclusion

I	Introduction	150
II	<i>How do You Know?</i>	150
	II-i Reflectively Appreciable Warrant	158
	II-ii Is an act of inference being in accordance with a meaning-constituting rule sufficient for its epistemic blamelessness?	162
	II-iii The set of inferences which are meaning-constituting of a concept <i>for a thinker</i>	165
	II-iv Epistemic Blamelessness is indeed <i>not</i> sufficient for Warrant	174
III	The Justification of Deduction	179
	III-i Begging the Question	181
	III-ii Bad Company	184
IV	Conclusion	187
	Reference List	190

Chapter 1

Introduction

The problem of the justification of deduction is a large and difficult philosophical problem, especially important now since the contemporary philosopher of logic is confronted by a whole range of systems of logic, many of which are quite overtly rivals of classical logic.¹ The question is whether it is possible to show that some basic rule of deductive inference is sound - necessarily truth-preserving in all cases²; and, if so, how? If we want to justify our *employment* of some such basic rule of logic, such as modus ponens, we will *need* to justify our belief that it is sound. But for any of our fundamental logical beliefs - the problem is - logic itself 'will inevitably be involved in any account of how we might be justified in believing it.'³ Thus, Hilary Putnam writes:

My own guess is that the truths of logic we are speaking of are *so* basic that the notion of *explanation* collapses when we try to 'explain' why they are true. I do not mean that there is something 'unexplainable' here; there is simply no room for an explanation of what is presupposed by every explanatory activity ...⁴

In chapter 2 I shall introduce the problem of the justification of deduction more fully, explaining that a justification of a basic rule of logic will inevitably be 'rule-circular', and therefore seemingly be open to the charges of begging the question and keeping 'bad company', as Boghossian puts it.⁵ This will lead us into a discussion of Lewis Carroll's famous tale of 'What the tortoise said to Achilles'⁶, and the question of what difficulties it raises for us. I consider an argument against the possibility of logically valid argument which might be gleaned from what the tortoise says to Achilles, going on to show that this argument against valid argument is unsound. The most challenging problem the story

¹ See Haack (1974). As she explains, there is some debate about whether 'genuine rivalry is possible'. She contends that it is. (1974, p. 23)

² Haack (1976, p.114)

³ Boghossian (1999, p. 229)

⁴ Putnam (1979, p. 110)

⁵ Boghossian (1999)

⁶ Carroll (1895)

brings to our attention, however, has to do not with logical *validity*⁷, but, as I see it, with the attainability of deductive inferential *knowledge*, as I'll explain shortly.

The main purpose of chapter 3 is to provide some motivation for a special interest in *deductive* argument. I argue that one can try to counter a general Cartesian sceptic in one of at least two ways: one can deny his claim that one ought to have doubt about each of one's beliefs; or one can grant it and maintain that some of the things one believes, one in fact *knows* - notwithstanding the fact that one ought to have doubt about them. Taking the second option would entail denying the possibility of us having any knowledge that's *certain* knowledge. Rather, I wish to pursue the first line of approach and to argue that certain knowledge is not an impossibility.

But one can only have certain knowledge of the truth of some proposition, *C*, I shall argue, on the basis only of knowledge that *P*, if: firstly, *P* entails *C* (thus, my interest in *deductive* argument); and, secondly, it is not the case that one ought to doubt that *P* entails *C*. It is plausible, in fact, that what we require is that one must *know* that *P* entails *C* if one is to be able to have certain knowledge that *C* on the basis of one's knowledge that *P*, as my Carroll's tortoise maintains. And so, we must face up to the problem of Carrollian regress, which, in its basic form, runs as follows. If I need to know that *P* entails *C* to be able to have certain inferential knowledge of the truth of *C*, then surely I need to have *P* entails *C* as a premise of whatever argument I give for *C*. But then I must still know that *this new argument for C is valid*, and have *this* proposition as a premise of whatever argument I give for *C*. And so on, *ad infinitum*. So it would appear that I could never have the sort of inferential entitlement we are after.

Chapter 4 takes us to the heart of the sceptic's (or the tortoise's) argument against the possibility of me having the sought for entitlement to carry out an inference. I argue that one's most natural response to a sceptic who claims that one doesn't know that a given inference one carries out is valid, is to say that one most certainly *does* know that it is valid, as long as one believes that one has an entitlement to carry out the given inference.

⁷ See Stroud (1979)

Thus, we implicitly accept the requirement that one must know that an inference is valid in order to be entitled to carry it out. The way to avoid Carrollian regress, then, I shall urge, is to maintain that though I need to know that *P* entails *C* in order to be entitled to infer *C* from *P*, I do not need to *supplement* my argument for *C* by adding *P* entails *C* as one of its premises. If we could make this claim good, then we could avoid Carrollian regress. In the remainder of the chapter, by noting that in practice we attribute knowledge of the validity of an inference always together with an entitlement to carry out the inference, but that we do so even when a subject does not and cannot know *explicitly* that her argument is valid, I shall argue that the requirement must be that a thinker must know, at least *implicitly*, that an inference she carries out is valid in order to be entitled to that act of inference. To conclude the chapter I present an argument from Gilbert Ryle⁸ against the propriety of the requirement in this form.

Chapter 5 takes as its starting point the observation that an account of inference must have two aspects. It must be an account of what it is to be *epistemically entitled* to an act of inference, and also an account allowing us to see acts of inference as *rational* acts. In this chapter I shall propose the view, which will not ultimately be fully endorsed, that insofar as we are concerned with what it is to be *epistemically entitled* to an inference, our account need make no reference to the *belief*, nor the *knowledge*, that an inference is valid which a reasoner has (if she has it) when she carries out that inference. The thought here is that what is required for one to be *entitled* to an act of inference, is only that one has an *entitlement* to believe that the inference is valid. Then, since we need make no reference to a belief or item of propositional knowledge which is had (that the inference is valid), we are not faced with the difficulty of explaining how said propositional knowledge is *deployed*.

However, as I shall explain, because the account of chapter 5 prescind from talking of the belief that *P* entails *C*, which I *do* maintain a reasoner has when she infers deductively from *P* to *C*, it is lacking in two respects. First, it leaves a reasoner's acts of inference as unexplained; it is not an account of the rationality of inference. Second, it is

⁸ Ryle (1946)

insufficiently internalist, leaving us with no understanding of how an entitlement to believe that *P* entails *C* is brought to bear on a warranted act of inference. With cognisance of the insights gained and lessons learned in chapter 5, then, in chapter 6 I shall develop my account of inference further, justifying the attribution of the (sometimes implicit) belief that *P* entails *C* to anyone who carries out a deductive inference from *P* to *C*. Then, by a conceptual role semantics along the lines of that of Paul Boghossian, and focussing on the question of how we are entitled to carry out inferences of the form of modus ponens, I shall argue that anyone who has the concept of the conditional is disposed to carry out inferences of that form. By the employment of Boghossian's notion of *epistemic blamelessness*, I shall argue that a thinker who carries out an inference by modus ponens is *epistemically blameless* in her associated *belief* that the inference is valid. This is because the inference is one which is meaning-constituting of the concept of the conditional. Thus, anyone who carries out an inference by modus ponens *with entitlement* has the belief that the inference is valid, she is blameless in having that belief, and the inference *is* valid (as we showed that it must be if she is to be entitled to the inference, in chapter 3). If this, together with the fact that her inference is a *meaning-constituting* inference, is sufficient for her in fact *knowing* that her inference is valid, as I shall argue in chapter 7, then anyone who carries out an inference by modus ponens in fact *knows* that inference to be valid.

In chapter 7 I shall return to explain why, by a conceptual role semantics, the above subject's modus ponens inference carried out *is* valid, as are all modus ponens inferences. Since inferences of the form of modus ponens are *constituting* of the meaning of the conditional, the thought is, it is impossible that the premise(s) of a modus ponens inference should be true while its conclusion is false. Before then, however, and *pace* Boghossian, I shall argue that we do not involve ourselves in Carrollian regress by remaining faithful to the idea that a thinker must know a deductive inference to be valid if she is to be entitled to carry it out. We are only led into regress if we neglect the fact that there is an intimate connection which exists between *propositional* and *practical* knowledge. As Boghossian himself allows us to appreciate, we do not attribute a given belief to a subject unless she *acts* – in particular, *infers* – in ways in which she is required

belief to a subject unless she *acts* – in particular, *infers* – in ways in which she is required to act in order to have an understanding of the proposition (purportedly) believed. But, similarly, I shall argue, the possession of some item of propositional knowledge goes hand in hand with an entitlement to *act* in certain ways – in particular, to carry out inferences. To carry this proposal through, I contend that anyone who carries out a deductive inference from *P* to *C* has the proposition *P* entails *C* as her reason for carrying out that *act*. And if she in fact *knows* that *P* entails *C*, then she is *entitled* to that act of inference.

The purpose of chapter 7 will be to fill in some of the details of my account of why any thinker who carries out an inference by modus ponens, say, in fact knows that inference to be valid. We'll revisit some of the objections already encountered along the way, and also meet with some previously unseen objections. I'll indicate how I think we need to proceed in response, after which we'll return to consider the implications of my account for the problem of the justification of deduction.

Chapter 2

The Justification of Deduction, Carrollian Regress, and Logical Validity

I The Justification of Deduction

According to the standard semantic definition, a deductively valid argument is an argument whose premise(s) and conclusion are such that the truth of the premise(s) *guarantees* the truth of the conclusion. That is, an argument is deductively valid if and only if it is impossible that its premise(s) should be true and its conclusion false.¹

Now consider the rule of inference known as modus ponens (here standardly abbreviated as ‘MP’) - the rule commonly taken to be the most basic rule of deductive inference - which can be formulated thus:

For any propositions P and C ,

from	If P then C , P
infer	C .

The problem of the justification of deduction, as conceived by Haack³ is to show that some basic rule of deductive inference such as MP is indeed truth-preserving. I shall focus on modus ponens in this essay, and must begin by asking whether a justification of modus ponens is even *possible*.⁴ Is it possible for us to show that MP is truth-preserving,

¹ Haack (1976, p.113)

² For now I follow Haack (1976) in treating the conditional involved in modus ponens as the *material* conditional of classical logic. Haack notes that the problem of the justification of deduction can be raised for any understanding of the conditional.

³ Haack (1976, p.114)

⁴ Boghossian, in (1999, p. 229), asks only whether ‘it [is] so much as *possible* for us to be justified in supposing that MP is a valid rule of inference, necessarily truth-preserving in all its applications?’ He is not there concerned with whether we are *actually* justified in supposing this. Also, and importantly, we’ll see later that he distinguishes between being justified in the *belief* that MP is valid and being entitled to *carry out* inferences in accordance with that rule.

that if an argument is of the form of modus ponens then it is impossible that its premise(s) should be true and its conclusion false? Indeed, there are considerations which quite powerfully suggest otherwise.

Since our belief that MP *is* truth-preserving is usually taken to be one for which we have maximal security, so that we think that MP, *at least*, will support us in making valid and justified inferences even when, or if, all else fails⁵, it might be startling to realise that we appear to be on particularly shaky ground when faced with the task of MP's justification. To realise this we need only appreciate that for any of our fundamental logical beliefs, as Boghossian puts it, 'logic will inevitably be involved in any account of how we might be justified in believing it.' How then could it be 'possible for us to be justified in our fundamental logical beliefs?'⁶

To elaborate, if a justification of MP is called for it is presumably an *argument* for the truth-preservingness of the rule we are after. That is, what we need is an *inferential* justification of MP.⁷ But any argument aiming to show that MP is truth-preserving must, one might naturally suppose, itself proceed according to some deductive⁸ rule or rules of inference. It must then either use MP, or some other rule of inference. But if the argument in support of MP itself uses that very rule, MP, it invites a charge of what Boghossian calls 'rule-circularity'— to be discussed further below. And it would seem that the situation can be made no better by the argument invoking some *other* rule of inference, for in that case we have surely not shown MP to be truth-preserving until we have shown that the *this other* rule is truth-preserving. And if we just continue to call upon another distinct rule each time it's pointed out that we haven't completed our justification, we'll clearly never complete it. The alternative is to come back round to using MP (or some

⁵ Consider the infamous problem of *induction* and how it is often highlighted by comparing inductive to deductive argument, and finding the former wanting in comparison. See Haack (1976) for argument that the problems of the justification of deduction and the justification of *induction* are in fact analogous.

⁶ Boghossian (1999, p. 229)

⁷ See Boghossian's (1999, p. 230) brief consideration of the possibility of us having a non-inferential justification. I shall return to discuss inferential versus non-inferential justification later in this essay.

⁸ See Haack (1976) for why an *inductive* justification of MP would be too weak. Mill's Inductivism, according to which logical truths are just inductive generalizations, fails for just this reason. Analogously, Haack argues, a *deductive* justification of an *inductive* rule of inference would be too *strong*.

other rule we've previously employed), in which case we'll have just as objectionable an instance of rule-circularity as if we'd used MP to start with.

II Rule-Circularity

One might try to justify MP by arguing as follows:

Assume that the premises, *if P then C* and *P*, of an MP argument are both true. Now, by the truth-table for the conditional, if the two propositions *if P then C* and *P* are both true, then *C* is true. So *C* must be true.⁹

This argument has taken at least one step which depends on the very rule, MP, which it is meant to justify. At least part of the argument is given by

<i>If P then C</i> and <i>P</i> are both true	(Assumed premise)
<u>If <i>if P then C</i> and <i>P</i> are both true, then <i>C</i> is true</u>	(By the truth-table)
<i>C</i> is true	

And *this* argument itself employs modus ponens. So we see that our argument given in an attempt to justify MP is *rule-circular*.

A rule-circular argument, as opposed to a grossly circular one, does not have or contain in any way its conclusion as, or in, any of its premises. A grossly circular argument, on the other hand, does just that.¹⁰ But, now what's so bad about rule-circular argument? Some have argued that whereas gross circularity is certainly in all cases unacceptable, the same is not true of rule-circularity.¹¹ This is indeed the line that Boghossian takes since

⁹ See Boghossian (1999, p.231); also Haack (1976) and Weiss (2002)

¹⁰ I borrow the terms 'rule-circular' and 'grossly circular' from Boghossian; see (1999, p. 245).

¹¹ For an example, see Black (1954) - cited by Haack (1976, p. 114). Dummett (1973) distinguishes between 'explanatory' and 'suasive' arguments and contends that a rule-circular argument is objectionable only if it is meant to be *suasive*: 'The philosopher does not seriously doubt the validity of the law and is therefore prepared to accept an argument in accordance with it. He does not seek to be persuaded of the conclusion; what he is seeking is an *explanation* of its being true.' I follow Boghossian (2002, p. 44-45) in

he agrees with Quine and Harman¹² that if we accept ‘the ban on the use of a logical principle in reconstructing our a priori warrant for that very principle, we would have to conclude that there can be no such reconstruction’. ‘[W]e cannot accept the claim that we have no warrant whatsoever for the core logic principles,’ and ‘[w]e cannot conceive what such a warrant could consist in ... if not in some sort of inference using those very core logical principles.’^{13, 14}

Boghossian points to two distinct reasons for *gross* circularity being unacceptable, in order to compare *rule*-circularity.¹⁵ First, a grossly circular argument, as we say, begs the question: it presupposes its conclusion by including it as (or as ‘part of’) one of its premises. And quite apart from any other considerations, Boghossian says, that just ‘seems wrong’. Second, a grossly circular argument, as Boghossian puts it, ‘keeps bad company’. If grossly circular arguments were acceptable, then *any* conclusion, including false and *patently* unjustifiable ones (thus, ‘bad company’), *could* be justified by an argument which simply included the desired conclusion as one of its premises.

Does rule-circularity, then, suffer from these same two problems? One might think, with Boghossian, that it does not. Let us first take a look at the question of whether or not rule-circular arguments keep bad company, and then at whether or not they are question-begging.

being interested in whether it is possible to have a *suasive*, or *justificatory* argument for the soundness of MP. See also Haack (1976) and Haack (1982).

¹² Harman cited by Boghossian (1999, p.232) as a reply entitled “Analyticity regained” to Boghossian’s article “Analyticity reconsidered”, see Boghossian (1996) and Harman (1996).

¹³ Boghossian (1999, p.232-234) also rejects the idea of proposing an *empirical* justification of logic in the face of the problems at hand, pointing out that *any* argument offered in an attempt to justify our ‘core logical principles’ would have to make use of those very principles. I shall come back at a later point to discuss in some depth Boghossian’s defence of (some) rule-circular arguments.

¹⁴ Boghossian (1999, p. 253)

¹⁵ Boghossian (1999, p. 245)

II-i Bad Company

Prima facie, it appears that rule-circular arguments keep bad company. Consider Arthur Prior's connective 'tonk'¹⁶ which has the following introduction and elimination rules.

Tonk-introduction:

For any propositions P and C ,

from	P
infer	$P \text{ tonk } C$.

Tonk-elimination:

For any propositions P and C ,

from	$P \text{ tonk } C$
infer	C .

Evidently, these rules license the inference of any proposition whatsoever from any other proposition, so they must be unjustifiable. But if rule-circular arguments were acceptable then we could justify tonk's introduction rule as follows:¹⁷

¹⁶ Prior (1967)

¹⁷ The argument to follow is given by Boghossian (1999, p. 247), drawn from Crispin Wright's commentary on that paper at the Stirling Conference on Naturalism.

1. 'P tonk C' is true iff 'P' is true tonk 'C' is true	Meaning Postulate
2. P	Assumption
3. 'P' is true	2, T-scheme
4. 'P' is true tonk 'C' is true	2, tonk-introduction
5. 'P tonk C' is true	4, 1, biconditional-elimination
6. P tonk C	5, T-scheme
7. If P, then P tonk C	6, logic

This appears to allow us a justification that *P tonk C* follows from *P* dependent on the meaning postulate in line 1.

Or, consider Haack's rule of modus morons:¹⁸

For any propositions *P* and *C*,

from	<i>If P then C</i> is true, <i>C</i> is true
infer	<i>P</i> is true.

(Someone I know likes to call it *modus bogus*.) Haack comes up with this rule in order to argue that rule-circular arguments keep bad company; that if rule-circular arguments were acceptable then we could justify patently invalid and unjustifiable rules of inference.¹⁹ We could, for instance, justify modus morons, as follows:

Assume that *C* is true and that *If P then C* is true. Given that the conditional is true we know, by the meaning stipulation for the conditional, that if *P* is true then *C* is true. So, by modus morons, *P* is true.²⁰

¹⁸ Haack (1982)

¹⁹ See Weiss (2002, p. 128-129)

²⁰ Weiss (2002, p. 129)

‘Clearly,’ Boghossian notes, ‘we cannot simply assert that rule-circular justifications are acceptable and leave it at that. The question is whether there is some intuitively plausible constraint that they should be made to satisfy that will repel the bad company they would otherwise keep while leaving in place the justificatory arguments that interest us.’²¹ I shall return in chapters 6 and 7 to the constraint on rule-circular arguments that Boghossian advocates.

II-ii Begging the Question

On the issue of begging the question, an argument at least does not seem to rely on ‘its implicated rule of inference *in the same way* ... [as] it relies on its premises.’²² (My emphasis.) One might even hold, as I read Dilman as holding²³, that an argument does not depend on any general proposition or rule of inference at all. One might agree with such a view as Wittgenstein expresses in the *Tractatus*:

The nature of the inference [from one proposition to another] can be gathered only from the two propositions. They themselves are the only possible justification of the inference. ‘Laws of inference’, which are supposed to justify inferences, as in the works of Frege and Russell, have no sense, and would be superfluous.²⁴

But the problem is that if the rule employed in an argument is not truth-preserving, then the truth of the premise(s) of that argument does not guarantee the truth of its conclusion. So then, to put it rhetorically, how can an argument which employs MP establish for us the very fact that MP is truth-preserving? If the conclusion of the argument is that MP is in fact truth-preserving, then, since that argument will be invalid, and therefore unacceptable, as long as its very conclusion is false, it seems in a clearly objectionable sense to be question-begging.

²¹ Boghossian (1999, p. 248)

²² Boghossian (1999, p. 245)

²³ Dilman (1973, p. 113)

²⁴ Wittgenstein (1922, 5.132)

As Boghossian explains²⁵, if it is true - as seems to be our intuition given what I have said above (though Boghossian denies that it's true, as we'll see later) - that an entitlement to *use* MP depends on an entitlement to *believe* that MP is truth-preserving, we must already be entitled to believe that MP is truth-preserving in order to be entitled to employ our above rule-circular argument for that rule. So that argument would not enable us to have any entitlement to believe its conclusion that we did not already have.

An easy way to appreciate these difficulties is to imagine someone who does not accept MP as truth-preserving; she either doubts it or in fact believes that it is *not* truth-preserving.²⁶ If we offer to this non-believer the rule-circular argument above which, as we recall, involved this argument:

If P then C and P are both true
If if P then C and P are both true, then C is true
C is true

then she would surely be quite right to object that we have begged her question²⁷. Anyone who doubts that MP is truth-preserving or believes that it is not, and who is vigilant enough, will of course not be convinced of the truth of some conclusion by an argument of the very form of modus ponens. But the starting point of our subject is precisely one of doubt with respect to the truth-preservingness of MP, and the argument is supposed to unburden her of this very doubt. In requiring that she infer from the premises to the conclusion of an MP argument in order to see that arguments of that very form are indeed truth-preserving, her question seems clearly to have been begged.

Indeed, it may seem that our imaginary doubter might well even accept the second premise (after also accepting the first) of the above argument while still refusing to make

²⁵ Boghossian (1999, p. 246)

²⁶ The fact that we might be faced with such a non-believer would not be problematic if we were only after an 'explanatory' argument for the soundness of MP (see note 11 above).

²⁷ See Boghossian (1999, p. 252)

the inference to the conclusion. “Yes, I see that according to the truth-table for the conditional, that **if** [*if P then C*] and *P* are true, **then** *C* is true. But I still don’t believe that *C* is true,” she might say. And, as we’ve seen, she might also add: “Your argument is of the form of MP. That’s why *you* think it’s valid, think that it establishes its conclusion and think that I ought to be convinced by it of its conclusion. But *I* don’t think that MP arguments are truth-preserving so I don’t think that the truth of the premises of your argument guarantees the truth of its conclusion. I therefore won’t infer the conclusion from the premises.”

But this is something very similar to what, according to Lewis Carroll, the tortoise said to Achilles.²⁸

III What the Tortoise said to Achilles

I present here a modified version of Carroll’s story.

For some particular propositions, *P* and *C*, Achilles tries to convince the tortoise, who accepts both the truth of *P* and of *if P then C*, to also accept that of *C*. Indeed, Achilles maintains that the tortoise *must* also accept the truth of *C*. To this end Achilles writes the MP argument

$$\begin{array}{l} P \\ \underline{\textit{If P then C}} \\ C, \end{array}$$

in his notebook. The tortoise, however, refuses to infer the conclusion.

In a renewed effort to induce the tortoise into accepting *C*, Achilles then says: “You *must* accept *C*. Look: If both *P* and *If P then C* are true, then *C* is true. But we have both the truth of *P* and of *If P then C*, so *C* is true.” He tries to convince the tortoise of the truth of

²⁸ Carroll (1895)

C by getting him to accept a conditional statement linking the two premises of the original argument, as antecedent of the conditional, with its conclusion, as consequent. (Following Thomson²⁹ I shall call this conditional statement the hypothetical associated with the original argument.) He then, as we'll see, asks the tortoise to infer C from the two original premises *together with* this new conditional statement. In doing so Achilles presents just the sort of rule-circular argument that we considered above. And in order to convince the tortoise of the relevant conditional, Achilles might also appeal to the meaning of the propositions involved, perhaps with reference to the truth-table of the conditional.

But the tortoise needs no convincing of the truth of that conditional statement, or, as we'll see, of any other proposition (apart from C itself), for he thinks that he can accept and believe anything Achilles asks him to without being 'compelled' to accept and believe that C. Without asking for any reason to believe it, the tortoise happily accepts what Achilles has asked him to, proposition (iii) below,

- (i) P,
- (ii) If P then C,
- (iii) If (P and if P then C) then C,

so that he has now accepted all three of these. Indeed, he accepts (iii) as a *logical* truth, and adds:

Whatever *Logic* is good enough to tell me is worth *writing down*... So enter it in your book please...Until I've granted [*this conditional*], of course I needn't grant [the conclusion]. So it's quite a *necessary* step, you see?³⁰

²⁹ Thomson, J. (1960)

³⁰ Carroll (1895, p. 693)

Achilles concedes and enters the new proposition into his notebook which now contains the above three propositions together with the conclusion, *C*, which he hopes the tortoise will now accept.

Predictably, however, the tortoise remains unwilling to infer *C* from the enlarged pool of propositions ((i), (ii), and (iii)), each of which he has now accepted. He persuades Achilles that until he, the tortoise, has granted yet another conditional, (iv) below, he needn't grant the conclusion.

- (i) *P*,
- (ii) *If P then C*,
- (iii) *If (P and if P then C) then C*,
- (iv) *If [(i) and (ii) and (iii)] then C*.

Since it is 'quite a necessary step' that he grants (iv), says the tortoise, this new proposition is then entered into the notebook of Achilles, while the tortoise remains unwilling to accept the conclusion. And so on, *ad infinitum*.

IV What of What the Tortoise Said to Achilles?

What is the moral to be drawn from this story? Does it bring to our attention a problem to do with logical *validity*, or to do with deductive inferential *knowledge*?

According to Stroud³¹, J. F. Thomson³² reads Carroll's story as a failed attempt to raise a genuine problem to do with logical *validity*. Thomson's idea is that Achilles is misled into thinking that his original argument has a suppressed premise which is 'quite necessary' for its *validity*. But herein, on Thomson's view, lies the mistake. The fact that the tortoise insists, even if rightfully, that Achilles must convince him of yet another proposition before he is compelled to *accept* the conclusion, does not show that the

³¹ Stroud (1979)

³² Thomson (1960)

original argument, which *lacked* that proposition as a premise, was *invalid*. Stroud agrees with Thomson on this point, saying with reference to the original argument that ‘the conjunction of [its premises] does imply [the conclusion], whatever the tortoise happens to think about it’.^{33, 34}

So, according to Stroud,³⁵ Thomson’s Achilles is after an argument which establishes its conclusion in the purely logical sense of being logically valid whilst having true premises. And Thomson argues that Achilles’ strengthening of the original argument, his addition of the conditional statement linking its premise(s) with its conclusion (its addition as a *premise*), is either futile or redundant for this purpose. Achilles should never have set off on the regress in the first place. Consider:³⁶

For starters, since an argument with a false premise cannot logically establish the truth of its conclusion, if we start off with an argument with a false premise, strengthening it will clearly be of no use (the strengthened argument will still have a false premise). So let us assume that Achilles’ original argument from *P* together with *if P then C*, to *C*, has only true premises. Now, this argument is either logically valid, or invalid.

If it is valid then the premise, (iii), added in strengthening it will indeed be true: if the conditional involved is taken to be the *material* conditional then, since the original argument is valid, either its antecedent is false or both its antecedent and consequent are true; if ‘if, then’ is read as ‘entails’ then (iii) will *clearly* be true. Now, the strengthened argument will also be valid, since we cannot turn a valid argument into an invalid one by adding premises to it. But since Achilles already *had* a logically valid argument with true premises, the adding of the new premise was not required and that premise in the new argument is redundant.

³³ Stroud (1979, p. 180)

³⁴ I shall return shortly to the question of whether an argument against the possibility of valid argument might be made out in Carroll’s story, if we generously read the tortoise as claiming not that he must *accept* something else before he is compelled to accept the conclusion, but that something else apart from the premises must be *true*, before the truth of the conclusion is *guaranteed*.

³⁵ Stroud (1979, p.180-181)

³⁶ What follows is my formulation of Thomson’s (1960, p. 97-98) argument.

On the other hand, if the original argument was *not* valid, then an attempt to *make* it valid by strengthening it is futile. For, claims Thomson, the added premise of the strengthened argument is ‘unacceptable’ in some way and thus of no help in establishing the conclusion.

But *in what way* is that added premise unacceptable in this last case? Understanding for now the conditional involved as the material conditional, as both Stroud and Thomson do, we can appreciate easily that the conditional statement added as a premise would at least be unacceptable *to the tortoise*. For the tortoise (supposedly) believes the antecedent of that conditional, since he believes both the premises of the original argument, while he doubts its consequent, the conclusion of the argument. We *might* think that the tortoise ought to accept the relevant conditional statement when the original argument is valid, since in that case he *ought* (as Achilles is at pains to show) to accept its consequent when he has accepted its antecedent, which he has, and of course a conditional statement is true when both its antecedent and consequent are true. But we surely cannot think that the tortoise ought to accept (iii) when the original argument is *invalid*, given that at the outset he believes its antecedent but doubts its consequent. He ought rather to find (iii) *unacceptable* in such a case.

But, as Stroud points out³⁷, Thomson cannot be concerned with the fact that this proposition would in such a case be unacceptable *to the tortoise* - *epistemically* unacceptable - if he is only interested in whether or not the strengthened argument *logically* establishes its conclusion. An argument logically establishes its conclusion as long as its premises are *true* and it is logically valid, and any strengthened argument is indeed logically valid, for it is of the form of modus ponens. But then the only way for a premise of a strengthened argument to be ‘unacceptable’, so as to disqualify that argument from logically establishing its conclusion, is for the premise to be simply *false*. So Thomson must mean that the premise added in the strengthening of any invalid argument is *false*.

³⁷ Stroud (1979, p. 182)

But if that's Thomson's idea, Stroud goes on, then he is simply mistaken. An attempt, by strengthening, to turn an invalid argument into one which logically establishes its conclusion is not always futile – the strengthened argument *will* be valid and *may* have all of its premises including the one added *true*. Quite simply, if the original argument had all of its premises *and* its conclusion true despite being invalid, then the added premise, which, recall, is the (material-) conditional statement linking those (true) premises to that (true) conclusion, would of course be *true*, not false. But then the strengthened argument would also have *all* of its premises (and its conclusion) true and would be, as we have said, valid. It would therefore be one which *did* establish its conclusion in the purely logical sense.

So Stroud claims that the story in fact 'shows something interesting about *inference* and *belief*,'³⁸ (my emphasis) and its intention is not, as Thomson claims, 'plainly enough, to raise a difficulty about the idea of valid arguments'.³⁹ '[T]he lesson can be fully appreciated only if those puzzling notions [of inference and belief] are clearly distinguished from the better-understood notions of *logical consequence* and *truth*.'⁴⁰ (My emphasis.)

This view is lent support, as Stroud finds, by a closer reading of Carroll's story. The tortoise wishes to understand why he *must accept* the conclusion of the original argument when he has accepted its premises, not why that argument is valid. Indeed, he might well even accept that the argument *is* valid or that its premises *entail* its conclusion, but still remain unwilling to *infer* the conclusion. For logic, after all, 'simply states what is the case, or that something implies something else; it does not state that people must ... accept certain things if they already accept certain other things.'⁴¹

³⁸ Stroud (1979, p. 180)

³⁹ Thomson (1960, p. 95)

⁴⁰ Stroud (1979, p. 180)

⁴¹ Stroud (1979, p. 182)

When the tortoise talks above of a ‘quite necessary step’, he is not saying that the truth of the conditional statement at hand is quite necessary for the *validity* of the argument, he is saying that it is quite a necessary step that he *accepts* that conditional statement in order to be compelled to *accept* the conclusion. “Until I’ve granted [*this conditional*], of course I needn’t grant [the conclusion]. So it’s quite a *necessary* step, you see?”

In section IV-iii below I shall set out, as I understand it, the real problem raised by Carroll’s story, and this problem will be the main focus of the chapters to come. In chapter 6 I shall return to a consideration of some of Stroud’s further reflections. For now, however, I wish to stay with Thomson for a moment to explain another of his arguments - one which is quite illuminating.

IV-i We Want What We Cannot Have

Before I proceed to that argument, note that if we understand the conditional involved in the premise added in the strengthening of an argument, not as the *material* conditional, but as we understand ‘entails’, or ‘*necessarily*, if, then’, then Thomson’s argument that strengthening is either futile or redundant for the purposes of finding an argument which *logically establishes* its conclusion, *does* go through. For, again, if an argument is already valid with true premises then another premise added to it is redundant in the new argument – the original was already an argument which logically established its conclusion. Then, if an argument, from *P* to *C* say, is *invalid*, then *P entails C* is plainly *false*, so adding it as a premise will be futile as an effort to turn the original argument into one which logically establishes its conclusion. For the remainder of this, essay, therefore, I shall use the conditional of strict implication, and write it as ‘entails’. I do so also because I think it more natural to say that I must first accept that *P entails C* (as opposed to that *if P then C* ⁴²) before it is the case that I *must* accept that *C* when I have accepted

⁴² Material conditional.

that P ⁴³; to say that someone must know that P entails C in order to be entitled to infer *deductively* from P to C ; etc.

To proceed as indicated, for simplicity, imagine that P and C are two propositions such that the argument from the former to the latter is obviously valid, that the tortoise has accepted that P , and that Achilles tries to convince him to accept that C by telling him that P entails C . Thomson's insightful argument⁴⁴ is nicely introduced by Stroud in his consideration of it:

Thomson rightly emphasizes that to point out that someone has good or conclusive reasons for believing something is not to state another reason he has to believe it. The remark is a comment on the original premises; it says something about them, but it does not give a part of the original argument. So when Achilles is led to remark at each point that if everything the tortoise has accepted so far is true then [the conclusion] must be true he should not be seen as offering the tortoise yet another reason to believe [the conclusion]. If he were doing that, he would never find enough reasons.⁴⁵

We might think that Achilles rightfully asks the tortoise to accept that P entails C because it is when, and only when, the hypothetical associated with his argument is itself one of the *premises* of that argument that he *must* infer according to the argument. But this thought is confused and unacceptable. Why? The commonly-accepted view is that if we accept this then we are led into regress.⁴⁶ For if the thought is right then Achilles must ask the tortoise not only to accept that P entails C , but also to include that proposition in his set of *premises*. And then the tortoise must accept yet another hypothetical before he is compelled to infer, and so on. The decisive mistake is made in thinking that the proposition newly accepted must be included as a *premise* of the argument. (Whether or not it might yet be in some other way required as part of the tortoise's *reason* for believing the conclusion I shall return to in chapter 6.)

⁴³ Indeed, this *is* what the tortoise says in Carroll's story. Thomson (1960) and Stroud (1979) both remove the 'must'. See Stroud (1979, p. 181).

⁴⁴ Thomson (1960, p. 100-102)

⁴⁵ Stroud (1979, p. 184-185)

⁴⁶ See, for one example, Ryle (1954).

But Thomson shows that the idea that the tortoise is only compelled to infer the conclusion when the hypothetical associated with his argument is itself one of the *premises* of that argument is even ‘more seriously confused’⁴⁷ than has been made out above. This idea is not unacceptable because we are led into regress if we *do* accept it; indeed, we should see that it is unacceptable *without* setting out on the path of regress at all. For if we present to the tortoise the argument from *P* to *C* and then ask him to accept that *the argument from P to C is valid*, or that *P* entails *C*, then *that* argument (from *P* to *C*) is already there, and neither of the propositions we might have asked the tortoise to accept is a premise of it; nor even *can* either of them be.

It is one thing to put forward an argument, even a valid one, and another to say that you are arguing validly. It is one thing to propose for acceptance propositions which (you hope or believe or know) entail another proposition and another thing to say that they do. In arguing, you may need to point out that you are. You then (as it were) step aside from what you are doing and comment on your own performance. But then the performance must be there, independently of the comment, to be commented on.

The proposition that such-and-such an argument is valid ... cannot [itself] be a premise in the argument to which it refers.⁴⁸

No matter what the propositions *P* and *C* are, *P* entails *C* is patently not a premise of the argument from *P* to *C* - the only premise of that argument is *P*, and *P* is not the same proposition as *P* entails *C*. It *cannot* be because ‘a conditional proposition cannot be part of its own antecedent.’⁴⁹ A similar argument shows that *no* argument, no matter how many premises it has, can have as one of its premises the conditional statement relating *all* of its premises to its conclusion⁵⁰ - that conditional statement would have to be self-referential and be a ‘part of its own antecedent’.

⁴⁷ Thomson (1960, p. 104)

⁴⁸ Thomson (1960, p. 100-101)

⁴⁹ Stroud (1979, p. 184)

⁵⁰ If an argument has *n* premises, none of those premises can be the proposition stating that the conjunction of all of the premises entails the conclusion. Call the set of all of the premises $S = \{P_1, P_2, \dots, P_n\}$. Then we would need to have:

If we accept that it is only when the hypothetical associated with the tortoise's argument is itself one of the premises of that argument that he must infer according to the argument, then we ought to conclude that it is *never* the case that he *must* so infer. But we ought to conclude this not because we are as a result led into regress, but by the above observations. *No* argument has the conditional proposition related to it in the given way as one of its premises, and this is all we need to see. To be led into regress we need to have *not* seen this and then mistakenly added (or asked Achilles to add) as a premise a proposition which we falsely thought would yield, when so added, the sort of argument we were after – but there is *no* argument of this sort at all.

IV-ii An Invalid Argument Against Valid Arguments

I agree with Stroud that the moral of Carroll's tale has more to do with inference and belief than logical validity; as I explain in the next section of this chapter, I shall focus on the problems the story raises for the idea of deductive inferential knowledge. But I wish first to dispel a lingering worry that a sound argument for the impossibility of *logically valid* argument can be found in the story.

I shall present a sceptical argument inspired by the tortoise against the possibility of logically valid argument. This will lead us to an idea analogous to the confused 'thought' I discussed above - the analogous idea is that an argument is *logically valid* only if one of its premises is the hypothetical associated with that same argument. Thomson's

$$((P1 \& \dots \& Pn) \text{ entails } C) \in S = \{P1, \dots, Pn\}.$$

But this is an impossibility: Assume that $(P1 \& \dots \& Pn) \text{ entails } C$ is an element of S ; then it has to be one of the propositions in the list $P1, P2, \dots, Pn$. It doesn't matter which it is, so let's say that it is Pn . Then we have

$$Pn = (P1 \& \dots \& Pn) \text{ entails } C.$$

The right hand side of the equality is Pn , and it also has Pn as part of its antecedent, as one of the conjuncts. But that means that one of the conjuncts of the antecedent of Pn is as long as Pn itself (namely, the conjunct which *is* Pn). So the antecedent of the above conditional is at least as 'long' as the whole conditional, which is absurd because we have to *add* a logical constant (the conditional) and a consequent to the antecedent to get the whole conditional.

reflections discussed above shall then be brought to bear on *this* thought. But how to deal with the sceptic's argument is, as we'll see, a separate question which I think needs further discussion.

Imagine a sceptic who claims that the tortoise shows us that there is no such thing as a valid argument. She claims that for arbitrary and distinct propositions, P and C , the truth of P is not sufficient for the truth of C . First, she points out that, for arbitrary P and C ,

if it is not true that P entails C , then the truth of P is not sufficient for the truth of C .

We accept this, so that we accept the seemingly paradoxical statement that only if P entails C is true is the truth of P (alone)⁵¹ sufficient for the truth of C . P alone is sufficient for C only if something else is also true. How can that be? The sceptic claims that it follows from this that the truth of P is in fact *not* by itself sufficient for the truth of C ; that the argument from P to C is not valid.

In a move which is reminiscent of Achilles' move above to ask the tortoise to accept a new proposition and include it as a premise, we might think that if, for arbitrary P and C , the argument from P to C is not valid then at least the argument from P , together with P entails C , to C is valid.⁵² But this, of course, will not do. We might think that it will not do because the sceptic's next move is just to point out that

if it is not true that [P and (P entails C)] entails C , then the truth of [P and (P entails C)] is not sufficient for the truth of C ;

in other words, because regress is unavoidable if we make that move.

⁵¹ 'Alone' is really redundant here – if P is sufficient for C then it is of course *alone* sufficient for C . I add it for emphasis.

⁵² One might think that Thomson's arguments above can be applied here to show that this move would be 'futile'. This is not the case – Thomson's arguments show that the strengthening of an invalid argument yields an argument with a *false premise* (as long as the conditional involved in that premise is understood as that of strict implication), and is therefore futile as an attempt to get an argument which *logically establishes* its conclusion. But we are not at present concerned with whether or not premises are true, only with whether or not arguments are valid (and which).

But the ‘move’ in question is more obviously confused than that. Making it might lead us into regress, but if we so much as make the move in the first place then we don’t fully understand what position we are in. For if we agree with the sceptic that for *arbitrary* P and C the argument from P to C is not valid (because P entails C is not a premise of that argument), then what good will it do to add premises to some argument in a defiant attempt to *get* a valid argument? To be in that position and make that move would be to think that ‘no argument is valid, but, given an argument, which will of course be invalid, we can always obtain from it an argument (its strengthened form) which will be valid.’⁵³

The sceptic might present her argument more perspicuously as follows:

For *any* proposition P , and *any* proposition C , if P entails C is not true then the truth of P is not sufficient for that of C . So, if P entails C is not a premise of the argument from P to C , then that argument is not valid. But P entails C is patently *not* a premise of the argument from P to C (‘a conditional proposition cannot be part of its own antecedent’), so the argument from P to C is not valid.⁵⁴

Now, I think that neither a regress argument, nor an appreciation of the fact that the hypothetical associated with an argument cannot be a premise of *that* argument, can help us to counter the sceptic. Rather, as we see above, one of those arguments (the latter, as I’ve presented it) is itself a part of the *sceptic’s* argument.

The question is: Is it true that, for any P and C , if P entails C is not a premise of the argument from P to C , then that argument is not valid? The first approach was to say that we can see that it is not true, that the thought is confused, by seeing that we are led into regress if we accept it. The second, preferred, approach was to say that we can see that it is not true by noting that the hypothetical associated with an argument is *never* a premise of that *same* argument. But that is precisely the sceptic’s point! We cannot reject a claim

⁵³ Thomson (1960, p. 104)

⁵⁴ This argument can easily be adapted to apply to an arbitrary argument having any number of premises. See note 50 above.

as *false* on the basis that our acceptance of it leads us into trouble - to sceptical or otherwise undesirable conclusions; just as we cannot say that Jonnie needn't tell his headmaster the truth just because his doing so would get him into trouble.

The reasoning underlying our answer, "No," to the above question seems to me to be as follows. If we answer "Yes" then we must accept that there are no valid arguments; but we cannot accept that there are no valid arguments, so we must answer "No". We must not accept that the hypothetical associated with an argument needs to be a premise of that argument if it is to be valid, because it follows from this that there are no valid arguments.

To be *justified* in answering "No", however, we must show that the first step in the sceptic's above argument is invalid; that it does not follow from

if *P* entails *C* is not true then the truth of *P* is not sufficient for that of *C*

that

if *P* entails *C* is not a premise of the argument from *P* to *C*, then that argument is not valid.

We agree that apart from having *P*'s truth, we must have the truth of some other proposition, distinct from *P*, in order for the truth of *C* to be guaranteed. But is it not then required that this other proposition must be a premise of the argument for *C* if the truth of that argument's premises is to be a guarantee of its conclusion's; if that argument is to be valid? The answer here again is "No", and we can see why. We think otherwise only because we neglect the fact that a proposition is never needed as a premise of an argument in order for that argument to be valid as long as that proposition is either a logically *necessary* truth, or guaranteed to be true by the truth of one of the *other* premises of the argument.

Most of us (though not the sceptic at hand) would agree that the argument Achilles originally presents to the tortoise is valid and that the hypothetical associated with any valid argument is necessarily true. So let us show that a logically necessary proposition is never required as a premise of an argument in order for that argument to be valid – if an argument having a logically necessary proposition as a premise is valid then so is the same argument without that proposition as a premise.

Assume that for some propositions P , Q , and C , the truth of P and Q guarantees the truth of C , and that Q is a logically necessary truth. We must show that the truth of P by itself guarantees the truth of C ; that there are no circumstances in which P is true and C is false.

Assume there is a circumstance in which P is true and C is false. Then, in those circumstances Q is true, since it is necessarily true. So, in those circumstances P and Q are true and C is false. But, by hypothesis, there are no circumstances in which P and Q are true and C is false - the truth of P and Q guarantees the truth of C . We have a contradiction, so our assumption is false: there are no circumstances in which P is true and C is false; the truth of P guarantees the truth of C .^{55 56}

I do not claim to have shown here that there are indeed any valid arguments; but I do think that we have shown our sceptic's argument against the possibility of valid argument to be unsound. It does not follow from

if P entails C is not true then the truth of P is not sufficient for that of C

that

if P entails C is not a premise of the argument from P to C , then that argument is not valid,

⁵⁵ I am grateful to Bernhard Weiss for assistance in the presentation of this argument.

⁵⁶ A similar argument can of course be given for the case in which Q is not *necessarily* true, but the truth of P guarantees that of Q .

for it might well be the case that the former is true and that the argument from *P alone* to *C* is valid, if *P entails C* is necessarily true. To show that we require, in addition to the truth of *P*, the truth of some distinct proposition in order for the truth of *C* to be guaranteed, is not to show that the given proposition is required *as a premise*, alongside *P*, of the argument for *C*, in order for that argument to be valid. So the sceptic does not get as far as she would like to by pointing out what we all know: For arbitrary *P* and *C*,

if it is not true that *P entails C*, then the truth of *P* is not sufficient for the truth of *C*.

Unless she has another argument that, for arbitrary *P* and *C*, *P entails C* is not a logically necessary truth (and also not logically entailed by *P*), we can easily accept this. We can accept it, that is, without thinking that an argument is only valid if it has, *per impossible*, the hypothetical associated with it as one of its premises; and therefore that there are *no* valid arguments.

Earlier, we realized that if we want to maintain that there are logically valid arguments then we *must* not answer in the affirmative to the question: “Is it true that if *P entails C* is not a premise of an argument from *P* to *C*, then that argument is not valid?” But now we see why we *need* not do so.

IV-iii A Challenging Argument Challenging Deductive Inferential Knowledge

The above reflections, I believe, certainly merit the few pages I have dedicated to them, and will also be valuable for my treatment, in chapters to come, of the more challenging problem raised by Carroll’s story. This problem has indeed, as Stroud has said, more to do with inference and belief than logical validity. When at first presented with Achilles’ original argument, the tortoise doubts the truth of the hypothetical associated with it, the falsity of which would entail that the truth of the premises of that argument was insufficient for that of its conclusion. But ‘the conjunction of [its premises] does imply

[the conclusion], whatever the tortoise happens to think about it,⁵⁷ and we have now seen how this can be the case. Even if the tortoise, for some P and C , doubts the proposition (3) below,

- (1) P
- (2) P entails C
- (3) $[P \ \& \ (P \text{ entails } C)]$ entails C ,

and even if the truth of (3) is required in order for the truth of P together with that of P entails C to be sufficient for C 's truth, it might still be the case that the argument for the conclusion, C , from P together with P entails C (but not (3)) is logically valid.

But now, firstly, it does *prima facie* seem that if the tortoise doubts that (3), then he is not under any compulsion to infer the conclusion when he believes that P and that P entails C . Indeed, anyone who genuinely doubts or even *ought to* doubt that (3), it would seem, really *ought to* doubt that C when his 'evidence' for it is only that P and that P entails C .⁵⁸ And I think that a sceptic might present an argument against the possibility of deductive inferential knowledge along similar lines.

For any P and any C , if you have only P and P entails C as reasons for believing that C , then if you ought to doubt that (3), you ought to have doubt about C (even if you in fact *know* that P and that P entails C). But at any point in time at which you don't actually know of the truth of some proposition, you ought to have *doubt* about its truth. So when you know that P and that P entails C , you ought still to doubt that C unless you *also* in fact *know* that (3).⁵⁹

⁵⁷ Stroud (1979, p. 180)

⁵⁸ See my argument for this in the next chapter.

⁵⁹ Note that the conclusion of the sceptic's argument is that when I know that P and that P entails C , I ought still to *doubt* that C unless I also in fact know that (3). Treating this as the claim that in such a case I can't *know* that C unless I know also that (3) is to subscribe to the view that if I ought to doubt something, then I can't know it, which we would do if we thought that knowledge is necessarily certain. See chapter 3.

And now, if all this is right, we have a real difficulty. If I am required to have some item of propositional knowledge in order to be able to have *inferential* knowledge that *C*, how can it not be further required that I have that knowledge as part of the *basis* on which I believe that *C*? How can my possession of some item of propositional knowledge help me to know some conclusion unless the fact that the relevant proposition is true is one of my *reasons* for believing the conclusion? And if it is one of my reasons, is it not one of my *premises*?

We could accept that (3) needs to be true in order for *P* together with *P entails C* to be sufficient for *C*, without concluding that the argument for *C* from only *P* together with *P entails C* was not valid. If (3) is necessarily true, we said, then it is redundant as a premise in any argument insofar as we want only for that argument to be *valid*. But a proposition is not superfluous as a premise of an argument of mine just because it is a necessary truth insofar as I want that argument to provide me with *inferential knowledge* of its conclusion. If an argument from *P* to some *C* is valid but not at all *obviously* valid, then I might be able to have inferential knowledge that *C* by inferring it from *P together with P entails C*, but not by inferring it from *P* alone – even if *P entails C* is indeed a *necessary* truth.

But if we accepted that, for any propositions *P* and *C*, *P entails C* must be a premise of my argument (from *P* to *C*) in order for me to be able to have inferential knowledge that *C* by the employment of that argument, then we would have to accept that I cannot have inferential knowledge at all. We might think that I could know inferentially that *C* by inferring it from *P* together with *P entails C* - and then find ourselves led into regress - but that would be to forget that no argument has the hypothetical associated with it as one of its premises. If I need to have the hypothetical associated with an argument as one of the premises of *that* argument for it to allow me to have inferential knowledge of its conclusion, then I can never have inferential knowledge.

This, then, is the central problem I shall engage with in this essay. To look ahead, I shall argue that I do indeed need to know that *P entails C* in order to be able to have inferential

knowledge that *C* on the basis of my knowledge that *P*; that *P* entails *C* does, in such a case, and in a certain sense, need to be one of my *reasons* for believing that *C*; but that in the sense in which it is required as one of my *reasons* for believing that *C*, it is not required to be one of the *premises* of my argument for that conclusion. To argue for and work out the details of this will be, in the main, the task of chapter 6.

In chapter 3 I shall explain the motivation for the special interest I have in *deductive* inferential knowledge, and for accepting an internalist constraint on the attainability of deductive inferential knowledge. I shall argue that an argument must be deductively *valid* and I must *know* that it is deductively valid if its employment is to provide me with inferential knowledge - or at least inferential knowledge that is *certain*.

Chapter 3

Descartes' Doubts, Infallibilism, and Internalism

I Introduction

The question we are faced with is this: when does my knowledge that P , for some proposition P , provide me with sufficient warrant to believe some other proposition, C , so that I am entitled to infer deductively from P to C ? In chapter 2 two general constraints were suggested. For any P and any C , firstly, P must entail C , and, secondly, I must *know* that P entails C . Since we could not see how my knowing that P entails C could help me to have inferential knowledge that C except if P entails C were a premise of my argument for C , we thought that the legitimacy of these two constraints implies the impossibility of deductive inferential knowledge – the hypothetical associated with an argument is *never* a premise of that same argument.

The aim of this chapter is to provide motivation for a special interest in *deductive* inferential knowledge and to outline the form of the solution I am after with respect to the problems at hand. Though we are of course interested in deductive inferential knowledge simply because deduction is one of the most important and prominent ways in which we reason, I think that deductive argument deserves our special attention because if we can ever have *certain* inferential knowledge, it can only be by *deductive* argument. What I mean by 'certain knowledge' will be clarified below.¹

In section II of this chapter I shall argue that I can have certain knowledge that C on the basis of my knowledge that P *only if* P entails C . Thus, though I cannot have *certain* inferential knowledge via an inference which is not deductively valid, such as an inductive inference, there is hope that I *can* via an inference which *is* deductively valid. I shall also argue in that section that even if P does entail C , it must further not be the case that I ought to *doubt* that P entails C if I am to be able to have certain inferential

¹ 'Certain knowledge' is somewhat ambiguous, and might better be expressed as 'knowledge that is certain'. I stick to the former for ease of expression.

knowledge that *C* on the basis of my knowledge that *P*. Now, because anyone who is willing to infer *C* from *P* has at least the *implicit* belief that *P* entails *C*, as I'll argue in chapters to come, I think that what the second requirement, together with the first, finally boils down to, is that I must in fact *know* (at least implicitly) that *P* entails *C*.²

The first constraint (that *P* must entail *C*) is one an infallibilist would place on inferential knowledge in general, and the second (I must know that *P* entails *C*) is one an infallibilist who was also an internalist of some kind would place on inferential knowledge in general. I shall discuss and provide some support for infallibilism and internalism in section III. In the final section I shall urge that even if we are not convinced by the general infallibilist and internalist arguments, and even if we think that knowledge is not *necessarily* certain, we should think that I can at least have certain *deductive* inferential knowledge in some clear cases.

But, as I have said, for me to be able to have certain knowledge of a conclusion I arrive at inferentially, the argument I employ for that conclusion must be valid and I must *know* that it is valid. One major objective of this essay, then, is to show that we can accept these requirements without accepting further the constraint that the hypothetical associated with my argument must itself be a premise of that same argument. In the previous chapter we saw that this constraint is unsatisfiable, so if it were indeed a constraint on the possibility of having certain inferential knowledge, then nobody could ever have such knowledge - so I must show that it is *not* a proper constraint. I must argue that we can accept that I must know that *P* entails *C* if I am to have certain knowledge of the truth of *C* on the basis of my knowledge that *P*, while denying that *P* entails *C* must be another premise of my argument for *C*. This, I aim to do in the chapters to follow.

I shall now proceed to the arguments for the claim that I can have certain knowledge that *C* on the basis of my knowledge that *P* only if *P* entails *C* and only if it's not the case that I ought to *doubt* that *P* entails *C*. Before I do so, however, I should begin to clarify my

² We will also see a (suspect) argument I attribute to a sceptic for the claim that requiring that it must not be the case that I ought to doubt that *P* entails *C* implies requiring that I must in fact *know* that *P* entails *C*.

understanding of *certain knowledge*. Firstly, and most importantly, I can have certain knowledge of something only if I believe it, without any doubt³, and it is not the case that I *ought* to doubt it to any degree. Later, I shall ask whether the satisfaction of these constraints is *sufficient* for my possession of the relevant item of certain knowledge.

II Descartes' Doubts as Motivation for a Special Interest in *Deductive Inferential Knowledge*

In his first meditation, Descartes presents a sceptical argument - against the possibility of knowledge of any kind, as I see it - as follows. He argues that he can (and does) discover reasons to have some doubt about each of his beliefs, so he really *ought* to have some doubt about each of his beliefs. Then, that if he ought to have some doubt about each of his beliefs, then each of those beliefs is one he really ought not to *have*.

Reason ... convinces me that I must withhold assent no less carefully from what is not plainly certain and indubitable than from what is obviously false,

So,

the discovery of some reason for doubt as regards each opinion will justify the rejection of all.⁴

And since, as I've said, Descartes purportedly *does* discover 'some reason for doubt as regards each [of his] opinion[s]', he concludes that he really ought not to have any given one of his beliefs – he is not *entitled* to any of them. But being entitled to a belief is necessary for the belief being a case of knowledge, so Descartes draws the conclusion

³ Below I'll consider briefly whether any sense can be given to the idea of believing something while also doubting it (presumably, to some small extent).

⁴ Descartes (edition 1971, p. 61)

that he has no knowledge whatsoever⁵ (and it would follow by the same reasoning that nor does anybody else).

But now, it might be argued that what Descartes succeeds in demonstrating in his first meditation is only that we cannot have any knowledge *that is certain*. And even that this is all that he in fact *means* to demonstrate.

I am obliged to admit in the end that none of my former ideas are beyond legitimate doubt... So I must carefully withhold assent from them just as if they were plainly false, if I want to find any *certainty*.⁶ (My emphasis.)

We might, the thought is, grant that Descartes shows that I really ought to have some doubt about each of my beliefs, since he establishes that there is ‘some reason for doubt as regards’ *each* of them, but deny that it follows from this that none of the beliefs I do in fact have, are cases of *knowledge*. The important question is evidently whether it follows from the fact that I ought to have some doubt about a belief, that I ought not to *have* that belief at all. Is it really the case that if I ought to have *some* doubt about a belief, no matter how small, then I am *ipso facto* not entitled to that belief? Or is it only the case that ‘if I want to find *certainty*’, I must ‘withhold assent from’ anything which I ought to have some doubt about? Need we endorse the following principle?

(K) If one ought to have any doubt whatsoever that *C*, then one ought not to believe that *C* (and therefore cannot know that *C*).

It seems clear to me that Descartes does indeed endorse such a principle, for he is convinced that he ‘must withhold assent no less carefully from what is not plainly certain and indubitable than from what is obviously false’. But it is not my aim here to provide an interpretation of Descartes’ first meditation. And neither, as we’ll see, is it my aim to defend principle (K).

⁵ Descartes does not explicitly draw this conclusion, but I think it is implicit in his argument. See below for another possible interpretation.

⁶ Descartes (edition 1971, p. 64)

that he has no knowledge whatsoever⁵ (and it would follow by the same reasoning that nor does anybody else).

But now, it might be argued that what Descartes succeeds in demonstrating in his first meditation is only that we cannot have any knowledge *that is certain*. And even that this is all that he in fact *means* to demonstrate.

I am obliged to admit in the end that none of my former ideas are beyond legitimate doubt... So I must carefully withhold assent from them just as if they were plainly false, if I want to find any *certainty*.⁶ (My emphasis.)

We might, the thought is, grant that Descartes shows that I really ought to have some doubt about each of my beliefs, since he establishes that there is 'some reason for doubt as regards' *each* of them, but deny that it follows from this that none of the beliefs I do in fact have, are cases of *knowledge*. The important question is evidently whether it follows from the fact that I ought to have some doubt about a belief, that I ought not to *have* that belief at all. Is it really the case that if I ought to have *some* doubt about a belief, no matter how small, then I am *ipso facto* not entitled to that belief? Or is it only the case that 'if I want to find *certainty*', I must 'withhold assent from' anything which I ought to have some doubt about? Need we endorse the following principle?

(K) If one ought to have any doubt whatsoever that *C*, then one ought not to believe that *C* (and therefore cannot know that *C*).

It seems clear to me that Descartes does indeed endorse such a principle, for he is convinced that he 'must withhold assent no less carefully from what is not plainly certain and indubitable than from what is obviously false'. But it is not my aim here to provide an interpretation of Descartes' first meditation. And neither, as we'll see, is it my aim to defend principle (K).

⁵ Descartes does not explicitly draw this conclusion, but I think it is implicit in his argument. See below for another possible interpretation.

⁶ Descartes (edition 1971, p. 64)

My aim here is to show that we can try to avoid scepticism in one of at least two ways. We can either grant that I really ought to have some doubt about each of my beliefs, and therefore that I cannot have any knowledge that is *certain*; while denying (K) - maintaining that I am in fact entitled to some of the beliefs I have (though I really ought to have *some* doubt about them)⁷, and that some of these beliefs in fact amount to *knowledge*. Alternatively - if we accepted (K) – we would have to deny that I really ought to have some doubt about each of my beliefs if we wanted to avoid scepticism. We could then hope to show that some of the beliefs I have, and which I have no doubt about, are in fact cases of knowledge, and, moreover, knowledge that *is* certain.

Looking at the first alternative, it might well seem hard to see how we could maintain that though I ought to have some doubt about a given belief of mine, I am yet *entitled* to that belief. Yet this does seem to me to be the only way that someone could maintain that there is a kind of knowledge that is not *certain* knowledge.⁸ I shall not take it for granted, however, that such a view is incoherent. Perhaps someone who took this option would have to maintain that to believe something is not necessarily to have *no doubt* about it. For if believing something is having no doubt about it, then if I *ought* to have some doubt about something, I clearly ought *not* to believe it. Then, the view would be that some of the things I believe, I believe with *some* doubt (“I believe that I am sitting at my desk; but, yes, I am not entirely sure that I am”); and that some of *these* beliefs are in fact cases of knowledge – I am entitled to them, and they are in fact true, say.⁹ Because I do not dismiss this view out of hand, I shall not take for granted that to believe something is to have no doubt about it at all. Doubting something, as I will speak of it, is not simply refusing to accept it, though anyone who refuses to accept something does of course doubt it. Having some doubt about something is more like either refusing to believe it, or believing it while at the same time thinking that one “might just” be mistaken in that belief (assuming for the sake of argument that a belief of this kind is possible). I shall

⁷ See below for whether we can make sense of this idea.

⁸ Someone who held this would have to hold that some of my beliefs, which I really ought to have doubt about, are cases of knowledge all the same – which requires that they are beliefs to which I am *entitled*.

⁹ I’ll provide some argument against the possibility of such a belief of mine (even if I can have this kind of belief) amounting to knowledge; in section III-i below

also, therefore, sometimes speak below of whether someone in a given situation could be “entitled to believe something with doubt”. Though this might come across somewhat strangely, I hope to have made clear why I am willing to consider such slightly odd possibilities.

Otherwise - if it was *not* held that one can believe something while doubting it to some extent - I take it that a rejection of (K) would be even harder to sustain. Again, if believing something is having no doubt about it, then how can one be entitled to a belief even though one really *ought* to have some doubt about it? Surely, one can't be. I can only imagine that someone who held this would be thinking that though, indeed, I cannot be entitled to any of my beliefs, some of my beliefs I “can be forgiven” for having. She might say: “Some of your beliefs are ones you ought to doubt to such a small extent that I will grant you an entitlement to them, even though they are beliefs you in fact have *no* doubt about, where you *ought* to doubt them.” But we surely would not want to grant someone knowledge of something when her relevant belief was one she “could be forgiven” for having, if it so happened that her belief was, luckily for her, true. I won't consider this view any further below.

This been said, I wish to pursue the second of the alternatives for countering the sceptic suggested three paragraphs above – the option of maintaining that we can and do have *certain* knowledge some of the time. I do so, however, not because I am fully committed to (K) – we might be able to have knowledge which is *not* certain *as well as* knowledge which *is* certain. I shall *not* argue for the truth of (K) or concern myself with the question of whether or not knowledge is necessarily certain. I shall show only in this section that there are infallibilist and internalist constraints on the attainability of *certain* knowledge. In the next section I shall briefly present what I think are quite persuasive arguments for general forms of infallibilism and internalism. But, as I'll explain later, we might not accept infallibilism and internalism in general and yet want to accept infallibilist and internalist constraints on the attainability of *deductive* inferential knowledge. This is because we might naturally think that we can, even easily, have *certain* knowledge by the employment of a *deductively* valid argument; for quite often we know that we have

conclusive evidence for the truth of some conclusion, and in these cases we think that we are entitled to have no doubt about that conclusion. “At least when I *know* that I have *conclusive* evidence for some conclusion, I can surely have *certain* knowledge of that conclusion,” – as I have said, there are infallibilist and internalist constraints on the attainability of *certain* knowledge. So while in this essay I am after an infallibilist and internalist solution of the problems with which we are faced with respect to *deductive* inferential knowledge, the desire for such a solution can be motivated independently of whether or not we are infallibilists and internalists in general, and whether or not we think that *all* knowledge is certain knowledge.

Let’s consider, then, what is required in order for me to be able to have certain knowledge. Firstly, with Descartes, I’ll argue that if I infer *C* from *P* alone where *P* does not entail *C*, then I cannot thereby come to have certain knowledge that *C*. Secondly, I’ll argue that I cannot have certain knowledge that *C* if I infer *C* from *P* though I ought to doubt that *P* entails *C*.

Descartes’ arguments show that many of our beliefs are based on inconclusive evidence, evidence which does not guarantee their truth. For simplicity and brevity, let’s construe my ‘evidence’ for a belief as including both my propositional knowledge, if any, as well as my perceptual evidence, if any, (where the latter may be part of the former) on which I base that belief. Descartes’ argument is that for many of my beliefs, that *C*, say, my evidence for *C* is no guarantee of its truth. For such a *C*, where *P* is a codification of all of the evidence I have for *C*, *P* does not entail *C*. And if *P* does not entail *C*, the argument suggests, then I ought to doubt that *C* given that my only evidence in favour of its truth is that *P*. Now, because if I ought to doubt that *C* then I don’t have certain knowledge that *C*, it follows that if *P* does not entail *C* then I cannot have certain knowledge that *C* on the basis of my knowledge that *P* alone.

Many of our beliefs are based on perceptual evidence - the evidence of the senses - and Descartes first brings into question the perceptual ‘knowledge’ he has always taken himself to have. He shows, by the dream argument, that the evidence he has for his belief

that he is sitting in front of the fire does not guarantee that he is in fact sitting in front of the fire. The argument is that he could in fact be dreaming in bed, and be in possession of precisely the same (or near enough the same and no weaker) evidence (than) he now has while in fact sitting in front of the fire. If Descartes were in bed, though, he would of course *not* be sitting in front of the fire, so if we codify in *P* all the (perceptual) evidence he has for his belief, call it *C*, his argument is meant to show that *P* does not entail *C*. He concludes that he ought to doubt and doesn't know for certain that he is sitting in front of the fire. The principle of the argument would appear to be¹⁰ that for any propositions, *P* and *C*,

- (A) If *P* does not entail *C* then I ought to doubt that *C* when my only evidence for *C* is that *P*.

If *P* does not entail *C* then, in such a case as above, *C* could be false given only my evidence for it. The fact that *P* does not entail *C* is a reason for doubting that *C* in such a case, so I ought to doubt to at least some extent that *C*.¹¹

If this is right, then only beliefs which are based on *conclusive* evidence can be beliefs I am entitled to have no doubt about. So if we think that we *are* in some cases entitled to have no doubt about something, then, since one way (and perhaps the only way) to have conclusive evidence for something is to *infer* it from something else whose truth *guarantees* its truth, we should be especially interested in *deductive* argument.

Now we might naturally think that inferring *C* from *P*, where I know for certain that *P* and where *P* entails *C*, is *sufficient* for being entitled to have no doubt about *C*. The thought is that though I cannot have certain knowledge on the basis of perceptual evidence or by the employment of an *inductive* argument, surely I *do* come to have certain knowledge via a *deductively* valid argument when I am certain of its premise(s). For if I believe that *C* on the basis of my knowledge that *P* where *P* *does* entail *C*, then

¹⁰ See the next note.

¹¹ "What if I don't know and can't be expected to know that *P* does not entail *C*?" The possibility of such a scenario will be discussed below.

my belief that *C cannot* be false given my evidence for it. Am I not therefore entitled to have no doubt about the truth of *C*?

Unfortunately, this is not the whole story. Even if *P does* entail *C* and I have certain knowledge that *P*, it may still be the case that I ought to doubt that *C*. This is the case as long as I ought to doubt that *P* entails *C*, notwithstanding that it in fact does. An example is easy to come by. Imagine that a friend of mine is, as I well know, easily convinced of things and tells me one day that some *C* is, “without doubt”, true. When I ask him why he believes that *C*, he provides as his reason another proposition, *P*, and nothing more; so he evidently believes that *P*'s truth is a guarantee of *C*'s. Now, let's say that we both know for sure that *P* and that I go ahead and join my friend in his belief that *C*, even though I cannot see and have no reason to think that *P* entails *C*, apart from the fact that my friend evidently thinks that it does. It is clear that even if it turns out that *P*'s truth is a *guarantee* of *C*'s, so that the argument from *P* to *C* is deductively valid, I ought yet to doubt or have doubted that *C* in these circumstances, all the more given the fact that I am aware of my friend's gullibility. And the obvious diagnosis of this is that I clearly ought to doubt that *P* entails *C* in such circumstances.

Or, if there is an objection to this example on the grounds that in the case described I have not inferred *C* from *P* alone, since part of my reason for believing that *C* is given by the testimony of my friend, imagine the following. Practising my deductive inferences one day, I infer *C* from *P* with confidence though I know that I am not a particularly good reasoner and often carry out invalid inferences even when I feel confident about them. As it turns out, however, *this* inference *is* valid. That being the case surely does not make it so that I am entitled to have no doubt about the truth of my conclusion. The diagnosis is the same:

- (B) If I ought to doubt that *P* entails *C* then I ought to doubt that *C* when my only evidence for *C* is that *P*.

To say that I ought to doubt that P entails C is to say that I ought to doubt that P 's truth is a guarantee of C 's. But if that's the position I'm in, then I ought to think that C could, for all I know, be false, even though P is true.¹² So I ought to doubt that C .

So far I have argued that I ought to doubt that C when my only reason for believing it is that P as long as either P does not entail C , or I ought to *doubt* that P entails C . But there is a possible objection to (A) which runs as follows. Even if P does not entail C , it is not the case that someone ought to doubt that C when she knows for certain that P if she cannot be held epistemically irresponsible in *taking it* that P entails C , by inferring C from P , say. For instance, imagine that a young student is learning about deductive inference, and learning which are valid and which invalid inferences from the example of her teacher. She knows for certain that P , for some P , and follows her teacher in inferring some C from P , though her teacher on this occasion makes a mistake and infers invalidly. It could be argued that this student is entitled to have no doubt about the truth of C even though P does not entail C , since she is in no position to question or doubt the judgement of her teacher. If we agree with this then we might think that (A) needs to be supplemented in some such way as the following.

(A)' If P does not entail C and I ought to realise that it does not, then I ought to doubt that C when my only evidence for it is that P .

Then, since if I ought to realise that P does not entail C then I clearly ought to *doubt* that P entails C (because believing that something is false necessarily involves doubting it), (A)' could be subsumed under (B). It is when I ought to *doubt* that P entails C that I ought to doubt that C when my only evidence for it is that P , and a scenario in which I ought to realise that P does *not* in fact entail C is just a special case.

But if that's right, why does Descartes move so readily from having shown that some of his beliefs are based on inconclusive evidence to conclude that he ought to have doubt

¹² To say that I ought to believe that C could, *for all I know*, be false, though P is true, is not to say that I ought to believe that C could be false though P is true. That would be to say that I ought to believe that P does *not* entail C , which clearly does not follow from the fact that I ought to *doubt* that P entails C .

about those beliefs? If we go along with what I have just said above, we can explain this as follows. Descartes' arguments show not only that many of his (and many of our) beliefs are based on inconclusive evidence, but that he (and we) ought to appreciate this fact; we ought to recognise, for instance, that we might not be sitting by a fire even when it appears to us as if we are. From this it follows by (A)' that we ought to have some doubt about those beliefs.

This may well be the right view of the matter, but I wish to apply some pressure on it. I cannot in any case, I'll maintain, have *certain* knowledge on the basis of inconclusive evidence. For one thing, I am not entirely convinced by the above argument and the example of the young student that someone might be entitled to have no doubt about something though she has only inconclusive evidence for it. Even a young student in the process of learning to distinguish valid from invalid arguments from her teacher ought to know that her teacher, like anyone else, some of the time errs in what she believes and infers. Thus, she ought to doubt to at least some extent that a particular inference her teacher carries out is good, and that a particular conclusion she reaches is true. And, of course, 'to some extent' is all we need.

Furthermore, I think that even if (A) is unacceptable, even if I can be entitled to have no doubt about something which I believe on the basis of inconclusive evidence, the following is surely right.

(A)" If P does not entail C then I cannot have *certain knowledge* that C when my only evidence for C is that P .

Even if I can be entitled to have no doubt about C when I infer it from P alone, though P does not entail C , and even if C happens to be true in one such case - it *could* have been, given only my evidence for it, false. I don't think we can rightfully say that I can have certain knowledge that C here. Imagine that a second, older and brighter student in the class sees that the teacher above has inferred incorrectly and therefore does not carry out the same inference. What we have is one student who has inferred C from P where P

does not entail *C*, and one who has not, who does not believe that *C*. But let's say that *C* happens in fact to be true. I think that even if the first student was entitled to have no doubt about *C*, we could not say that she *knew* for sure that *C*. What I am suggesting here is that being entitled to have no doubt about something (it not being the case that I ought to doubt it), is not sufficient, even with that something's truth, for having certain knowledge of it. What is also at least required is that my belief could not have been false given my evidence for it.^{13 14}

What I have argued in this section is that if I am to have *certain* knowledge that *C* on the basis of knowledge that *P*, it must be the case that *P* entails *C* and not the case that I ought to doubt that *P* entails *C*. But all of this leaves untouched the question of whether this is required if I am to be able to come to *know*, perhaps *without* certainty, that *C* on the basis of knowledge that *P*. I shall now turn to brief discussions of infallibilism and internalism.

III Infallibilism and Internalism

III-i Infallibilism

As we've seen, I have many beliefs which are based on inconclusive evidence. I believe that I am sitting at my desk on the basis of the perceptual evidence I have for the truth of this proposition; I believe that the next raven I come across will be black on the basis of knowing that all previously observed ravens have been black; I believe that it will rain tomorrow because I have read that it will in the newspaper. In none of these cases does my evidence (again broadly construed to include both propositional knowledge and perceptual evidence) rule out the falsity of what I believe on that basis. I could be asleep and dreaming, and therefore not sitting at my desk; though nobody has ever observed a non-black raven, there could be one which I am about to come across in my back yard; I

¹³ This would of course make (A)" trivially true.

¹⁴ See (III-i Infallibilism) below for related discussions.

could have read the wrong newspaper, or there could have been a misprint, or, of course, the weather-man could just have been wrong.

Now, I often claim to know such things as that I am sitting at my desk and that it will rain tomorrow. But how can I know something without having eliminated all possibility of it being false? As Williams expresses the infallibilist intuition, there is surely “a kind of incoherence in combining a claim to know that *P* with an admission that one might nevertheless be in error.”¹⁵

The infallibilist holds that beliefs such as these, based on inconclusive evidence, never amount to knowledge. For a belief, that *C*, based on (propositional and/or experiential) evidence *P* to count as knowledge, it must be the case, she says, that *P* entails *C*. So knowledge that *P* does not provide me with sufficient warrant to believe and thereby *know* that *C* unless *P*'s truth guarantees *C*'s. On this view, if one shows (as Descartes does with his dream argument) that some *P* is compatible with the truth of some other proposition, *Q* (that I am asleep in bed, say), which is itself *incompatible* with *C* (that I am sitting in front of a fire), one thereby shows that one cannot come to *know*, in any sense, that *C* on the basis of knowledge that *P* alone. For one demonstrates in that way that the truth of *P* is compatible with the falsity of *C*, that is, that *P* does not entail *C*.¹⁶

One way to argue for infallibilism is to point out that if I infer *C* from *P* where *P* does *not* entail *C*, then I am in at least one sense lucky¹⁷ to have a true belief that *C* when *C* is indeed true. Given all and only my reasons for believing that *C* (given only the truth of *P*), *C* could (logically) have been false. It just so happened that it was true so that my belief that *C* was true. And if I was lucky that my belief was true, the argument goes, I did not know that *C* was true. If *C* is for me only luckily true and we allow that in such a case I *know* that *C*, then since truth is of course required for knowledge, what I have is lucky *knowledge*. But there is, one might think, no such thing.

¹⁵ Williams (1991, p.48)

¹⁶ See chapter 5 - (III All that My Knowledge Entails, p. 97) for related discussion.

¹⁷ See Bach (2003) for types of 'lucky' knowledge.

Another argument is to the effect that I have been epistemically *irresponsible* if I have come to believe that *C* on the basis of *P* alone, where *P*'s truth does not guarantee *C*'s. As Boghossian has put it, 'being justified is, at least in part, a matter of being epistemically blameless,'¹⁸ and being justified is necessary for having knowledge. I have provided some argument in the previous section for the claim that I ought to doubt to some extent that *C* if my only evidence for it is that *P*, where *P* does not entail *C*. If this is right and if *I ought not to believe that C* follows from *I ought to doubt that C*, then we have that if *P* is my only reason for believing that *C*, where *P* does not entail *C*, then I ought not to believe that *C*; I believe irresponsibly if I do. Now, I think it is indeed plausible that if I ought to doubt that *C* then I ought not to believe that *C*.¹⁹ Shouldn't reason convince us 'that I must withhold assent no less carefully from what is not plainly certain and indubitable than from what is obviously false'?

But even if I can be entitled to believe something though I ought to doubt it, how could I have knowledge in such a case? If I believed it *without* doubt, I would surely believe irresponsibly, since I *ought* to doubt it. So, then, what if I believed it with some doubt (if that is possible)? *Perhaps* I could be entitled to this belief, but we would surely not find coherent someone who said, for some *C*, "I believe that *C* though I am not absolutely sure of it," and then, perhaps after she had discovered that *C* was in fact true, "See, I *knew* it to be true." There is 'a kind of incoherence in combining a claim to know that [*C*] with an admission that one might nevertheless be in error.'²⁰

The possibility of a charge of epistemic irresponsibility also brings out a second sense in which, when *C* is true, I am lucky to have a true belief that *C* when I have based that belief on *P* alone, where *P* does not entail *C*. If what I have said two paragraphs above is right, then I ought really not to have believed that *C* on the basis of my knowledge that *P* in such a case. Now, someone else, June, say, who also knew that *P* but was perceptive or vigilant enough to realise that *P* did not entail *C*, and therefore did not infer *C* from *P*,

¹⁸ Boghossian (2003, p. 228)

¹⁹ See my discussion in section II above.

²⁰ Williams (1996, p. 48)

would end up not knowing that *C* in such a case because, of course, she didn't even believe it. If we allow that I, on the other hand, *know* that *C* even though I acted irresponsibly by inferring it from *P* (certainly more irresponsibly than our other believer, or non-believer if you like) then I end up knowing something as a direct result of having been irresponsible in what I believe. But if I know something not because I have believed responsibly but precisely because I have, indeed, believed *irresponsibly*, then I am certainly in a sense quite lucky to have that knowledge. (Incidentally, I also think that if I can be entitled to believe something though I really ought to doubt it to some (small) extent, and if I do believe it with some doubt (if that is possible) and end up *knowing* it, then I am in a clear sense lucky to know it.)

III-ii Internalism

Following Alston, using 'the term 'epistemizer' to range over anything that affects the justification of a belief, positively or negatively,'²¹ internalism in epistemology can be stated as the view that

- (PI) The only thing(s) that can epistemize my belief that *C* is (or are) some of my *other* justified belief(s).²²

The intuition is that any fact of which I am unaware, which is external to my 'perspective', makes no difference to *my* epistemic state. Such a fact is something I don't know of, so its obtaining (or not) has no effect on the justification I have for my belief that *C*.

This form of internalism can be called 'perspectival internalism' – only what is within my *perspective* can epistemize a belief of mine, where my perspective is circumscribed by the totality of my justified beliefs. Alston rejects perspectival internalism for reasons I

²¹ Alston (1989, p. 192)

²² Alston, on behalf of the perspectival internalist (which he is not), goes on to refine this formulation of 'perspectival' internalism as he sees fit. This simple formulation of it, however, suffices for our purposes.

shall not go into, and goes on to develop and discuss²³, but also reject, a more favourable version of internalism which he calls 'access internalism'. I shall state this as the view that

- (AI) The only thing(s) that can epistemize my belief that *C* is (or are) some of my *other* justified belief(s), and other thing(s) I *can come* by reflection alone to believe with justification.

According to (AI), my justification for a belief is a function only of other things I either believe with justification, or am able to believe with justification, by reflection alone - in my current internal epistemic state. If a fact obtains though I cannot at present, without any further gathering of evidence from things external to me, come to have the justified belief that it does, this fact's obtaining neither positively nor negatively affects the justification I have for any one of my beliefs.

I think that it is intuitively quite a compelling idea that, loosely stated, only things I *know* or am able to *know* can epistemize a belief of mine – the degree of justification of a belief *of mine* depends only on what justification *I have*, internally, so to speak, for it. I have in fact presented, earlier in this chapter, arguments which rest on some such internalist intuition. Having certain knowledge that *P*, I said, is not sufficient for being entitled to have no doubt about *C*, even when *P* entails *C*. (Remember me and my gullible friend.) I advocated a further condition, motivating it by the provision of such examples which would have seemed persuasive only to an internalist.²⁴ The further constraint was that it must at least also not be the case that I ought to doubt that *P* entails *C*.

But, as an aside, and to look ahead to the development of my ideas in chapters to come, let's look at how this internalist constraint I introduced connects with (PI). To be able to have certain knowledge that *C* on the basis of knowledge that *P*, I must firstly, of course, be willing to infer *C* from *P*. But that involves believing, I shall argue, at least implicitly,

²³ See Alston (1989, p. 211-23)

²⁴ See chapter 4 for further, even more compelling, examples.

that *P* entails *C*. Then, it must be the case that *P* *does* entail *C*. Lastly, it is required that I am *entitled* to believe that *P* entails *C* (it must not be the case that I ought to doubt that *P* entails *C*). But all of this would seem to indicate that what we require is that I must in fact *know* that *P* entails *C*.²⁵

Turning back to another argument of the previous section, I suggested at one point that the fact that *P* does not entail *C* does not by itself make it the case that I am not entitled to believe that *C* on the basis of my knowledge that *P*; the fact that *P* does not entail *C* does not epistemize my belief that *C* if it is a merely external fact. If it is further the case, however, that I ought to *realise* that *P* does not entail *C*, then my belief that *C* is indeed negatively epistemized. And since if I ought to realise something then I *can* realise it, the added fact that I ought to realise that *P* does not entail *C* turns what is otherwise a merely external fact into an *internal* one, as characterized by (AI) above.

Because it seems on the face of it that some version of internalism is correct, and because internalism is evidently well in line with our actual practices of the ascription of justification and knowledge, I rather think that it is up to anyone opposed to it to somehow appease the uneasiness felt at a rejection of it. I'm not sure, though, that this can be done, because I think that for any externalist account of justification an appropriate example can be thought up in which a belief of some subject is justified according to the account, but quite intuitively *not* genuinely justified, there being some 'merely external' fact which supposedly, somehow, *makes* it justified. There is an example, that is, which brings us back to our internalist commitments. Indeed, because it is simply natural to operate within an internalist framework, support is most commonly provided for internalism by providing examples such as I have given in this chapter. A scenario is described in which something 'good' about a subject's belief or formation of belief holds (such as it being impossible for her belief to be false given her evidence for it, or there being a general reliability in her formation of beliefs, or whatever) but where this fact is unknown and possibly unknowable to the subject. When we have such a hypothetical scenario in mind we imagine the 'external' fact *not* obtaining and naturally

²⁵ See chapter 7.

think that in such a case the subject at hand is yet equally (*un*)justified in her given belief. Thus, Bonjour describes the following scenario.

Norman, under certain conditions which usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against the belief. In fact the belief is true and results from his clairvoyant power under circumstances in which it is completely reliable.²⁶

As Boghossian writes, '[o]ur robust response to this case is that Norman is not justified' and he notes that our natural reaction to this example and others like it is to say that 'if a subject is to be justified in believing some proposition *p*, he must have to hand a *reflectively accessible* warrant for the proposition that *p*'.²⁷ But Boghossian thinks that this 'natural reaction' of ours leads to insurmountable difficulties – the problems of regress we have seen - and he therefore rejects it. Indeed, I think that a withdrawal from internalism is generally motivated by the appearance of internalism being highly problematic, and not by any independent appeal of some opposing externalist account. Thus, in attempting to show how an internalist can deal, to her satisfaction, with the supposedly intractable problem of regress with respect to inferential warrant, I hope in this essay to undermine such motivation for any theory of inferential entitlement which is unpalatably externalist.²⁸

Before closing this section, I shall show that if access internalism is right then so is a particular version of infallibilism: the reasons one has and *can have* for a belief must guarantee its truth if that belief is to amount to knowledge. Assume the truth of AI and assume that the only reason I have and can have reflectively accessible to me for or

²⁶ Bonjour (1985, p.41)

²⁷ Boghossian (2003, p. 228)

²⁸ My account will indeed have a small externalist component, but one which will be acceptable except to a quite extreme internalist; it will be in large part internalist, and will allow us to maintain that we do sometimes have deductive inferential knowledge that is certain. See chapter 7.

against believing that *C* is that *P*, where I do know that *P*, and where *P* does *not* entail *C*. According to access internalism, the fact that I know that *P* is the only fact relevant to the questions of whether or not I am justified in believing that *C* on that basis, and whether or not I know that *C* when I infer it from *P*. But now, since *P* does not entail *C*, there is a possible world in which I have inferred *C* from *P*, where I know that *P* is true, and in which *C* is *false*. Since the falsity of *C* entails that I don't know that *C*, in this possible world I don't know that *C*. But, as we've said, my epistemic status with respect to the belief that *C* is precisely the same in the actual world (in which, let's say, *C* is in fact true) as in the possible world we have imagined. Therefore, just as I don't in this possible world know that *C* is true, since it's *false*, I don't in the actual world know that *C* is true. The internalist idea is that I know precisely the same with the respect to *C* in both of these worlds because the reasons I have and *could have* for believing that *C* are precisely the same in both. The mere fact that *C* is true in the actual world does not, by hypothesis, affect in any way what reasons are available to me for *believing* that it is true.

Though my arguments here are somewhat cursory, since I don't have the space to give these issues a more detailed treatment, I have tried to make evident the persuasive force of the infallibilist and internalist arguments. In light of such considerations, I find fallibilism and externalism at the very least highly counter-intuitive. With respect to fallibilism, as Lewis puts it, "to speak of fallible knowledge just sounds contradictory... If you are a contented fallibilist, I implore you to be honest, be naive, hear it afresh. 'He knows yet he has not eliminated all possibility of error.'"²⁹

As for externalism, it's hard to see how something of which I am incapable at this time of being aware (something I '*know* nothing of', we ordinarily say) can affect in any way my justification for a belief or my knowledge. As I've said, I suspect that fallibilism and externalism are, most commonly, moves made by theorists in an effort to avoid sceptical

²⁹ Lewis (1996)

conclusions, once it *appears* that these are unavoidable for the *infallibilist* or *internalist*. I hope in this essay to go some way to showing, therefore, that such conclusions are *not* unavoidable even if we remain to a very large degree faithful to our infallibilist and internalist intuitions, with respect to inferential knowledge at least.

IV Know it or Not

The *analysis* of knowledge is not, however, a central concern of mine here. The aim of this chapter is to motivate a special interest in *deductive* inferential knowledge, and to outline the shape of the solution I seek to the problems raised in chapter 2.

What we have learned, at least, is that if Descartes is right, a great deal of what we believe ought to be doubted, or at least cannot amount to *certain* knowledge. Once we are shown, by his arguments, say, that many of our beliefs are based on inconclusive evidence, we ought to accept that those beliefs could, for all we know, be false - we can't be sure of them. This realisation might then lead us to hold out hope for the attainability of certain knowledge in the case of beliefs inferred from known (with certainty) premises in accordance with what are known (with certainty) to be *valid deductive* arguments. At least, the thought is, if I know for sure that *P* and that *P* entails *C*, then I can also know for *sure* that *C*.³⁰

But this idea raises for us the difficulties raised in the previous chapter. The thought is that we *can* have inferential knowledge that is certain; so, as I have argued, if we subscribe to it we must accept both an infallibilist and an internalist constraint in some form. *P* must entail *C* if I am to be able to have certain knowledge that *C* on the basis of knowledge that *P* (the infallibilist constraint), and I must also *know* in some sense that *P* entails *C* (the internalist constraint). In the next chapter I shall provide further support for

³⁰ I think that the considerations of this chapter explain why we feel the idea of deductive inferential knowledge to be far less problematic than the idea of *inductive* inferential knowledge - even if we think that induction must surely allow us to gain knowledge, in some sense, and even if we agree with Haack (1976) that the problem of the *justification* of deduction is in fact analogous to the problem of the *justification* of induction.

this latter, internalist, constraint, and consider what (as evidenced by our actual practices) this constraint does *not* amount to. The question then is whether the satisfaction of these two requirements is *sufficient* for me to be able to have certain inferential knowledge, for these are the two requirements that seem to get us into the predicament of requiring that the hypothetical associated with an inference of mine be a premise of that very inference. In chapter 4 I shall also, therefore, examine arguments directed against the propriety of these two requirements more closely, to see just how persuasive they are.

Chapter 4

The Sceptic Up Close

I Introduction

Let's consider more carefully the sceptical argument, introduced in chapter 2, against the possibility of (certain) deductive inferential knowledge, which began as follows.

For any P and any C , if you have only P and P entails C as reasons for believing that C , then if you ought to doubt that (3), you ought to have doubt about C (even if you in fact *know* that P and that P entails C).

To move from the claim above that in such a case, if I ought to *doubt* that (3),

- (1) P
- (2) P entails C
- (3) $(P \ \& \ (P \ \text{entails} \ C)) \ \text{entails} \ C$

then I ought also to doubt that C , to the claim that in such a case, if I don't in fact *know* that (3), then I ought yet to doubt that C , I said that the sceptic might argue as follows.

At any point in time at which you don't actually know of the truth of some proposition, you ought to have *doubt* about its truth. So when you know that P and that P entails C , you ought still to doubt that C unless you *also* in fact *know* that (3).

But I have also proposed, in chapter 3¹, another view of the matter according to which, similarly, I cannot come to know for certain that C when my only reason for believing it is that P , unless I also in fact *know* that P entails C . I said that to be able to have certain knowledge that C on the basis of knowledge that P , I must firstly be willing to infer C

¹ Section III-ii, starting on p. 46.

from *P*. And that involves believing, at least implicitly, that *P* entails *C* (I shall argue for this below and in the chapters to follow). Then, it must be the case that *P* *does* entail *C*. Lastly, we have seen that it must *not* be the case that I ought to doubt that *P* entails *C*, in which case I would have, arguably², sufficient *entitlement* for my belief that *P* entails *C* so that all of this adds up to requiring that I must in fact *know* that *P* entails *C*.

Now, I must remind the reader that the sceptical argument we have on our hands is indeed an argument against the possibility of *certain* inferential knowledge. I have argued, along with the sceptic, that if I ought to doubt that *P* entails *C* then I ought to *doubt* that *C* when my only evidence in favour of *C* is that *P*. Therefore, I cannot have *certain* knowledge that *C* in such a case; but I have not argued that I cannot know at all that *C* in such a case. But, as I hopefully made clear in chapter 3, the question of whether knowledge is necessarily certain is not one I propose to answer. I aim only to show that I can indeed have *certain* deductive inferential knowledge in some circumstances. Given this as my aim, I shall, for the sake of simplicity, write of knowledge as if it *is* necessarily certain for the remainder of this essay. So from here on, in my sense, I am entitled to believe a given proposition if and only if it's not the case that I ought to doubt it. Further, when I say that I know that *C* I mean that I have *certain* knowledge that *C*, so that for one thing it's not the case that I ought to doubt that *C*; when I say that I believe that *C* (or infer *C* from *P*), I mean that I believe (or come to believe) that *C*, and have *no doubt* about it; when I say that my knowledge that *P* provides me with sufficient warrant to believe that *C* (or that I am entitled to infer *C* from *P*), I mean that I do (or can) have certain knowledge that *C* if I infer it from *P*; etc.

In section II of this chapter I shall present the sceptic's argument in simplified form: as incorporating in part the outright claim that if I don't in fact *know* that *P* entails *C*, then my knowledge that *P* does not provide me with sufficient warrant to believe that *C*. By a consideration of our actual practices of ascribing entitlement to an inference I shall show that we do in practice accept this as a general constraint. Then, again by looking at cases in which we do and in which we don't ascribe entitlement to an inference to a subject, I

² See chapter 7.

shall show what we *don't* in practice take this constraint to amount to. The constraint we accept in practice is not that I must in fact know *explicitly* that *P* entails *C* if my knowledge that *P* is to provide me with sufficient warrant to believe that *C*; and it is also not the constraint that I must *be able to* know *explicitly*, by reflection alone, that *P* entails *C*. Therefore, by elimination, it would seem that what we accept in practice is that I must know *implicitly* that *P* entails *C* if my knowledge that *P* is to provide me with sufficient warrant to believe that *C*.

In section III I shall present arguments from Gilbert Ryle against the propriety of the constraint in this form. In chapter 5 I shall present the first part of my account of inference, to be refined and developed more fully in chapter 6. The idea in chapter 5 will be to return to the sceptic's original argument given above – to maintain that the requirement is that it must not be the case that I ought to doubt that my argument is valid, while denying that it follows from this that I must in fact *know* that it is valid. The reflections of chapter 5 will lead into the next chapter, where I take them further so that we once more endorse the view that anyone who is entitled to an inference knows, at least implicitly, that the inference is valid. The challenge will be to show that remaining faithful to this view does not lead us to the conclusion (by regress-related difficulties, say) that nobody can ever have (certain) inferential knowledge.

II The Sceptic's Argument, the Natural Reply, and the Problem of Regress

The sceptic we have on our hands doesn't just deny every claim to knowledge that we make. If she were that kind of sceptic we wouldn't find her 'argument' at all convincing. That's because *that* kind of sceptic doesn't *have* an argument. She simply says: "You don't know that," to everything I claim to know. But we are not persuaded at all, by her merely making assertions, that those assertions are true. Of course, she's given us *no reason* to be.

Our sceptic, however, *does* have an argument. Let's say that for some *P* and some *C* I have carried out the inference from *P* to *C*, where the argument from *P* to *C* is of the form

of, say, modus ponens³, so that *P* *does* entail *C*. The conclusion of the sceptic's preliminary argument is that my knowledge that *P* does not provide me with sufficient warrant to believe that *C*. Now, as is always the case, if this argument is valid and sound, then its conclusion is true - then I am not entitled to the inference from *P* to *C*⁴, and cannot come to have knowledge that *C* as a result of an inference of *C* from *P*. Therefore, if we want to hold fast to the idea that I *can* have deductive inferential knowledge by the employment of the MP argument at hand, we must show either that her argument is *not* valid, or that it is *not* sound. We cannot simply elect to ignore such a sceptic.

II-i The Sceptic's Argument and the Most Natural Reply

I have said that a sceptic could, perhaps, argue for her claim that I must know that *P* entails *C* in order to be entitled to infer *C* from *P* by arguing that it follows from the fact that if I ought to doubt that *P* entails *C* then I am not entitled to that inference. I shall now provide further support for her claim by showing that we quite naturally accept it, both explicitly and implicitly in our practices, and are quite strongly disinclined to give it up.

We have set up a scenario in which I have inferred some given *C* from some given *P*, where the argument from *P* to *C* is a modus ponens argument. Now imagine that the sceptic challenges my entitlement to that inference in a quite natural way by objecting: "But you don't know that *P* entails *C*." Her argument is as follows,

If you don't know that *P* entails *C*, then your knowledge that *P* does not provide you with sufficient warrant to believe that *C*.

You don't know that *P* entails *C*.

Your knowledge that *P* does not provide you with sufficient warrant to believe that *C*,

³ *P* would be of the form: *Q* and (If *Q* then *C*). The fact that this *P* would be the *conjunction* of the two premises of a modus ponens argument does not affect my argument in any way, so I will speak below of modus ponens arguments as if they had only one premise – a conjunction of the above form.

⁴ As I've said above, I treat 'my knowledge that *P* provides me with sufficient warrant to believe that *C*' as equivalent to 'I am entitled to infer *C* from *P*'.

and it is clearly valid; so, since we cannot accept its conclusion, we must show somehow that the argument is not sound; that at least one of its premises is false. Now, I think it most natural, and correct (see coming chapters), to accept the first premise while rejecting the second in such a case as this – where my argument from P to C is of the form of modus ponens (and, indeed, in any case in which I do in fact know that the argument at hand is valid).

With respect to the first premise, a plausible defence of it (simpler than the ones we have just rehearsed at the start of this chapter) might run as follows.

If you know that P , and P provides you with sufficient warrant to believe that C , then there must be some sort of relationship which holds between P and C . In this case, since we're looking at *deductive* arguments, the required relationship is one of logical entailment. But the converse is not true: even if P *does* entail C , if you are unaware of that fact, then you don't know enough to be justified in believing that C . In other words, if you don't *know* that P entails C , then your knowledge that P does not provide you with sufficient warrant to believe that C (premise 1).⁵

I've provided some support for this form of argument in my discussion of internalism in chapter 3.⁶ The basic (perspectival) internalist idea underpinning the argument is that the fact that P entails C can have no effect on the justification I have for the belief that C , unless I in fact *know* that P entails C ; for how, we asked rhetorically, can the mere fact that something obtains, where I *know* nothing of it, have any effect on what I otherwise do or don't *know*? We saw that support for internalism is usually given by the description of scenarios in which some given fact obtains, and where a certain belief of a subject would be positively 'epistemized' if she *knew*, but not if she did not know, that fact to obtain (recall Bonjour's clairvoyant example). Below we will see some such examples as these, examples which are relevant to us here.

⁵ See Alston (1989)

⁶ See section III-ii, starting on p. 46.

In any case in which the truth of some P convinced me of the truth of some other C , that is, in which I inferred C from P , I would reject the sceptic's second premise that I don't know that P entails C . The following reply to the sceptic would surely be incoherent: "You are right, I don't know that P entails C . All the same, the truth of P convinces me of the truth of C ." The claim that I don't know that P entails C seems to be a direct challenge to my entitlement to carry out this inference (which is why the sceptic often won't make the argument explicit in the way I have above, she will simply state premise two - "You don't know that P entails C " - leaving the rest implicit). She's said an equivalent of: "You don't know that the inference you've just carried out is valid," implying that *by* carrying it out I've indeed endorsed it as valid. So replying: "You are right, I don't know that P entails C ," would seem to amount to: "Yes, I am *not* entitled to that inference."

In the case of a simple argument proceeding according to modus ponens, I can't imagine anyone faced with the sceptic's challenge naturally doing anything other than emphatically denying her second premise, and thereby implicitly accepting her first. We all know, after all, for instance, that

This mug is red and If this mug is red then this mug is coloured

entails

This mug is coloured.

We would, if we had carried out the inference with which this hypothetical is associated, and been challenged in the above way by the sceptic, retort that we do (of course) know that the inference is valid – that it's premises jointly entail its conclusion. And if we found that someone, somehow, did not know that this argument was valid, we would indeed say that she had no entitlement to the given inference.

I, along with many of us, would do the same if my premise was, for some particular, perhaps large, values of x , y , z , and n ,

(*P-Fermat*) x , y , z and n are natural numbers and n is greater than or equal to three,

and my conclusion

(*C-Fermat*) $x^n + y^n$ is not equal to z^n ,

when I have calculated neither the value of $x^n + y^n$, nor of z^n . (Below, I will call the inference from *P-Fermat* (*P-F*) to *C-Fermat* (*C-F*) the Fermat-inference⁷). Here I would claim that my knowledge that *P-F* provides me with sufficient warrant to believe that *C-F* because I know that *P-F* entails *C-F*. And I know this because I know that Andrew Wiles has proved Fermat's Last Theorem, which states that any inference of this form is valid. On the other hand, a schoolboy who has never heard of this proof and has no other good reason for believing that *P-F* entails *C-F* (such as, just perhaps, that his teacher has told him that it does) is not entitled to make such an inference, even though as a matter of fact *P-F* does entail *C-F*. His knowledge that *P-F* does not provide *him* with sufficient warrant to believe that *C-F* simply because he *doesn't* know that *P-F* entails *C-F*.

Such considerations lead us to believe that it is indeed a requirement that a reasoner must know that *P* entails *C* if her knowledge that *P* is to provide her with sufficient warrant to believe that *C*. At least, that this is our intuition, and in accordance with our actual practices. And the intuition is of course an *internalist* intuition – it is not good for my inference to be valid unless I actually *know* that it is valid.

Furthermore, these examples suggest that it makes no difference whether or not the inference is 'obviously' valid, as one might describe it in the case of a modus ponens inference. If someone claims that I don't know that my inference is valid in such a case, I

⁷ This 'Fermat-inference' example is Boghossian's (2003).

reply that of course I do, thereby accepting the legitimacy of the challenge and the sceptic's first premise above, just as I do if I have carried out the 'non-obvious' Fermat-inference. And in any case in which I have carried out an inference from some *P* to some *C* and decided, upon reflection, that I in fact don't know that *P* entails *C*, I would then accept the sceptic's conclusion that my knowledge that *P* does *not* provide me with sufficient warrant to believe that *C*. As I have said above, any other response would seem to be incoherent. I can't accept that I don't know that *P* entails *C* and yet maintain that I am justified in believing (that is, being *convinced*) that *C* on the basis of *P*. It seems, therefore, that any warrant I have, or lack, for believing that *P* entails *C* carries over as a corresponding entitlement I either have, or lack, respectively, for inferring *C* from *P*. And this is true in both the case of an obviously valid and a non-obviously valid inference.

II-ii The Problem of Regress

Then why don't we just reject the sceptic's second premise and reply: "Of course I know that *P* entails *C*"? The worry here - apart from that of *how* I know that my inference is valid in the case of a simple *modus ponens* inference, say⁸ - is that this kicks off a vicious regress. The idea is that I need to know of more than just the truth of this given proposition, *P*, in order to have sufficient warrant to believe that *C*. I need to also know that *P* entails *C*; but then, surely, I need both *P* and *P entails C* as *premises* of my argument for *C*; the single premise I had was not enough. But now, when I infer *C* from these two premises,

(1) *P*,

and

(2) *P entails C*,

the sceptic repeats: "But you don't know that the two premises of this *new* argument together entail your conclusion! That is, you don't know that

⁸ See chapters 5, 6, and especially 7.

(3) $(P \& (P \text{ entails } C)) \text{ entails } C$.

And just as $P \text{ entails } C$ was required as a premise of your argument for C , so is (3), that is, you need a *third* premise.” And so on *ad infinitum*. Then it would seem that knowledge of the truth of my premises, whatever and however numerous they may be, will never provide me with sufficient warrant to believe my conclusion; I can never have the sort of deductive inferential knowledge we are after.

It is worth discussing a certain thought we might have at this point, though in truth nothing much hinges on it. We might think, in light of the considerations of chapter 2, that to present a *regress* argument in the above way is not the best way for our sceptic to set out her argument. Isn't the sceptic's argument that in order to have deductive knowledge via some inference, the hypothetical associated with that inference must itself, *per impossible*, be one of its premises? No - in the way that we have set out the sceptic's argument here, this is not right. The sceptic has not argued that for *arbitrary* P and C I must know that $P \text{ entails } C$ and have it as one of my premises alongside P in order to be able to have inferential knowledge that C ; she has argued for a *particular* P from which I have inferred a *particular* C , that I must know that $P \text{ entails } C$ and have it as a premise of my (thereby *distinct*, strengthened) argument for C if I mean to acquire knowledge that C .

From this it would indeed follow that I am not entitled to this *particular* inference from P to C , since $P \text{ entails } C$ is of course not a premise of that argument; but it would not follow that I am not entitled to *any* inference whatsoever. For perhaps the sceptic cannot apply her argument to *any* given inference I carry out. For example, we *just might* accept that I cannot have inferential knowledge that $C-F$ by inferring it from $P-F$ alone - that $P-F \text{ entails } C-F$ must be included as a premise of a new argument for $C-F$ - but not accept that any further premises are then required. “The new inference is *obviously* valid,” we might say.⁹

⁹ See Thomson (1960, closing pages).

When (or if) the sceptic has shown for particular P and C that I am not entitled to the inference from P to C , she has still not shown that I am not entitled to another inference, the one from P together with P entails C , to C , say. It is, indeed, as we can see, not quite so silly to show someone who claims that I don't know that a given inference is valid, that I in fact do, and even sometimes to include the relevant proposition as a premise of a new argument for C , if I want to convince her that I do in fact have justification for believing that C . It is quite common that someone challenges my entitlement to an inference in the above way, but not common at all that someone challenges, in this way, my entitlement to *all* inferences I am inclined to carry out. So it is quite natural to think that the regress – if it begins at all - comes to a halt at some point. But the problem is to say at *what* point, and *why*?¹⁰ The regress must surely come or already have come to a halt when I have the modus ponens inference from our particular P to our particular C , my entitlement to which the sceptic challenges above. Therefore, it is clear that we must not take a step along the path of regress from *this* point. At *this* point, we must not agree with the sceptic that I am not entitled to infer C from P , thinking that I might be entitled to infer C from P together with P entails C ; we must not add P entails C to our list of premises *here*. That means we must reject one, or both, of the two claims made by the sceptic which do indeed force us into a strengthening of the argument: (a) Your knowledge that P does not provide you with sufficient warrant to believe that C unless you know that P entails C ; (b) If you need to know that P entails C in order to be able to have inferential knowledge that C , then P entails C must to that end be one of the premises of whatever argument you give for C .

Now, I have already provided ample support for (a), so it is clear where, on my view, we need to attack the sceptic's argument: we must show that we can reject (b). Though I need to know more than just that P if I am to have sufficient warrant to believe that C - I

¹⁰ Note that what commonly happens is that a challenger rests content when I have convinced her that I know that a given inference is valid. More importantly, note that she at that point rests content that I was entitled to *that given* inference, not to any strengthened version thereof. So, as we'll see, I think that the regress *need* never get started, though I have explained why it is some of the time natural, and not objectionable, to take just a step or two along its path.

need also to know that *P* entails *C* – we must maintain that I don't need *P entails C* as a *premise* of my argument for *C*.¹¹

Let's go back to the sceptic's original argument (which, recall, challenged a given inference I carried out, from *P* to *C*, by modus ponens):

If you don't know that *P* entails *C*, then your knowledge that *P* does not provide you with sufficient warrant to believe that *C*

You don't know that *P* entails *C*.

Your knowledge that *P* does not provide you with sufficient warrant to believe that *C*

Without change of meaning we can reword premise one slightly as follows.

If you don't know that *P* entails *C*, then your knowledge that *P* does not *alone* provide you with sufficient warrant to believe that *C*.

I add the word 'alone' to the first premise, even though it is redundant in the given sentence, because doing so highlights the fact that the sceptic's claim here, which I endorse, is that *if* I don't know that *P* entails *C*, then I am not entitled to infer *C* from *P* *by itself*. But only *if* I don't know that *P* entails *C*; if that's not the case then, of course, the sceptic's conclusion cannot be drawn from *this* argument of hers (above). So we must maintain that if I *do* know that *P* entails *C* then I *am* entitled to infer *C* from *P*, *alone*. And now note that, firstly, I surely *do* know that *P* entails *C*, since *P* and *C* are related by modus ponens¹²; and, secondly, inferring *C* from *P* *alone* is precisely what I do in such a case – I infer the conclusion of a modus ponens inference from its two premises alone, even if I have pointed out to someone (or explicitly formed the justified belief) that the argument is valid.

¹¹ See Wright (2002, closing pages)

¹² See note 8 above.

In chapter 6 I shall argue that for any P and any C , if I know that P entails C then I am entitled to infer C from P *alone* - we do not need to take a step along the path of regress by strengthening the original argument from P to C . Thus, my view is that no matter how 'non-obviously valid' an argument is, as long as I know that it is valid I am entitled to infer according to it; just as I was entitled to infer C -F from P -F *alone* because I knew that inference to be valid. Furthermore, I shall argue in chapter 7 that anyone who believes that the premises of some modus ponens argument are true, and infers the conclusion, *does* know that argument to be valid. Indeed, anyone who has the concept of the conditional, knows of any modus ponens inference that it is valid whenever she *believes* that it is.

At the end of the present chapter, by the consideration of arguments from Ryle, we will see two major difficulties that will have to be dealt with in support of this view. First, there is the question of how my knowledge that P entails C can have any bearing on the justification I have for the belief that C without P entails C being a premise of my argument for C . In other words, how do we satisfactorily avoid the regress in this way while maintaining that I *do* yet need to know that P entails C in order to be entitled to infer C from P ? Second, I will have to defend my view that even an unsophisticated and unreflective reasoner has the belief that an inference is *valid* when she carries it out, and even quite often *knows* that it is valid.

III The Form of the Requirement - Explicit vs. Implicit Knowledge

Section III Preliminaries

To see that we are not led into regress (or to the unsatisfiable requirement of a self-referential premise) if we remain faithful to the internalist view that a reasoner must know that P entails C in order to be entitled to infer C from P ; and to see that anyone who infers C from P together with P entails C , say, also as a matter of fact knows that inference to be valid, it will help us to gain clarity on what form of knowledge is at issue. I'll start by showing what form of knowledge is *not* required.

In this section I shall argue, mainly by the provision of examples, that it is not the case that I must know *explicitly*, or be able to know *explicitly*, that *P* entails *C* in order to be entitled to infer *C* from *P*. So if I do need to know that *P* entails *C* I only need to know it *implicitly*. In the next and final section of this chapter I shall explain that the problems we face are just as much difficulties for the requirement that I must have *implicit* propositional knowledge, as for the requirement that I must have *explicit* propositional knowledge, but that focussing too narrowly on explicit knowledge can easily lead us astray.

I should begin by clarifying the explicit/implicit distinction which will be important below. A believer has *explicit* propositional knowledge if and only if she has an *explicit* belief which counts as knowledge; and a believer has *implicit* propositional knowledge if and only if she has an *implicit* belief which counts as knowledge. So the important difference is the difference between explicit and implicit *belief*.

An explicit belief that *C*, for some proposition *C*, is just what a believer has if and only if she has sincerely explicitly accepted, either publicly or privately, that *C*. She might straightforwardly accept that *C* by asserting it (again, either publicly or privately), but there are also other ways of doing so. She might, for instance, assent when presented with the question of whether *C*. One thing is clear, however: anyone who either has or *can* have the explicit belief that *C* must be able to understand some sentence which expresses *C*, and therefore must be able to understand the component expressions of at least that sentence. If she could not, there would be no way for her to accept that *C*. But a believer might come to believe explicitly that *C* without having the ability to formulate for herself the content of that proposition in an expression of (any of) her language(s). For she might yet be able to *understand* a sentence which is such a formulation of *C*'s content, and go ahead to accept it.¹³

¹³ My conception of implicit knowledge differs from that of some other authors. It is not, however, to any large degree unconventional.

The notion of *implicit* belief is a more difficult and complex one, but we can see that it has an important and surely legitimate role to play in our rationalising of the actions of intentional beings (such as, importantly, believers and inferers). We might say of a dog that goes to the kitchen *in order to* eat that she believes that her food is in the kitchen; she goes to the kitchen *because* she believes that her food is there.¹⁴ Similarly, as I'll argue in chapter 6, a person sincerely asserts some proposition because she believes that proposition to be true, and deductively infers from a premise to a conclusion because she believes that the truth of the premise guarantees the truth of the conclusion. But in none of these cases (and certainly not in the case of the dog) is it necessarily the case that the doer has explicitly accepted any formulation of the content of the proposition which we nevertheless rightfully, we think, take them to be committed to - to *believe*. These, then, are cases of *implicit* belief, and we'll see in section III-iii that we *treat* an intentional being as if she can have an implicit belief without even being *able* to have that belief *explicitly* (our dog, again, is an example, but we will see others). In chapter 6 I shall argue that to do so is perfectly in order.

So, when is it right to attribute a given *implicit* belief to a subject?¹⁵ Firstly, an implicit belief, that *P* entails *C*, say, might be manifested in an assertion that *P* entails *C*, as I might do when I am asked *why* I carried out an inference from *P* to *C*. Secondly, as suggested by the considerations of two paragraphs above, an implicit belief that *C* might be manifested in a recognition and acceptance of some formulation of the content of *C* when it is presented to the subject, even if she lacks the linguistic resources necessary for expressing that proposition herself. (In both of these cases a subject comes to have an *explicit* belief which she previously only had *implicitly*.) A young child might infer validly with entitlement from *P* to *C* but be unable to provide *P* entails *C* as her reason for having done so, simply because she doesn't know how to express that proposition. Yet, if faced with the question, "Does *P* entail *C*?" in one form or another - "If *P* is true, *C* must be true?" - she may well understand the question and reply to it in the affirmative, thereby showing that she *does* believe that *P* entails *C*. Thirdly, I think that an intentional

¹⁴ This example was given to me by Bernhard Weiss, in conversation. See a reference to "Chrysippus' Dog" in a note in Wright (2002, p. 72).

¹⁵ See Weiss (2002, p. 28-31) for a discussion of implicit belief and knowledge.

agent might have an implicit belief without even being *able* to have that belief *explicitly*. Again, as I'll argue below, a young child might carry out a basic deductive inference from *P* to *C* though she cannot even understand any expression codifying the content of *P* entails *C*. It is, I think, yet right to ascribe to this child the implicit belief that *P* entails *C*, if she can demonstrate an understanding that her inferential behaviour is constrained by certain norms; for instance, by correcting and accepting a teacher's corrections of her mistakes.¹⁶ Indeed, I shall argue later that anyone who cannot demonstrate in some such way as this that she has at least the implicit belief that *P* entails *C* cannot rightfully be said to have committed a deductive *inference* from *P* to *C* at all – for (almost) the whole point of a deductive inference is to infer a conclusion which *follows from* one's premises.

Section III Overview

What precisely is the internalist requirement at hand? I have been speaking loosely of it as the condition that I must *actually* know that *P* entails *C* if I am to be entitled to infer *C* from *P*, which would be the requirement prescribed by a *perspectival* internalist, as discussed in chapter 3. But in that chapter I also discussed *access* internalism, the corresponding requirement of which would be the requirement that I must merely be *able* to know, by reflection alone, that *P* entails *C*. As I've already said and as I'll argue below and in the chapters to follow, because I think that anyone who infers or is willing to infer deductively from *P* to *C* does as a matter of fact have the *implicit* belief that *P* entails *C*, and that the question at hand is just whether this belief is warranted (and true) or not, I think that the perspectival internalist's requirement is what we ought to accept in this context (even if not in all).

I've argued that our intuition is that I must know that *P* entails *C* in order to be entitled to infer *C* from *P*, and that this is in keeping with our actual practices of ascribing entitlement to an inference always *together with* knowledge of its validity. So, to focus in on the form that the requirement should take, we can take note of the circumstances in which we in fact ascribe entitlement to an inference, along with knowledge of its validity.

¹⁶ Weiss (2002, p. 31)

The primary aim of this section is to point out that we do so even when a subject does not and also when she *cannot* have *explicit* knowledge of the validity of her inference.

In section III-i I shall show, by the provision of clear cases, that a reasoner can be entitled to an inference from *P* to *C* though she doesn't actually have *explicit* knowledge that *P* entails *C*, because she hasn't explicitly accepted that *P* entails *C*. In III-ii I'll consider and counter an objection to the effect that a reasoner *does explicitly* accept that *P* entails *C*, if not before then in the course of inferring *C* from *P*, and that the question is whether or not this explicit belief is a case of knowledge. In section III-iii I shall argue that a justified inferer might even be *unable* to have such knowledge *explicitly*, because she might not be able to *understand* any formulation of the content of *P entails C*, in which case she can't explicitly accept it.

Yet I think that in all of these cases we remain inclined to attribute propositional knowledge of the validity of the inference carried out. If we are right to do so, the knowledge possessed must therefore be *implicit*. In section III-iv I'll explain briefly why I think that we *are* indeed right to do so, elaborating on this in the chapters to come. Part of what I'll do in section III is to explain Ryle's two objections to this view

III-i Explicit Knowledge; Explicit Premise

We might think that the sceptic's view is that when I know that *P*, for some *P* which entails *C*, I must have further the *explicit* knowledge that *P* entails *C* prior to inferring *C* if I am to be entitled to that inference; then that, accordingly, *P entails C* must be made explicit as a premise of my argument for *C*. If that's her idea, then we can refute it by describing clear cases in which someone is entitled to an inference though she doesn't even have *explicit* knowledge of its validity, let alone have that knowledge deployed as a premise of her argument. To see this is to take a step forward because it suggests that an inference from *P* to *C* is often just that – an inference of *C* from *P alone*, without suppressed premises. It is also an important observation because if the requirement at hand were indeed the requirement that I must know *explicitly* that *P* entails *C* in order to

be entitled to infer C from P , then the sceptic's next move, which kicks off the regress, would *seem* fairly uncontentious. To know explicitly that P entails C is to have accepted it with entitlement when P entails C is in fact true. But, the sceptic would argue, I might well have done this while now, when I infer C from P , I take no account of the fact that P entails C at all. If I must know *explicitly* that P entails C in order to be entitled to an inference, then surely I need also to have brought that knowledge to bear on the inference as, the sceptic would argue, a *premise* of the inference.

But though we intuitively and in practice accept the loosely stated internalist requirement of *knowledge* that P entails C , we do not in practice accept that it is required that a reasoner must have *explicit* knowledge that P entails C prior to inferring, and nor, therefore, that P entails C must be employed *explicitly* as a premise of her argument for C . Imagine that Andrew Wiles himself has carried out an inference from the premise, P -Fermat, given above, to the conclusion, C -Fermat, and that this particular inference is one he has never before carried out, or explicitly accepted the validity of. Clearly Wiles could very easily be entitled to this inference (unless, perhaps, he has suddenly forgotten that he has proved Fermat's theorem, say) even though he does not have the explicit belief, let alone knowledge, that P -F entails C -F, and therefore does not have that proposition as an explicit premise of his argument.

We would not ordinarily call into question the entitlement Wiles has for such an inference, but if a sceptic challenges that entitlement by claiming that Wiles doesn't know that P -F entails C -F, we reply immediately by simply denying the sceptic's claim. Wiles knows, in whatever sense we require him to know, that P -F entails C -F, if anyone does – he was responsible for proving Fermat's theorem! And, we think, he infers *because* he believes that P -F entails C -F, and he infers *validly* and with entitlement because he *knows* that P -F entails C -F's, notwithstanding that proposition not being an explicit premise of his argument in this particular instance.

To provide another example, imagine that a friend of mine tells me that Sam is a bachelor (P), and that I infer that Sam is a man (C) (or that I carry out an inference of the form of

modus ponens). I am certainly entitled to this inference - my knowledge that P provides me with sufficient warrant to believe that C - even though I don't have the explicit belief that P entails C prior to carrying out the inference. The most I have in the way of relevant *explicit* beliefs is a belief that for any person, S , S is a bachelor entails S is a man, not that Sam is a bachelor entails Sam is a man. So, again, since we don't require that I have such an explicit belief in order to be entitled to the inference, we don't require that P entails C must be an *explicit* premise of my argument for C .

As with the Wiles example, however, I think that it would be hasty of us to retract our acceptance of the claim that I must know that P entails C if I am to be entitled to infer C from P . For, again, if someone challenges my entitlement to the above inference by claiming that I don't know that P entails C , I reply by saying that I most certainly do. What these examples show is just that when I do this I am neither claiming that I know *explicitly* that P entails C , nor that P entails C was, somehow, an *explicit* premise of my argument.

III-ii Explicitly Accepting that P entails C in the Course of Inferring

But it may be claimed that I do in fact (privately) explicitly accept that my premise entails my conclusion *in the process* of an inference to C , so that the above counter-examples are invalidated. We know that there is often a certain 'feeling' associated with reasoning or inferring. I consider my premise, the proposition P , which I believe to be true, trying to 'work out' what *follows* from it, and then it dawns on me, as if a light has gone on. I say to myself, perhaps, " P , ...Aha! So C ." Isn't this the experience of somehow realising, of coming to believe or know, that P entails C , so that I do in fact have the explicit belief that P entails C before I form the belief that C ? In that case P entails C might even be an explicit premise of my argument for C . Perhaps it's like this: " P , ... Now, Aha! - P entails C , so C ."

In opposition to this view, firstly, this 'feeling' is often absent when I infer. Secondly, even when it's present, I might have no such *thought*, while making the inference, about a

relationship of entailment obtaining between *P* and *C*. I might have no thoughts at all, apart from the thought that *P* is true, so, *C* is true. If something like an explicit acceptance of *P* entails *C* goes on when I infer I'm certainly often not conscious of it, but how can I explicitly accept something without being conscious of doing so? It would be all too mysterious, to say the least, to hold such a view. The feeling referred to is just that – a *feeling* we sometimes have, a feeling of confidence, not an acceptance or appreciation of the truth of some proposition.

I act quickly, with perfect certainty, and the lack of reasons does not trouble me.¹⁷

Thirdly, against the idea that *P* entails *C* could be an explicit premise of my argument, an inference from *P* to *C* is just that - an acceptance of *C* (alone) on the basis of *P* alone. If whenever I accept *C* on the basis of a belief that *P* I also in fact explicitly accept that *P* entails *C*, we would have something far more complex and mysterious than a simple inference from *P* to *C*. We would have either two inferences, one from *P* to *C*, and another inference of *P* entails *C* from some other belief I have (in which case we would surely have a regress on our hands), or just one inference from *P* to *C* together with, say, a 'seeing' or 'intuiting' that *P* entails *C*. But, aside from the undue complexity that such a view introduces for the process of a simple inference, whatever else we can coherently explain as going on 'behind the scenes', if I have *P* entails *C* as an explicit premise of my argument for *C*, we are then talking about an inference from both *P* and *P* entails *C* to *C*, not one from *P* alone to *C* (and we can easily see the regress looming).

To back this up, and fourthly, consider that if *P* entails *C* is always another explicit premise of my argument for *C*, then when I am carrying out a proof in a system of logic, I would write down *P* entails *C* as well as *P* before I wrote down *C*. But this I do not do, even if I first consult my rules of inference to work out what, by those rules, follows from my premise before I infer. And if I wrote down *P* entails *C* before I wrote down *C*, then I would also write down the proposition stating that these two jointly entail *C*, and so on. I would never get around to writing down my conclusion, *C*. So, often, even when I *do*

¹⁷ Wittgenstein (1953, remark 212)

explicitly accept that P entails C , I do not include P entails C as a premise of my argument for C . (I shall return to this observation in chapter 6.)

Fifthly, to provide further argument, I might well, upon reflection, explicitly accept that P entails C *after* I have inferred C from P . Or, more importantly, I might consider whether P entails C and realise that it does not. For instance, imagine that I am taking a course in logic and have fallen into the bad habit of affirming the consequent - when I have proofs for two propositions of the form C and *If P then C* , I often infer P . Now, however, someone has shown me that such inferences are invalid, and I have recognized that they are indeed invalid. But old (and sometimes new) habits die hard, so that I now go ahead and infer invalidly in this way again in a new case. Someone then asks me whether or not the premises of my argument jointly entail my conclusion, and I remember then that they do not. So at that point I have the explicit belief that my premises do *not* jointly entail my conclusion. But this surely can't be a case of changing my mind if I have just carried out the inference, and therefore, on the view at hand, just explicitly accepted that my premises *do* jointly entail my conclusion. What do I know now that I didn't know a moment ago that has made me change my mind? It seems far more reasonable to suppose that I didn't consider whether or not my premises jointly entailed my conclusion at all. We might say that I *felt* confident that the inference at hand was good, that is, valid, but I could not have *explicitly* accepted that it was valid - if I had stopped to consider whether or not it was valid, I would surely have realised that it was not.

... in a proof you advance from one proposition to another; but do you also accept a check on whether you have gone right? –Or do you merely say “It *must* be right” and measure everything else by the proposition you arrived at.

For if *that* is how it is, then you are only advancing from one picture to another. ¹⁸

I often carry out an inference without having any *explicit* belief, one way or the other, about its validity.

¹⁸ Wittgenstein (1956, I-137 & I-138)

III-iii Being Able to Know *Explicitly* that *P* entails *C*

If it's not the case that I need to *actually* know explicitly that *P* entails *C* in order to be entitled to infer *C* from *P*, perhaps I must only be *able* to know, *explicitly* again, that *P* entails *C*, as an accessibility internalist might have it.¹⁹ This idea might suggest itself if we consider the examples I gave in section III-i above. Though I don't in fact know explicitly that *Sam is a bachelor* entails *Sam is man* even though I am entitled to the inference from the antecedent to the consequent of that conditional, I am certainly *able* to know this *explicitly*. I *could* know this, explicitly (perhaps by an inference from *all bachelors are men*), just as Wiles would be *able* to have *explicit* inferential knowledge that *P-Fermat* entails *C-Fermat* even if he did *not* have that explicit knowledge when he inferred *C-Fermat* from *P-Fermat*. So, the thought might be, what is clearly required is not that I *in fact* know explicitly that an inference is valid if I am to be entitled to carry it out, but only that I am *able* to know *explicitly*, by reflection alone, that the inference is valid.

But to continue with our examination of how we in fact ascribe entitlement to an inference, from *P* to *C*, and ascribe propositional knowledge *that P entails C*, we see that we do both, together, in some cases in which a reasoner has neither the linguistic resources available for formulating the content of the proposition *P entails C*, nor even the *ability* to understand any such formulation. Since, as I explained in the preliminaries to this section, anyone who can't understand any such formulation can neither have nor *come* to have the belief that *P entails C explicitly*, and since we still do ascribe this propositional knowledge in such cases, and would, I think, do so even knowing that the subject cannot explicitly accept that *P entails C*, our idea cannot be the one proposed in the previous paragraph.

¹⁹ Such an accessibility internalist might say that I need to be able to know, *by reflection alone*, that *P* entails *C*. One question I don't want or need to delve into here is that of *how* I would be able to know such a thing when I carry out an inference with entitlement. For instance, would it be inferential or non-inferential knowledge? In this section I argue that it is in any case *not* required that I am able to know *explicitly* that *P entails C*, but I shall return in later chapters to the question I have raised in this note in connection with my view expressed there – the view that to be entitled to an inference I must at least *in fact* know *implicitly* that the inference is valid.

Imagine that a child is looking for her bicycle, which she knows has been hidden by a friend behind one of the two trees in her garden.²⁰ Let's say that the child knows that either her bike is behind this tree, or it is behind that tree. She then discovers, by checking, that her bike is *not* behind *this* tree, and infers that it is behind *that* one. Such a child is entitled to this inference though she may be quite unable to formulate the content of the proposition stating that her premises jointly entail her conclusion. If we assume that she doesn't have in her usable vocabulary the expression 'entails', or 'if, then', or any other which could be used to encode the content of the relevant proposition, then she certainly can't come to believe it *explicitly* without the aid of someone else who provides such a formulation for her. Yet, even if we know this about her, when we witness her behaviour we quite naturally say that she knows very well that

(*P*) (Either the bike is behind this tree or the bike is behind that tree) and (the bike is not behind this tree)

entails

(*C*) The bike is behind that tree.

When we require of the child that she must know that *P* entails *C* we are therefore requiring neither that she knows explicitly nor that she is *able to know explicitly*, without assistance, that *P* entails *C*.

²⁰ The idea of this example comes from Boghossian (2002, p. 36).

To take us closer to issues at the heart of the problem we are discussing - issues relating to rule-following, language-use, concept-formation, etc., to which we shall return later - imagine now the following. I am *training* a young child in the use of our English expressions ‘or’ and ‘not’ in the disjunctive syllogism, the rule that licenses, for any P and C , the inference

from $(P \text{ or } C) \text{ and not-}P$ ²¹
to C .

I have on the table two cups labelled P and C , respectively, and a ball.²² I put the ball in cup C while the child is not looking and cover both cups. Now I tell her that either the ball is in cup P , or it is in cup C . But, I tell her, it is *not* in cup P . I have, of course, told her two things from which it obviously follows logically that the ball is in cup C . The child understands that she is to try and find the ball, but she doesn’t yet understand ‘or’ and ‘not’. At first, let’s say, she guesses and misses often enough. With a little training, though, she starts to ‘guess’ correctly quite often. And later she gets it right ‘all the time’ (in a number of consecutive tests).²³ At this point we naturally say that she’s no longer guessing, that she now *knows* that $[(P \text{ or } C) \text{ and not } P]$ entails C . And also that she’s entitled to the inferences she’s now making. But, again, this child might not have available the linguistic resources required for expressing that a relationship of entailment exists between her premises and her conclusion, so when we say that she knows that some inference she carries out is valid, we are not saying that she knows this *explicitly*, and not saying that she can *come*, unaided, to know this explicitly.

This shows that someone can be entitled to an inference from P to C though she is unable to know *explicitly* that P entails C without someone else formulating the content of that proposition on her behalf. But what I aim to show here is that someone can be entitled to such an inference though she cannot know explicitly that P entails C at all. That follows from the obvious fact that our inferers above may, apart from not having resources

²¹ For simplicity, I conjoin the premises.

²² Or I might have more than two cups.

²³ See Wittgenstein (1953, remark 143).

adequate for formulating the content of *P entails C*, not even be able to *understand* any such formulation. We can easily imagine such a child inferring according to the disjunctive syllogism, and she is clearly not disqualified from being entitled to her inferences just because she can't understand any formulation of *P entails C*.

Paul Boghossian has argued that such examples show that someone can be entitled to an inference without being able to know at all that her premise entails her conclusion – she can be entitled to an inference though she has no logical ability.²⁴ But I think that our considerations here only demonstrate that someone might be entitled to an inference though she does not and cannot know *explicitly* that her inference is valid. 'A man knows that if p, then q if, when he knows that p, he is able to see that, consequently, q.'²⁵ And 'if the-fact-that-p convinces [someone] that q, then he thinks that if p, then q even though he never formulates it.'²⁶ As I've said, I shall argue later that anyone who infers *C* from *P* believes *implicitly* that *P entails C*, 'though she never formulates it', and even if she *cannot* formulate it or even explicitly accept it.²⁷ She shows us that she has this belief not by accepting some formulation of *P entails C*, for this she cannot do, but by her norm-governed behaviour.

This is again in line with our actual practices. We would not automatically withdraw an attribution of the belief or knowledge that *P entails C* just because we realise that one of the above children can't understand any formulation of *P entails C*. After all, we do not even consider whether or not she can understand some such formulation when we at first attribute that belief or knowledge. We do not, for instance, find it necessary to ask her whether *P entails C*, or ask her to assent to 'If *P* is true then *C* must be true', or any other formulation of *P entails C*.

²⁴ Op. cit. See also Alston (1989) and Wright (2002, p. 73-74).

²⁵ Brown (1954, p. 175)

²⁶ Brown (1954, p. 177)

²⁷ If we think that to believe a proposition one must have all of its 'component' concepts, then my view will have to be that though someone might not understand any equivalent or close relative of the expression 'entails', she can still have the *concept* of entailment. I don't think this is too far-fetched - I'll defend the view later. See chapters 6 and 7.

III-iv Implicit Knowledge

We have eliminated a number of candidates for what requirement we intuitively and in practice subscribe to, loosely stated as the requirement that I must know that *P* entails *C* in order to be entitled to infer *C* from *P*. It is the requirement neither that I must have *explicit* knowledge that *P* entails *C* nor that I must be *able* to have that knowledge *explicitly*.

Yet, as I argued above, even when we are aware that a subject who has inferred *C* from *P* has neither the explicit belief that *P* entails *C* nor the *ability* to have that belief explicitly, we in practice naturally ascribe this propositional knowledge to her just as long as we think that she is entitled to the inference. We think that knowing, in some sense at least, that a given inference is valid, is required for being entitled to carry it out. But in what sense?

The only idea left open to us, as I've alluded to, is that we require that a reasoner must have *implicit* knowledge that *P* entails *C* if she is to be entitled to infer *C* from *P*. As I've argued, it would be incoherent to be unwilling to accept that *P* entails *C* or to admit that I don't know that *P* entails *C* while remaining fully committed to the inference from *P* to *C*, being willing to carry it out. We would most likely say of anyone who refused to accept '*P* entails *C*' while inferring *C* from *P* that she did not understand '*P* entails *C*'. But if I am willing to accept '*P* entails *C*' when I infer *C* from *P*, since *I do* understand the expression '*P* entails *C*', then, by our above characterization of implicit belief, I believe *implicitly* if not (yet) *explicitly* that *P* entails *C*. Are we not likewise right to attribute this implicit belief to someone who inferred *C* from *P* even if she could not understand any formulation of the content of *P* entails *C*, and therefore could not explicitly accept that *P* entails *C*? According to what I said in the preliminaries to this section, we are indeed right to do so as long as the subject can demonstrate an understanding that her inferential activity is norm-governed.

I'll elaborate and provide further support for this in chapter 6. Suffice it to point out for now that deductive reasoning is, of course, a rational, intentional activity. When I infer *C* from *P*, my intention is to come to believe something, *C*, which *follows* from something else I believe, *P*, so my intention is for my inference to be *valid*. That's why anyone who is fully committed to an inference from *P* to *C* thereby endorses *implicitly* the truth of *P entails C*, so *believes* it implicitly, whether she is able to have the belief *explicitly* or not. Anyone who 'inferred' *C* from *P* but did not believe that *P entails C* would not really have inferred at all.²⁸

I think that it is this implicit belief that *P entails C* which is required to be a case of knowledge if I am to be entitled to the inference from *P* to *C*, and that having this knowledge is also sufficient for being entitled to *that* inference. I am entitled to an inference I carry out from *P* to *C* if and only if my associated belief that the inference is valid is in fact a case of knowledge. I shall argue in chapter 6 for the sufficiency part of this claim, contending that this idea does not lead us into regress – knowing that *P entails C* is *sufficient* for being entitled to infer *C* from *P*, and from *P alone*; *P entails C* need be neither an explicit nor an implicit premise of my argument for *C*. For now, however, I wish to allow the objections to this view a fair hearing in the final section of this chapter.

IV Telling Myths to get into Regress?

To be entitled to infer some proposition, *C*, from some *other*, *P*, I must know that *P entails C*. I have argued above that, pre-theoretically, we do in practice accept this idea implicitly as well as explicitly in appropriate circumstances – we attribute knowledge that *P entails C* to anyone who we take to be entitled to infer *C* from *P*, and we *say* that if someone is to be entitled to an inference, she must know that it is valid.

²⁸ A possible objection to this in fact, I think, succeeds in illustrating the point nicely. Mightn't I carry out an inference while knowing full-well that it is invalid, for instance, if I mean to test the vigilance of my students? No, I don't think so. For to *infer* some *C* is to come to *believe* that *C*, or to reflect that one's belief that *C* is true 'because' some other proposition from which it *follows* is true. And in the case described, I might not believe that my conclusion is true, but even if I do my reason for believing it is not the premise of the 'argument' I present to my students.

Furthermore, I think that we accept both implicitly and explicitly that anyone who *does* know that *P* entails *C* is entitled to infer *C* from *P*. Before we read Carroll's story, at least, it seems absurd indeed to say that someone knows that the truth of *P* is a *guarantee* of the truth of *C* and yet is not entitled to infer *C* from *P*. Therefore, we accept that knowing of the validity of an argument is both necessary and sufficient for being entitled to infer from its premise(s) to its conclusion.

To argue that we need not and ought not to abandon these pre-theoretical commitments, I must show that they do not deliver us into regress, and I have explained that the main task is to show that we can rightfully attribute knowledge that *P* entails *C* to someone who infers *C* from *P* with entitlement without thinking of *P* entails *C* as a further *premise* of her argument for *C*. To carry this out will be the burden of the chapters to follow, and there I shall also provide support for another central element of my view: anyone who believes that *P* and that *P* entails *C*, and infers *C*, for instance, also knows *at least* implicitly that her inference is valid.

Now, to turn briefly to the question of how, for arbitrary *P* and *C*, knowing that *P* entails *C* can be *sufficient* for being entitled to infer *C* from *P*, it would *seem*, as I have suggested, that we would certainly be in trouble in this regard if we thought that I need to in fact know *explicitly* or be *capable* of knowing explicitly that *P* entails *C*. For the original thought would then be that no matter what the (explicit) premises of my argument are, I am not entitled to infer a conclusion from them unless I *also* know, explicitly again, that the relevant inference is valid. But how then can *explicit* knowledge of the validity of my argument help me at all without this knowledge being deployed as part of my basis for believing my conclusion?²⁹ I am certainly in a more secure epistemic position when I take account of what I explicitly know to be the case than when I don't – for I can accept and know something to be the case and then, as it were, put it to the back of my mind – and to take account of it is surely, the claim would be, to include it as a premise. But, from what we have just said, a belief that *C* will also not be justified when

²⁹ True, it must be part of my 'basis' for believing that *C*, but *P* entails *C* need not be a *premise* of my argument for *C* – as I'll explain in chapter 6.

C is inferred from *P* together with *P* entails *C* unless I also know that (3) - that these two premises jointly entail the conclusion - so it can't be justified if it is inferred from *P* alone when I simply also happen to know that *P* entails *C*, unless, again, I know that (3). (And it is clear also that merely being *able* to know that *P* entails *C* puts me in an epistemic position of still *less* security.)

Thus, Stroud, who has approached the problem at hand from a slightly different angle, by asking what it is for one belief to be based on another, writes that

... for every proposition or set of propositions the belief or acceptance of which is involved in someone's believing one proposition on the basis of another there must be something else, not simply a further proposition accepted, that is responsible for the one belief's being based on the other.³⁰

It *seems* that whatever explicit propositional knowledge I have, including both that which is deployed as premises of an inference and that which is not, this does not provide me with sufficient warrant to believe some *other* proposition. What difference could it make, then, if some of my propositional knowledge is *implicit* rather than explicit? Again, if I infer *C* from *P* I am surely in at least no *stronger* an epistemic position with respect to doing so if I know also *implicitly* rather than *explicitly* that *P* entails *C*.

Getting at the correct response to this line of thought is part of the task of chapters to come. Here, however, I'll provide a glimpse of what I'll have to say there. Though the *argument* of the previous paragraph is valid, its premise, what I said *seems* to be the case, is in fact false. It is not true that whatever *explicit* propositional knowledge I have, I don't have sufficient warrant to believe some distinct proposition. If that were the case then it would indeed help me no more to have *implicit* propositional knowledge instead of or as well as some *explicit* propositional knowledge. It is not that requiring that I have *explicit* propositional knowledge leads us into regress whereas requiring it to be only *implicit* does not - the reason for rejecting the requirement of *explicit* propositional knowledge is

³⁰ Stroud (1979, p. 187)

not that the requirement leads to regress, but that we can easily see, as we did in the previous section, that such knowledge is quite simply *not* required. But taking a closer look at when and why we *do* attribute implicit propositional knowledge will allow us to see our way clear of the mistakes we have made above, in the argument of the previous paragraph. We often attribute *explicit* knowledge to someone on the basis only of them accepting some formulation of the content of the relevant proposition. But believing something is more than doing that, which is why we automatically attribute the implicit belief that *P* entails *C* to someone who *infers C* from *P*, but do *not* if she does not. Too much focus on *explicit* knowledge easily leads us to make the mistake of separating out *propositional* from *practical* knowledge, *because* we are easily led to think that accepting a formulation of the content of some proposition *is* sufficient for believing that proposition. Then we ask, with the tortoise, “Why *must* I infer *C* from *P* just because I’ve accepted ‘*P* entails *C*’?” The answer is that it’s not the case that I *must* infer in such a case, but then it’s not necessarily the case that I actually *believe* that *P* entails *C* if I have merely *accepted* ‘*P* entails *C*’.

Having said this, however, the argument of two paragraphs back can be, and *has* widely been taken to be, persuasive. I shall therefore allow it its fair due and fuller presentation now. We have seen Stroud speaking of our problem without reference to the implicit/explicit distinction at all. A few pages later he writes: ‘There must always exist some ‘non-propositional’ factor if any of [a subject’s] beliefs are based on others’,³¹ and since an *implicit* belief is a *propositional* factor as much as an *explicit* belief, the distinction is irrelevant here. Ryle has the following to say about the ‘shuffle’ of which I am in favour.

Principles of inference are not extra premises and knowing these principles exhibits itself not in the recitation of formulae but in the execution of valid inferences and in the avoidance, detection, and correction of fallacies, etc. The dull reasoner is not ignorant ; he is inefficient. A silly pupil may know by heart a great number of logicians’ formulae

³¹ Stroud (1979, p. 189)

without being good at arguing. The sharp pupil may argue well who has never heard of formal logic.

There is a not unfashionable shuffle which tries to circumvent these considerations by saying that the intelligent reasoner who has not been taught logic knows the logicians' formulae "implicitly" but not "explicitly" ; ... the skilful but untheoretical chess-player "implicitly" acknowledges a lot of strategic and tactical maxims, though he never formulates them and might not recognise them if they were imparted to him by some Clausewitz of the game. This shuffle assumes that knowledge-how must be reducible to knowledge-that, while conceding that no operations of acknowledging-that need be actually found occurring. It fails to explain how, even if such acknowledgements did occur, their maker might still be a fool in his performance.

All this intellectualist legend must be rejected, not merely because it tells psychological myths but because the myths are not of the right type to account for the facts which they are invented to explain. However many strata of knowledge-that are postulated, the same crux always recurs that a fool might have all that knowledge without knowing how to perform ...³²

So Ryle directs two complaints against the idea of implicit knowledge. Firstly, he thinks that we subscribe to a 'myth' if we are willing to attribute *implicit* propositional knowledge to someone who *acts* intelligently, in reasoning or chess or whatever, even if she *does not* and *cannot* accept those propositions *explicitly*. 'The sharp pupil may argue well who has never heard of formal logic'; a chess player may play skilfully who 'never formulates [any tactical maxims] and might not recognise them if they were imparted to him by some Clausewitz of the game.'

Secondly, Ryle contends that attributing this knowledge does not even succeed in explaining what it *aims* to – intelligent *behaviour*. For if '[a] silly pupil may know by heart a great number of logicians' formulae without being good at arguing', if the tortoise may know that *P* entails *C* and yet not infer *C* from *P*, then she or he may as well, presumably, know many more things and yet remain a silly pupil who cannot *in practice* reason well. And, as I have argued above, if no amount of *explicit* belief or knowledge

³² Ryle (1946, p. 7-8)

can ‘compel’ one to carry out a particular inference in appropriate circumstances, then certainly any amount of further *implicit* belief or knowledge can do no better.

I have made clear my view, which is as follows: For any *P* and any *C*, to be able to know that *C* on the basis of my knowledge that *P*, I must also know, if only implicitly, that *P* entails *C*; and when I *do* know (explicitly *or* implicitly) that *P* entails *C*, I am entitled to infer *C* from *P* alone.

Our most recent reflections lead us to believe that this will not do. The argument, which I will *not* ultimately endorse, is that to think that knowing at least implicitly that *P* entails *C* is necessary and sufficient for being entitled to infer *C* from *P*, is to subscribe to a ‘psychological myth’ which is in any case just as good at leading us into regress. I shall turn in chapter 5 to the first part of my account of inference and entitlement to inference which will ultimately enable us to see our way clear of these apparent difficulties.

To see our way clear of them I shall proceed as follows. With respect to our appraisal of a norm-governed act such as an inference, there are two distinct matters with which we are concerned. First, we are concerned with the question of my *epistemic entitlement* to an inference of *C* from *P*. In chapter 5 I shall suggest that the fact (if it is a fact) that I believe that *P* entails *C* when I make the inference has no role to play *here*. The question is whether I have an *entitlement* to carry out the inference, and I can have such, it can be plausibly argued, by having an entitlement *to* believe that *P* entails *C* - whether or not I in fact have that belief.

Analogously, with respect to the question of my epistemic entitlement to an *assertion* that *C*, the fact (and it surely is a fact³³) that I believe that *C* as long as I am willing to assert that *C*, might be thought to have no role to play - what is important for the question of whether I am *entitled* to an assertion that *C*, is whether or not I have an *entitlement* for the belief that *C*. Thus, in chapter 5 I shall explore the possibility of maintaining that for me to be entitled to assert that *C*, it is necessary and sufficient that I am entitled *to* believe

³³ See chapter 6.

that *C*. And of maintaining that for me to be entitled to infer *C* from *P*, it is necessary and sufficient that I am entitled *to* believe that *P* entails *C*. On this view, the role played by my belief that *C*, and that *P* entails *C*, respectively, is another matter.

But, on the other hand, it would be strange to say that I am entitled to an inference from *P* to *C* as long as I have an entitlement *to* believe that *P* entails *C*, and whether or not I *in fact* believe that *P* entails *C*. How is the entitlement I have, for the belief that *P* entails *C*, brought to bear on my inference from *P* to *C* if I don't in fact believe that *P* entails *C*? "Of what use is it to have an entitlement to believe something," one might ask, "if you don't *make use* of the fact that you have that entitlement?"

The account I shall develop following on from chapter 5 is as follows. When I infer some *C* from some *P*, I am epistemically entitled to that inference only if I have an entitlement for the belief that *P* entails *C*, and only if, as we saw in chapter 3, the inference is in fact valid. But when I carry out this inference I do as a matter of fact *have* the belief that *P* entails *C*³⁴ (so I have it with entitlement, and it is true, as long as I am entitled to the inference). This view of the matter allows us to explain, as we surely must, an inference of mine as *rational* - the second matter with which an account of inference must be concerned, as mentioned just above. My account allows us to see acts of inference as rational acts - to have an explanation of the fact that I carry out an inference when I do. But it also enables us to explain how my *entitlement* to believe that *P* entails *C*, when it is had, is *brought to bear* on the act of inference from *P* to *C*. I shall argue in the coming chapters that when I am entitled to some inference I carry out from *P* to *C*, the warranted and true belief that *P* entails *C* which I have, as I have explained, is in fact a case of *knowledge*. And that *P* entails *C* is in fact my *reason* for carrying out the inference from *P* to *C*. Knowing that an inference is valid is necessary and, we'll be able to see, also *sufficient* for being entitled to carry out that inference.

What's brought back to the table, of course, are the two challenges Ryle has presented us with: how does the fact that I believe that *P* entails *C* when I infer *C* from *P* explain that

³⁴ See chapter 6 (VI Entitlement to *Act*).

act of inference – how can we avoid regress; and how can we be justified in ascribing some belief or item of propositional knowledge to a subject who has not and even cannot accept the truth of that proposition?

Chapter 5

No Need to Know?

I Introduction

I shall indeed maintain in chapter 6 that anyone who infers from some P to some C with entitlement *does* know that P entails C , and that we can see that knowledge to be deployed in a certain way which does not lead us into regress. My view is that we are right to ascribe the implicit belief that P entails C to anyone who infers or is willing to infer from P to C . But we have just seen, in chapter 4, arguments against this view. The view is allegedly unacceptable for at least two reasons. First, because a reasoner might carry out an inference from P to C without even being able to accept that P entails C . Second, because ascribing the relevant belief to anyone in the given circumstances in an attempt to explain her behaviour leads to regress – it supposedly does *not* explain her behaviour at all.

I shall therefore, in the present chapter, consider the possibility of an account of what is necessary and sufficient for me to be entitled to an inference which does not focus on the belief I have that P entails C , if I have it, when I infer C from P . The challenge here is to present an account according to which I can be entitled to an inference from P to C whether or not I know that P entails C . Remembering that I am especially interested in the attainability of *certain* inferential knowledge, the challenge is to provide an account of what is necessary and sufficient for being able to have certain knowledge that C , for some C , by inferring it from some other P , without reference to the belief that P entails C which I have (again, if I have it) when I infer C from P .

I have previously argued that I must in fact know that P entails C in order to be entitled to infer C from P . Apart from considering how we actually ascribe entitlement to an inference together with knowledge of the inference's validity to a subject, I argued in chapter 3 that

- (B) If I ought to doubt that P entails C then I ought to doubt that C when my only evidence for C is that P ,

so I cannot have certain knowledge that C in such a case if I ought to doubt that P entails C . I then said that we can try to move from (B) to the requirement that I must in fact *know* that P entails C in such a case if I am to be able to have certain knowledge that C ; and we can do so in one of two ways.¹ First, I presented a sceptic as maintaining that if I don't at a particular point in time *know* that P entails C , then I ought at that point in time to *doubt* that P entails C . This is a clearly questionable claim which I shall not endorse in this chapter. Second, I said that we might think that anyone who infers C from P believes implicitly that P entails C , so if we require that it must not be the case that she ought to doubt that P entails C , and that P *does* entail C ², if she is to be entitled to that inference, the requirement is plausibly that the given implicit belief of hers must amount to implicit *knowledge*. I shall return to defend this view in the chapters to follow, but, as I've said, in the present chapter I shall attempt to construct an account of inferential *entitlement* which prescind from employing the notions of implicit belief and knowledge, to see how far such an account can get us.

A way of rescuing the possibility of *certain* inferential knowledge now suggests itself. Perhaps it is only required not to be the case that I ought to *doubt* that P entails C if I am to be entitled to infer C from P ; and requiring this does not entail requiring that I must in fact *know* that P entails C . The suggestion now is that we can reject, for at least some values of P and C , the *first* premise of a sceptical argument:

If you don't know that P entails C , then your knowledge that P does not provide you with sufficient warrant to believe that C .

You don't know that P entails C .

Your knowledge that P does not provide you with sufficient warrant to believe that C .

¹ See the end of chapter 2; and the beginning of chapter 4.

² Both of these constraints were argued for in chapter 3.

As I explained in chapter 4³, this is one of at least two ways in which an attempt to avoid Carrollian regress can be made. The other idea, to which I shall return in the coming chapters, is the idea that I can be required to have knowledge that *P* entails *C* without being required to have that proposition as a premise of my argument for *C*.⁴ But the thought now is that it is not important for me to in fact know that *P* entails *C* in so far as I am to be *epistemically entitled* to infer *C* from *P*, so it is certainly not important for me to have knowledge that *P* entails *C* deployed as a premise of my argument for *C*, or as, in any other way, another one of my reasons for believing that *C*.

Of course, if the account of this chapter is acceptable, it also allows us to circumvent the alleged problem of attributing a belief to someone who couldn't understand any sentence formulating the content of that belief. If it is only important that she has an entitlement *for* the given belief, and not that she in fact *has* the belief, then we don't need to make this supposedly illegitimate move.

In section II of this chapter I shall present this account, according to which it is only required to not be the case that I ought to doubt that *P* entails *C* if I am to be *epistemically entitled* to the inference from *P* to *C*. In section III we'll see that the account seems plausible by looking at what position we would be in if, in the general case, I was required to in fact *know* anything whose falsity would cast doubt on the truth of *C*, in order to know that *C*. In section IV I shall explain that we must maintain that for some arguments, such as modus ponens arguments, knowing of the truth of their premises is sufficient for having an entitlement to believe that their premises entail their conclusion. Since the sceptic will clearly claim that this is never the case, this suggests that we are in a state of deadlock with the sceptic. I'll consider how we might try to undermine the sceptical arguments which keep us at this impasse, before indicating the best way to proceed.

³ In section II-ii ('The Problem of Regress') – I presented two claims, (a) and (b), made by the sceptic.

⁴ This is the idea that we can reject the sceptic's claim, (b).

II My Beliefs are Not the Matter?

In this section I shall propose an account of what it is to be epistemically entitled to an inference which prescind from talking of the *belief* that an inferer has (that the inference is valid) when she carries out some inference. The account will later be developed to be quite satisfactory with respect to our internalist intuitions, and to maintain that I can in some cases have *certain* inferential knowledge. As far as it goes in *this* chapter, the view will have a fairly pronounced *externalist* component, but will also be partly internalist. On this view, even though *P* must entail *C* if I am to be entitled to infer *C* from *P*, the question of whether or not I know that *P* entails *C* is not a question we need ask in determining whether or not I am entitled to that inference. Therefore, we need also not ask whether I have knowledge that *P* entails *C* brought to bear in some way as part of my reason for believing that *C*. The question we need to ask, and this is the account's internalist component in basic form, is whether or not I am entitled *to* believe that *P* entails *C*. Leaving alone the question of whether or not I *do* believe that *P* entails *C*, is it or is it not the case that I *ought to doubt* that *P* entails *C*?⁵ If I am not entitled to believe that *P* entails *C*, as I've already argued, then I ought to doubt that *C* when my only 'evidence' for it is that *P*. Then, on the present view, the hope is that as long as I *am* entitled to believe that *P* entails *C* I can be sure of the truth of *C* when I know for sure that *P*. The difference between someone who is entitled to an inference and someone who is not is that the former has an entitlement *to* believe that the inference is valid, whereas the latter has none. The hope is that this can be seen to be, in a way acceptable to an internalist, a difference in their 'internal epistemic state'.

⁵ Remembering that, because I am interested in the attainability of *certain* knowledge, when I speak of someone believing that *C* I mean that she believes it, and has no doubt about it, I use 'I am entitled to believe that *C*' interchangeably with 'it is not the case that I ought to doubt that *C*'.

In this connection, there are some important questions of which I must make mention here. It would seem⁶ that for the requirement that I must be entitled to believe that *P* entails *C* in order to be entitled to infer *C* from *P* to be appropriately internalist, it must be something like the requirement that I must have a *reflectively appreciable* warrant to believe that *P* entails *C*. According to an internalist, unless I have such a warrant in such a way, I ought to doubt that *P* entails *C*. Or, to put it another way, if I can be entitled to an inference when it is because of some ‘merely external’ fact that I am entitled to believe that it is valid, then the requirement at hand – that I must have an entitlement for that belief – will not be to the liking of an internalist. But now, if to be entitled to an inference from *P* to *C* I must have a reflectively appreciable warrant to believe that *P* entails *C*, in what way am I required to have that warrant? For instance, am I required to have some item of propositional knowledge which counts as a good reason to believe that *P* entails *C*; or must I by an act of ‘rational insight’, say, be able to justify this belief? In chapter 6, I shall come back to these questions and briefly consider argument, which I see as convincing, against the idea of rational insight. But we can also already see that we cannot require that a reasoner must have such warrant for the belief that *P* entails *C* in the form of some item(s) of propositional knowledge from which she can infer *P* entails *C*. We cannot require this on pain of regress again – to be entitled to this *latter* inference the reasoner would again have to have sufficient warrant, in the same way, to believe that it was valid; and so on. But surely these two alternatives - having some item(s) of propositional knowledge which count as good reason to believe that *P* entails *C*; and being able to ‘see’ that *P* entails *C* – are exhaustive of how I could have the required reflectively appreciable warrant. What then becomes of the supposedly internalist component of our account? In chapters 6 and 7 I shall show how we can carry out a satisfying rescue operation on it.

⁶ See chapter 7.

Even with these questions left for the moment unanswered, however, I think we can see something right about the current proposal. Often, when I infer according to modus ponens – from some P together with some P entails C , to C – I do so without verifying, or even accepting, or even considering whether my premises jointly entail my conclusion (i.e., whether (3)).

- (1) P
- (2) P entails C
- (3) $(P \ \& \ (P \ \text{entails} \ C)) \ \text{entails} \ C$
- (4) *etc.*

Now, we are indeed inclined to ask here whether or not I all the same believe that (3) in such a case. But isn't that question only important for the purposes of attempting to *rationalize* my behaviour - my act of inference here? We wonder whether we can explain the fact that I have carried out the given inference by attributing to me the belief that it is valid. But for the moment we are not interested in whether or not my act of inference was *rational*, we are interested in whether or not it was *epistemically respectable*, given my current epistemic state. To answer that question we surely need not focus on what *beliefs* I have; my *entitlement* to carry out the inference comes from the *entitlement* I have for the belief that P entails C , when I have such an *entitlement*. It is, I must agree here, important for any account of inference to be in part an account of the *rationality* of inference, but for now we are concerned with the question of what it is to be *epistemically entitled* to an inference. In the next chapter I shall move on, firstly, to the question of whether the account of epistemic entitlement to an inference presented here is quite satisfactory, and then to the issue of the *rationality* of inference. There, by preserving the intimate connection which we can see to exist between *belief* and *action*, I shall explain how a belief, when warranted, can play the dual role of both *entitling* and *explaining* an act, such as an act of *inference*.

Consider again the child discussed in chapter 4⁷ who inferred from

(P) (Either the bike is behind *this* tree or the bike is behind *that* tree) and (the bike is not behind *this* tree)

to

(C) The bike is behind *that* tree,

who was a well-skilled reasoner, but couldn't understand any formulation of the content of *P* entails *C*. This child certainly cannot explicitly accept that *P* entails *C*, so there may be some doubt about whether we can attribute the belief that *P* entails *C* to her in order to *explain* her act of inference. But again, the question of whether or not this child implicitly believes that *P* entails *C* though she cannot believe it *explicitly* seems to miss the point in so far as we are concerned with the question of her *epistemic entitlement* to this inference. We noted in chapter 4 that someone who wanted to defend this child's entitlement to this inference would be inclined to say, "She *knows* that *P* entails *C*," and I do indeed think that this is right – she *does* know that *P* entails *C*.⁸ But what is important about the fact that she knows that her inference is valid in so far as her *entitlement* to the inference goes is what follows from this fact, what her defender implies by what he says: it is not the case that she ought to doubt that *P* entails *C*. A challenge to the child's entitlement to the inference could just as satisfactorily have been countered by: "You think that she ought to doubt that *P* entails *C*? But no – she *does* have an entitlement to believe that *P* entails *C*, because of ..."

But again - you might ask here - because of what? If we think, with Ryle, (as I do *not* think) that it is obscurantist to hold that someone who could not acknowledge the truth of some proposition could yet believe it; we most certainly must think (as I *do* think) that it would truly be to obfuscate matters to maintain that such a subject can have an entitlement to believe the relevant proposition by having some other item(s) of

⁷ (Chapter 4, III-iii Being Able to Know *Explicitly* that *P* entails *C*, p. 73)

⁸ See chapters 6 and 7.

propositional knowledge which count as good reasons for believing it (or, for that matter, by the possession of some faculty of rational insight which allows her to *see* that the proposition is true). A subject cannot justify a belief that *P* entails *C* by providing argument for it if she cannot understand any codification of the content of *P entails C*. So we see again that we cannot think that anyone who is entitled to some inference has an *inferential* warrant for the belief that the relevant inference is valid.

Also, to add argument that which will come in chapter 6 against the employment of a rational insight theory for dealing with the problems here at issue, it is indeed hard to *see* how someone who could not understand any formulation of the content of a proposition, *P entails C*, could yet somehow ‘see’ that proposition to be true. I can’t argue here that such a story could not coherently be told, but I do think that it would be quite a difficult one to tell, or hear.

All the same, granted that we are all agreed that the above child is the sort of reasoner who is often entitled to the kind of inference she has carried out, and that she is entitled to *that* inference, I don’t think we should want to say that she ought to doubt that *P* entails *C*; that she is not entitled *to* believe that *C follows from P*. Not to prejudge the question of whether or not she in fact believes *implicitly* that *P* entails *C*, consider that if she *were* able to believe it explicitly, and *was* to believe it explicitly, that belief surely would be warranted. That is what I mean by her having an entitlement *to* believe that *P* entails *C*, whether or not she in fact believes or is able to believe it. If we teach this child the meaning of ‘entails’, in the course of half a day or so, say, and then take her back to her tree and her bike and ask her whether

(*P*) (Either the bike is behind this tree or the bike is behind that tree) and (the bike is not behind this tree)

entails

(*C*) The bike is behind that tree,

she *will* reply in the affirmative (or we'd probably say that she hasn't yet understood 'entails'⁹), and surely *will* be entitled to the explicit belief she comes thereby to have - that *P* entails *C*. But we have not in such a case given her any reason for believing that *P* entails *C*, or trained her somehow to see it, so she surely had an entitlement to this belief all along, whether or not she in fact had the belief.

The present suggestion is that the sceptic's claim above, the first premise of her argument for the conclusion that I am not *epistemically entitled* to some given inference, misses the point. The claim, stated in other words, was that

If I don't *know* that *P* entails *C* then I am not entitled to infer *C* from *P*.

What is significant about someone who is entitled to an inference is that she has an *entitlement* to believe that the inference is valid.

If I am not entitled to believe that *P* entails *C* then I am not entitled to infer *C* from *P*.

Because this is true, as I have said, someone could legitimately challenge my entitlement to carry out an inference by pointing out that I am not entitled to believe that *P* entails *C*. Whether or not I in fact believe that *P* entails *C*, if I ought to *doubt* that *P* entails *C* then I ought to doubt that *C* when my only reason for believing it is that *P*.

Challenging my entitlement to an act of inference of *C* from *P* by claiming that I am not entitled to believe that *P* entails *C* can be seen as analogous to challenging my entitlement to an act of assertion that *C* by claiming that I am not entitled to believe that *C*, that I ought to doubt it. Similarly, this latter challenge would often naturally, but misleadingly on the present view, be expressed by the claim: "You don't know that *C* is true." Likewise, if someone challenged my entitlement to an assertion that *C*, another person taking my side might naturally defend my entitlement to the assertion by saying

⁹ See chapter 6.

that I *know* that *C*. But what is important, on the present view, about the fact that I know that *C* (assuming that it is a fact) in so far as we are concerned with my *entitlement* to the assertion that *C*, is what *follows* from that fact, what my defender *implies* by saying that I know that *C*: I am *entitled* to the belief that *C*. Now, it is indeed surely the case that I *believe* that *C* as long as I am willing to sincerely assert that *C*¹⁰, but the fact that I believe that *C* when I sincerely assert it *rationalizes* that act of assertion – when I sincerely assert that *C* at least part of my reason for doing so must be that *C*, if we are to see me as rational, and making here an *assertion* at all.¹¹ Then, as far as my *entitlement* to that act of assertion goes, what is important is whether or not I have an *entitlement* for the belief that *C*. It does of course follow from all of this that when I sincerely assert that *C*, I am entitled to that act only if I have a warranted belief that *C*¹², but this constraint might be better expressed as: the belief I have that *C* when I sincerely assert that *C* must be warranted if I am to be entitled to that act of assertion.

Imagine that in the course of a conversation where the question of John's mortality somehow becomes relevant, I claim that John is mortal, without providing justification. My conversational partner then challenges my entitlement to that assertion by claiming that I don't know that John is mortal.¹³ I might well respond to this challenge by saying: "Sure I know that John is mortal – John is a man, and all men are mortal." But what is significant about what I do here is that I show my challenger that I have an *entitlement* for the belief that John is mortal – I know that he is a man and that all men are mortal – and am for *that* reason entitled to assert that he is. The explanation of the fact that I *acted* by asserting that John is mortal is given (in large part) by the fact that I believe (or know) that he is; but the explanation of the fact that I was *entitled* to that act is given (in large part¹⁴) by the fact that I have an *entitlement* for the belief that John is mortal. After I

¹⁰ In chapter 6, I shall argue that this is true and that, analogously, I believe that *P* entails *C* as long as I am willing to infer *C* from *P*.

¹¹ Again, see chapter 6.

¹² I can be entitled to assert that *C* without *knowing* that *C* because I can be entitled to assert that *C* though *C* is false.

¹³ Never mind the fact, which I've just pointed out in a note above, that I don't need to *know* that *C* in order to be entitled to assert it – this challenge is a common and natural way to challenge someone who has asserted something. Perhaps, though, our current considerations can provide even further support for the claim that it is not a *proper* way to challenge my entitlement to an act of assertion.

¹⁴ See coming chapters.

assert that John is mortal. an opponent might more simply say that I am not *entitled* to believe that he is. If what she says is true then I now have the explicit belief that John is mortal *without* entitlement, and was certainly not entitled to my act of assertion. But if it is false then I now have the explicit belief that John is mortal *with* entitlement, and I was entitled to my act of assertion. The belief that John is mortal which I have when I am disposed to sincerely assert that he is, doesn't come into it.

To turn back to inference, imagine again¹⁵ that a friend tells me that Sam is a bachelor and I infer that Sam is a man, without providing any justification for carrying out this inference. My friend then challenges my entitlement to that inference by claiming that I don't know that *Sam is a bachelor* entails *Sam is a man*. I might well respond to this challenge by saying: "Sure I do – all bachelors are men." What is important about what I do here is that I show my friend that I have an *entitlement* for the belief that *Sam is a bachelor* entails *Sam is a man* – I know that all bachelors are men – and am for *that* reason entitled to the given inference. The fact that I have the *belief* (or *knowledge*) that *Sam is a bachelor* entails *Sam is a man* when I carry out this inference explains that *act* of inference, as I'll argue in chapter 6; but the fact that I have an *entitlement* for the given belief explains the fact that I am *epistemically entitled* to that act.

After my friend tells me that Sam is a bachelor and I infer that Sam is a man, a challenge might more simply come in the form: You are not *entitled* to believe that *Sam is a bachelor* entails *Sam is a man*. If that claim is true then I was certainly *not* entitled to my act of inference. But if it is false, and if *Sam is a bachelor* does indeed entail *Sam is a man*, then I *was* entitled to that act. My having, or not having, of the *belief* that *Sam is a bachelor* entails *Sam is a man*, doesn't *seem* to come into it.

¹⁵ See chapter 4 (III-i Explicit Knowledge; Explicit Premise, p. 68).

III All that My Knowledge Entails

If what is important for me to have an *epistemic entitlement* to an inference of *C* from *P* is that I have an *entitlement* to believe that *P* entails *C*, and the role that my belief that *P* entails *C* has to play in such a case is the distinct role of *explaining* that act of inference, then we are not led into regress by thinking that to be entitled to this inference it is important that I in fact know that *P* entails *C*, and then have that knowledge deployed as a premise of my argument for *C*.¹⁶ This helps us to understand our natural response to a challenge to an inference from *P* to *C* in the form of: “You don’t know that *P* entails *C*.” As I urged in the previous chapter¹⁷, it is sometimes natural and surely not always misguided to respond to such a challenge by saying that I *do* know that *P* entails *C*. I might even include *P* entails *C* as a premise of a *new* argument for *C* in such a case if I want to convince my challenger that I am indeed entitled to the belief that *C*.¹⁸ But we do not expect someone to claim for *any* inference we carry out that we must know that it is valid *and* have the proposition that it is valid as a premise of an argument for *C* to be able to have inferential knowledge that *C*. Rightfully so – on the present view the claim that I know that an inference of mine is valid is at least some of the time meant only to show that I have an *entitlement* to believe that it is valid. The fact that I do in fact know that it is valid, if I do, is not what matters as far as my epistemic entitlement for the inference goes. And if it is only important that I do indeed *have* an entitlement for the belief that *P* entails *C*, and not that I do *believe* that *P* entails *C*, then I cannot be required to include *P* entails *C* as a premise of a new argument for the conclusion, *C*.¹⁹

This also makes a good deal of sense if we consider the general strategy of a sceptical argument and reflect on what form it must take if it is to be a reasonable and defensible

¹⁶ We can reject the sceptic’s claim, (a). See note 3 above.

¹⁷ (II-ii The Problem of Regress, p. 60)

¹⁸ This would clearly not be the right move to show that I was entitled to the *original*, distinct, inference of *C* from *P* *alone*. By supplementing the original argument with additional premises, I imply that I was *not* entitled to the original inference.

¹⁹ In the next chapter I shall consider whether or not it is indeed sufficient, from an internalist point of view, to require that I must be *entitled* to believe that *P* entails *C* in order to be entitled to the inference from *P* to *C*. It would seem to be more acceptable to an internalist to require that I must *in fact* know that *P* entails *C* *and* somehow bring this knowledge to bear on the inference with which we are concerned.

strategy. Take P , C , and Q to be three propositions such that I believe that C on the basis of my knowledge that P , and such that either the falsity of Q would entail the *falsity* of C , or the *possibility* of C being false given only P (my reason(s) for believing it). In such a case the sceptic's strategy is to *show* me that the falsity of Q would entail the *falsity* of C , or the *possibility* of C being false given only P , and then, perhaps, to make the claim: "But you don't know that Q ." From this it is meant to follow that I don't know that C .

To clarify, for us, the three propositions are some P , some C (which I have inferred from the given P) and the proposition P entails C , respectively. If P entails C were false then C could be false given only P , and the sceptic claims that I don't know that P entails C . It is implied that I therefore don't know that C .

But on the present account the sceptic is not right to imply, in such a case, that if I don't actually *know* that P entails C then I don't know that C . And it would seem that if she *were* right to do so, then I would need to *actually* know a whole host of things related to any C in an appropriate way in order to know that C ; and many of these things I quite clearly do not need to know in order to know that C . For simplicity, let us turn to the kind of sceptical argument which chooses for its troublesome proposition one whose falsity would entail the *falsity* of something I believe. Somewhat unconventionally²⁰, I call such a proposition a 'possible defeater' of what I believe. So a possible defeater of some proposition, C , is any other proposition entailed by C – whose *falsity* would entail C 's *falsity*.

Now, for any proposition, C , there is an indefinite number of propositions which I know or can know to be possible defeaters of the truth of that proposition, and yet I don't need to know each of these propositions to be true in order to be entitled to believe that C . We know that any C entails each of the propositions in the following set: $\{C \& C, C \& C \& C, \dots, C \text{ or } Q, C \text{ or } R, \dots, \text{etc.}\}$, for any other propositions, Q and R . But I do not need to know, for each of these propositions, that it is true (and certainly need not have that

²⁰ Standardly, a defeater of C is a proposition whose *truth*, not whose *falsity*, would defeat the truth of C . Deviance from this convention makes no difference to my discussion.

propositional knowledge deployed as a reason for believing that C) in order to be entitled to believe C . Perhaps you will say: "Of course not, that is absurd," but, by the same token, the thought is, I do not need to *know* that P entails C (and do not need P entails C as a premise of my argument for C) to be able to know that C when I have P as a good reason for believing it.

The *correct* form of the above sceptical argument is as follows. When, on the basis of my knowledge that P , I believe that C , the sceptic argues that I ought really to have doubt about the truth of C in the following way. For some proposition, Q , which is entailed by C , she demonstrates to me that if Q were false then C would be false. Then, she argues that I *ought to doubt* that Q , given only my evidence for believing that C (i.e., given only my knowledge that P). She standardly does this by showing me that the truth of Q is not guaranteed by the truth of P : "Given only your knowledge which serves as your basis for believing that C , Q could yet be false. So, given only that knowledge, you ought to doubt that Q ." But since I ought to doubt that Q , and I know that if Q were false then C would be false, she argues, I ought also to doubt that C .

It can be seen that this argument depends on some such principle as the following.

Principle 1 For any propositions, Q and C , if you ought to doubt that Q and you know, or can know, that if Q were false then C would be false, then you ought to doubt that C .

And a related principle, which I think that all of us, together with the sceptic, likewise endorse, could be formulated as:

Principle 2 For any propositions, Q and C , if you ought to doubt that Q and you know, or can know, that if Q were false then C *could* be false given only all of your evidence for C , then you ought to doubt that C .

For clarification let us revisit Descartes' sceptic, who employs principle 1 by arguing as follows:

You believe that you are sitting in front of the fire (proposition C) on the basis of the perceptual evidence you have for this belief (codify this evidence in the proposition P). Call the proposition that you are not asleep in bed, Q . Now, you could have the same perceptual evidence you now have while you are in fact asleep in bed and dreaming that you are sitting in front of the fire (that is, P does not entail Q). So, given only that evidence, you ought really to have doubt about the claim that you are not asleep in bed (you ought to doubt that Q). But if you *are* asleep in bed then you are not sitting in front of the fire (if Q is false then C is false). So you ought really to doubt that you are sitting in front of the fire (that C), given only your evidence for it.

This is indeed, I think, a sound way to argue for my lack of entitlement to a belief that C . The sceptic shows me by this argument that, even given all of my positive evidence for the belief that C , there is something, Q , I ought to doubt, and that if this something is indeed false, then C is false. How could I not doubt that C after she has shown me that, and be entitled to do so?

But, again, if the sceptic presents her challenge by claiming that I don't *know* that Q , and therefore, since I do know that C entails Q , ought to doubt that C , her presentation of the argument is misleading. It is true that if I ought to *doubt* that Q in such a case, then I ought to doubt that C . But it is not true that if I don't in fact *know* that Q when I know (or can know) that C entails Q , then I ought to doubt that C . As I said above, if that were true it would follow that I must know of the truth of every proposition entailed by C (or at least of all those I *can* know to be entailed by C) in order to be entitled to believe that C . That would certainly do away with the possibility of deductive inferential knowledge – instead of being able to know that C when I know that P and that P entails C , say, I would need to know that C in order to be able to know that P in the first place. If I ever

knew anything, I would also know of the truth of each and every proposition that I knew (or could know) to be entailed by it.²¹

Our considerations of this chapter allow us to see why a sceptical argument presented in this last way would be misleading. A sceptic might indeed be inclined to challenge my entitlement to believe something, that *C*, by saying, for some *Q*: “But you don’t know that *Q*, and if *Q* were false then *C* would be false.” But it is not important that I in fact know that *Q* here in order to be entitled to the belief that *C*; what is important is that it should not be the case that I ought to doubt that *Q*, I must be *entitled* to believe that *Q*. I am inclined to defend my entitlement to the belief that *C* by claiming that I *know* that *Q* simply because it obviously follows from this that I have the required *entitlement* to believe that *Q*.

IV Modus Ponens Premises; Deadlock; Progress

Some of the time, for instance when an argument of mine is a simple one of the form of modus ponens, it is enough for me to know of the truth of the premises and infer the conclusion in order to come to have knowledge of the truth of that conclusion. In such a case, an inference from *P* to *C* of the form of modus ponens, say, I do not first check that *P* entails *C* before I carry out the inference. Similarly, if I were to witness someone inferring from some such *P* to some such *C*, as long as I was satisfied that she knew that *P* (which would require her, of course, to understand that proposition) I would grant her knowledge that *C*. But on the view at hand any subject must at least have an entitlement to believe that an inference is valid if she is to be entitled to carry out that inference. It is clear, therefore, that we must maintain that anyone who knows of the truth of the premises of a modus ponens argument, say, has the required entitlement to believe that the argument is valid.

²¹ See Stroud (1984, p. 108) and Williams (1991, p. 80-81) for related discussions I have since writing come across; on the question of whether knowledge is ‘closed under (known) entailment’.

This suggests that the sceptic and I are in a state of deadlock at this stage of the game. The sceptic aims to show that no matter what I know, I am not entitled to believe something else on the basis of that knowledge. I, on the other hand, would like to do more than simply undermine her argument; I would like to show that I *can* be entitled to believe something on the basis of some of my other knowledge.

Let's say that the sceptic accepts our account of this chapter – she accepts that for me to be entitled to an inference what is important is that I have an entitlement to believe that it is valid, and not important that I in fact have that belief, so we do not get ourselves into a crippling regress of *premises*. She now argues that it is all the same, for any *P* and *C*, not enough for me to know that *P* and infer *C* if I am to have knowledge that *C*. And this is not enough because, apart from knowing that *P*, I also need the entitlement to believe that *P* entails *C*.

I *do* think that we can *undermine* this argument in the way that has already been alluded to. The argument takes it for granted that, for any *P* and *C*, knowing that *P* is not sufficient for being entitled to believe that *P* entails *C* in order to show that, for any *P* and *C*, knowing that *P* is not sufficient for being entitled to believe that *C*. The sceptic's method (or madness), in essence, is to maintain that when she has granted me knowledge that *P*, for some *P*, she has *only* granted me an entitlement to believe that *P*. That seems alright on the face of it, at first²², but there seems to be a clear sense in which she begs the question by doing that – she purports, after all, to *show* that having an entitlement to believe one thing is never to have an entitlement to believe another. So how can she take for granted that for no *P* and *C* is knowledge that *P* sufficient for having an entitlement to believe that *P* entails *C*?

All the same, the sceptic wishes for me to show that my possession of some item or items of knowledge can be sufficient for me to be entitled to believe something else. So I present her with an obviously valid inference, of the form of modus ponens, say, and

²² It is indeed true that when she has granted me knowledge that *P*, she has only explicitly *granted* me an entitlement for the belief that *P*. but the question is whether having an entitlement for one belief is *in fact* sometimes sufficient for having an entitlement for *another* belief.

claim that I know of the truth of its premise(s), and therefore am entitled to believe its conclusion. But, as we have seen, for this to be true, I must also have an entitlement to believe that my argument's premise entails its conclusion. I have just said that the sceptic begs the question by claiming that I do *not* have this last entitlement, but surely, now, a similar charge can be leveled against me. To argue that when I know that *P*, for some *P*, I am entitled to believe something else, that *C*, I must maintain that when I know that *P*, I am also entitled to believe yet another proposition, that *P* entails *C*. But the sceptic will of course deny this.

There is, I think, another way to attempt to undermine the sceptic's argument, taking note of the very principles we saw the sceptic to employ above.

Principle 1 For any propositions, *Q* and *C*, if you ought to doubt that *Q* and you know, or can know, that if *Q* were false then *C* would be false, then you ought to doubt that *C*.

Principle 2 For any propositions, *Q* and *C*, if you ought to doubt that *Q* and you know, or can know, that if *Q* were false then *C could* be false given only all of your evidence for *C*, then you ought to doubt that *C*.

The sceptic maintains that, in some cases, if I ought to doubt one thing, I ought to doubt another. But then, quite straightforwardly, in such a case, if it's not the case that I ought to doubt the latter thing - because I know it, say - then it's not the case that I ought to doubt the former. Now, that's precisely what the sceptic wishes to deny! And *when* is it the case that if I ought to doubt that *Q*, say, then I ought to doubt that *C*? Well, for one thing, when *Q*'s falsity entails *C*'s falsity, and I have appreciated or can appreciate this fact.

If I know that *C* entails *Q* and I ought to doubt that *Q* then I ought to doubt that *C*.

So if I know that *C* entails *Q* and I know that *C*, then I'm entitled to believe that *Q*.

To consider more closely the sort of scenario we are interested in, imagine that I have inferred, for some P and some C , from P together with P entails C to C (a modus ponens inference). The sceptic grants that I know that P and that P entails C , but argues: if you ought to doubt that (3) then you are not entitled to this inference, because if (3) were false then C could yet be false given only all of your evidence for C (by principle 2).

(1) P

(2) P entails C

(3) $(P \ \& \ (P \ \text{entails} \ C)) \ \text{entails} \ C$

Now, I could present here one of two arguments which employ one of the sceptic's own principles – Principle 1 (and, as I have suggested, I think that any sceptic would endorse *both* principles, and even would *need* to in order to be coherent). I could argue that since I know that P entails C , and I know that P – so it's not the case that I ought to doubt that P – I am entitled to believe that C (by contraposition of Principle 1). Or, perhaps in addition, I could accept the sceptic's claim - if I ought to doubt that (3) then, by Principle 2, I am not entitled to the inference – and argue that it's just not the case that I ought to doubt that (3). I *am* entitled to believe that (3), and by contraposition of *Principle 1* again, because I know that if (3) were false then P entails C would be false²³; but I know that P entails C is *true*, as the sceptic has already granted.

I do think we have accomplished something here by showing that when the sceptic presents one of her standard arguments she maintains, as we do, that there is a certain relationship which holds between some two beliefs I might have – if I am not entitled to the first of them, say, then I am not entitled to the second. This seems in itself to refute the sceptic's argument, for it follows easily that for those two beliefs, if I *am* entitled to the second of them, then I *am* entitled to the first. But what if the sceptic denies this? Well, her story then would surely lose coherence. What if whenever I claim that I have an entitlement for some belief because of some other knowledge I have, which I know to be

²³ For any propositions, P , Q , and C , if P entails C then $(P \ \& \ Q)$ entails C .

conclusive evidence for the belief, the sceptic will not in that particular case endorse the above principle – Principle 1? Well, then it would seem that she is inconsistent in her application of that principle, choosing to employ it only when it suits her. But even if we can show that she does this, or that her position or argument is incoherent, or self-undermining, or something else not to our liking, I'm not too sure that this would be all that satisfying.

I have not worked out to my satisfaction how far we could get by the consideration of such arguments as those given above in opposition to the sceptic. But I suspect that we have the following difficulties in this regard.²⁴ It is well known that a sceptic undermines, perhaps in more than one way, her own arguments and the sceptical conclusions she comes to precisely *by* the presentation of those arguments. But if she abides by the same rules of the game as we do in presenting those arguments, she surely succeeds at calling into question the legitimacy of those 'rules'- by turning them back upon themselves, if you like, to wreak havoc. Even our sceptic above who seemingly 'chooses' when to adhere to and when not to adhere to the principle we have mentioned, seems to have good 'reason' for not sticking to it when she does not - she plays by *our own* rules, even then. For when I claim that, by Principle 1, I am entitled to believe that *C* because I know that *P* and that *P* entails *C*, the sceptic employs Principle 2, which I *also* endorse, to argue that I have *not* the entitlement I claim to. She refuses to grant me an entitlement to believe that *C* until she has checked that I am also entitled to believe that (3), because, *as we all agree*, if I am not entitled to that latter belief in such a case, then I am not entitled to the former.

We need, I think, to step back from this to and fro with the sceptic, attempting to get her to grant me an entitlement for some belief of mine, or to show her up in some way, and rather provide an account which is satisfying to *ourselves*; an account of how and why it is that for some values of *P* and *C*, believing that *P* is sufficient for being entitled to believe that *P* entails *C*, and of why, for those values of *P* and *C*, *P entails C* is in fact

²⁴ See Wright (2002, p. 54) for a related discussion, in that paper in connection with the question of whether it would be any good to argue that a belief in the validity of modus ponens is 'default-justified' - because 'anything we'd recognize as a case against it would have implicitly to rely on it...'

true.²⁵ Then, the account must explain why in such a case, being entitled to *believe* that *P* entails *C*, together with *P* entails *C* being true, is sufficient for being entitled to *infer* *C* from *P*. If we can do all of that, then we have an explanation of how it can be that knowing that *P* and inferring *C* can be sufficient for knowing that *C*.

But, as I have said above, an account of inference must also be an account of the *rationality* of inference. To this end, part of the burden of the next chapter will be to explain that anyone who infers *C* from *P* does in fact *believe* that *P* entails *C* – she infers from *P* to *C* *because* she believes that *P* entails *C*, and, indeed, her *reason* for inferring is that *P* entails *C* (as she takes it to). If we are unwilling to attribute the belief that *P* entails *C* to someone who infers *C* from *P*, we leave her act of inference as unexplained. This means that I must explain how seeing *P* entails *C* as her *reason* for inferring *C* from *P*, and therefore as, in some sense, one of her reasons for *believing* that *C*, does not lead us into regress. Also, I must argue that for a simple modus ponens inference from *P* to *C*, say, anyone who knows that *P* is willing to infer *C*²⁶, and that she therefore believes that *P* entails *C*. I'll do this by a conceptual role or inferentialist semantics, also explaining by such a theory of meaning why her belief would be *true* in such a case. I shall explain her *entitlement* to this belief in chapter 6, and more fully in chapter 7, by making use of Boghossian's notion of 'epistemic blamelessness'. Ultimately, we'll have an account of how a reasoner in fact *knows* that *P* entails *C* in such a case, and comes to know of the truth of *C* as a result of her inference of it from *P*, and from *P* alone.

This is not to rescind the views expressed in the present chapter, but to take cognisance of the insights we have been allowed here, and to make progress. Progress is made in at least two respects. Because my account is one according to which anyone who is entitled to an inference does *in fact* know that inference to be valid, first, it is more internalist in nature; and, second, it is an account of the *rationality* of inference. When a subject carries out a simple modus ponens inference, she *believes* that the inference is valid, and this fact explains the *act* of inference as a rational act. Furthermore, she is *entitled* to that belief, as

²⁵ So that the two constraints of chapter 3 are in such a case met.

²⁶ We will see an objection to this view in chapter 7.

she must be if she is to be entitled to the inference, and in fact *knows* that the inference is valid and brings this knowledge to bear on the *act* of inference, which explains why she is *entitled* to the inference, and explains it in a way which is sufficiently internalist to make room for the possibility of *certain* deductive inferential knowledge.

Chapter 6

Putting Theory into Practice

I Introduction

I think that the account of the last chapter is (almost) correct as far as it goes, but that it does not go far enough. The idea there was to explain what it is to be entitled to an inference without reference to the belief that the inference is valid which a reasoner (I will argue) does indeed have when she infers. The approach was appealing because it supposedly allowed us to circumvent the difficulty of having to explain how knowledge that *P* entails *C* might be *deployed* by a reasoner who infers *C* from *P*, in so far as we were only concerned with the question of her *epistemically entitlement* to the inference. The idea was that we attribute the belief that *P* entails *C* to such a reasoner in an attempt to explain her act of inference as *rational*, not as part of an explanation of why she is *entitled* to that act of inference.

But now, firstly, even if that idea was correct – if we do not need to make reference to a belief that *P* entails *C* in order to explain how someone is *entitled* to infer *C* from *P* – an account of inference *must* also be one which allows us to see inferences as *rational* acts. And if we do not attribute the belief that *P* entails *C* to someone who infers *C* from *P* we leave that act of inference as unexplained. Then, secondly, I do not think that the previous chapter provides a quite satisfactory account of what it is to be *entitled* to an inference. For there are elements of the view developed there which are at least quite unacceptable from an internalist point of view. Even though *P* must entail *C* if I am to be entitled to infer *C* from *P*, it is not important that I actually know that fact to obtain (and this would at least be unacceptable to a *perspectival* internalist¹). And though I must have an entitlement to believe that *P* entails *C*², we have not seen how my having of that entitlement is of any *use* to me when I infer *C* from *P* – how it has a bearing on my act of

¹ See chapter 3.

² I leave aside for now, until the next chapter, the question of whether the *way* in which a thinker is required to have the entitlement to believe that *P* entails *C* is acceptable to an internalist. I have already *raised* this issue in chapter 5.

inference – for it's not critical that I even *have* the belief that *P* entails *C*, let alone justify it. Then, a related point, though an internalist might well think that a subject who carries out an inference with entitlement *makes use* of knowledge that her inference is valid in some way when she carries out that inference, in chapter 5 we did not discuss how knowledge of the validity of her inference could be deployed in a way that would enable us to see her act of inference as justified, because, again, we did not think it to be even important that she should *have* that knowledge, only important that she should have the *entitlement* for the given belief.

I have already indicated what I wish to accomplish in this chapter. The full picture, to be refined in some of its details in chapter 7, will be one according to which anyone who carries out an inference with entitlement in fact has knowledge that the inference is valid, and has this knowledge deployed³ as her *reason* for performing that *act* of inference. (Seeing this *propositional* knowledge to be deployed as her reason for performing the *act* allows us to deal with the problem of regress, by a preservation of the tie between *propositional* and *practical* knowledge.) Furthermore, the view will be that anyone who believes that *P* is willing to infer *C* as long as *P* and *C* are obviously related by modus ponens, say; she thus *believes* that *P* entails *C* and, I'll argue, this belief does indeed in such a case amount to knowledge.⁴ She is therefore entitled to the inference. Maintaining that anyone who carries out an inference in fact believes that inference to be valid allows us to explain an act of inference as *rational*, and also enables us to explain in a more satisfying way that when the given belief is a case of knowledge, the act of inference is one to which she is *entitled*. The fact that I believe something can both explain why I *perform* some act, and, at least when I am entitled to the belief, explain why I am *entitled* to that act.

In section II, to get a sense of what our approach should be, I shall consider our intuitive reaction to the tortoise who refuses to accept some proposition, no matter what else he accepts, or who argues that no matter what I know, I am not entitled to believe that some

³ It would be better to say that the proposition of whose truth she knows is her *reason* for performing the act; I refrain from putting things in that way for ease of expression.

⁴ In chapter 7 I'll deliver a fuller and more refined treatment of these issues.

other distinct proposition is true. I shall intimate that we would not allow ourselves to be led into the tortoise's trap of regress if we were indeed to remain *faithful* to this intuition – which comes down to the intuition that propositional and practical knowledge are intimately connected - so that we should not think that we must in fact *withdraw* from it in order to avoid regress.

II Pointing the Way Forward

Would we really attribute the belief that P entails C to someone who accepted ' P 's truth is a guarantee of C 's,' if she did indeed believe that P but remained doubtful of the truth of C ? Surely not. But why, then, are we so generous to the tortoise? Isn't his claim that he believes that P and that P entails C , while doubting that C , incoherent?

If Carroll's tortoise says that he will never be willing to believe some distinct proposition, no matter what else he accepts, we naturally think that he is saying that no matter what he believes he will never be convinced that its truth *guarantees* the truth of a distinct proposition. That is, we think that if the tortoise is unwilling to believe that C when he is convinced that P , then he is *not* convinced that P entails C ; the tortoise cannot coherently say that he believes that P , but doubts that C , though he is also convinced that P 's truth *guarantees* C 's. That's why we move immediately to try to convince him, after he has accepted a given proposition, P , that P does in fact entail C . But the tortoise then readily 'accepts' that P entails C - he claims to believe it - and still, contrary to our expectations, remains doubtful of the truth of C . Our thought above was that this is incoherent. If the tortoise is doubtful of the truth of C though he is convinced of the truth of P , then he is not, no matter what he says, convinced that P 's truth guarantees C 's.

The tortoise tries, however, to explain his position as follows. He claims now to believe that $P \& (P \text{ entails } C)$ - he claims that he is *indeed* convinced both that P and that P entails C - it is just that he is not (yet) convinced that these two propositions jointly entail C (that is, that (3)).

- (1) P
- (2) P entails C
- (3) $(P \ \& \ (P \text{ entails } C))$ entails C

And it is indeed true that he is not convinced that (3) if he is unwilling to infer C from $P \& (P \text{ entails } C)$. This accords with what we have just said: If the tortoise is convinced that his premise is true *and* that his premise entails his conclusion, then he *is* willing to infer from premise to conclusion. But since the tortoise is *not* willing to infer, though he's convinced of the truth of his premise, he's not convinced that his premise entails his conclusion, C (that (3)).

The tortoise's position begins to seem coherent, perhaps even *reasonable*. He says that he still doubts that (3), and we're agreed that anyone who doubts that his premise entails his conclusion will not go ahead and infer the latter from the former.⁵ Now we, along with Achilles, are inclined to move again to ask the tortoise to accept that (3), that is, that his (*new*) premise entails his conclusion. Once more, we think that once he does *that* he will (and even *must*) accept the truth of C . For our intuition is, recall, that the tortoise cannot believe that his premise is true, and yet doubt that his conclusion is true while at the same time being convinced that premise entails conclusion.

But to remain faithful to this intuition we should of course not take this further step on the path of regress. We should say that as long as the tortoise doubts the truth of C though he believes that P , he does *not* believe that P entails C in the first place. This means that as long as he genuinely believes that P and that P entails C , he has no doubt about the truth of C ; so as long as he believes that $P \ \& \ (P \text{ entails } C)$ he *is* willing to infer C . But what does this mean for us? What are we committing ourselves to in remaining faithful to this intuition? Since, as I've just rehearsed, if someone is willing to infer from $P \& (P \text{ entails } C)$ to C then she has no doubt about the truth of

- (3) $(P \ \& \ (P \text{ entails } C))$ entails C ,

⁵ See chapter 4.

and I've said that believing that $P \& (P \text{ entails } C)$ entails being willing to carry out that inference, what we're committed to is that, for any P and C ,

If I believe that P and that P entails C ⁶, then I have no doubt about the truth of (3).

If I am convinced of the truth of both P and $P \text{ entails } C$, not only is it the case that I have no doubt about the truth of C , it's also the case that I have no doubt about the truth of (3). This gives us some idea of how it can be that having knowledge of one thing can entail having the knowledge that it entails something else. I said earlier that I would explain how this is possible. I *need* to explain this because I think that some of the time knowing of the truth of a premise of an argument and inferring the conclusion is sufficient for knowing of the truth of that conclusion. But also I think that I must *always* know that an argument is valid in order to be entitled to employ it, so I must maintain that some of the time knowing of the truth of the premise of an argument (perhaps together with inferring its conclusion) is sufficient for knowing that its premise *entails* its conclusion.

Now I've said that, for any P and C , anyone who is convinced of the truth of $P \& (P \text{ entails } C)$ has no doubt that this proposition entails C . If it is also true that such a person in fact *believes* that this proposition entails C , we might argue that she is automatically *entitled* to that belief. A way to flesh out the details of this is to argue, as I do below and in the next chapter, that anyone who has the concept of entailment, so anyone who believes that P and that P entails C , is convinced that (3). Then, that such a person cannot be *blamed* for being convinced that (3) on the basis that she, say, has not justified it or has no reflectively appreciable warrant for believing it. And she cannot on such a basis be *blamed* for being convinced that (3) because she cannot have such an 'antecedent warrant' for the belief that (3) *until* she has the concept of entailment (that concept is one which is ingredient in the proposition, (3)) - and she has to be convinced that (3) in the first place in order to have that concept.

⁶ Or that $P \& (P \text{ entails } C)$

When faced with the sceptic's argument against deductive inferential knowledge, the conclusion of which is that no matter what I know, I always ought yet to doubt the truth of some distinct proposition, we naturally think that the sceptic is claiming that I can never know, for some P and C , that P and that P entails C . How can the sceptic coherently claim that I ought to doubt the truth of some C when I know that P , for some P , without maintaining that I ought to doubt, and therefore do not know, that P 's truth guarantees C 's? He leads us to think that he *can* by (purportedly) granting that I know both that P and that P entails C , while arguing that I am still not entitled to be convinced of the truth of C , because I need to also know that (3). But there is a way to accept that I need to also know that (3) here while remaining faithful to our intuition - that to claim that I ought to doubt that C though I know that P is to *deny* that I know that P entails C . We must maintain that when I know that P and that P entails C I *do* know that (3), and we have some indication now of how we *can* maintain this.

III Boghossian's Conceptual Role Semantics

I have argued that, for any P and C , it must at least not be the case that I ought to doubt that P entails C if I am to be entitled to infer C from P . In the previous chapter it was suggested that when I defend my entitlement to an inference from P to C by saying that I *know* that P entails C , what I mean to show by this is really only that I am *entitled* to believe that P entails C - what matters as far as my *epistemic entitlement* to the inference goes is that I should have an *entitlement* for the belief that P entails C .

I have also argued in chapter 3, that P must *in fact* entail C if I am to be entitled to infer the latter proposition from the former. If it is furthermore true that when I infer C from P (or am willing to do so) I do in fact believe that P entails C , then when I am entitled to such an inference, I have the belief that P entails C with entitlement, and P *does* entail C ; so we might naturally think that in any case in which I am entitled to an inference, I in fact *know* that the inference is valid.

Now, I have said that I would explain how it is that, when an argument is obviously of the form of modus ponens, say, as long as I know of the truth of its premises, I do as a matter of fact have the all important entitlement to believe that the argument is valid. So, how will I explain this? Why and how is knowing of the truth of the premise(s), in such a case, sufficient for being entitled to believe that the premise entails the conclusion? In the previous section I already provided some idea of how I will explain this. But we might ask whether I have a *reflectively appreciable* warrant for the given belief in such a case, as we may wish it to be, given our internalist commitments. If so, in what form is this warrant accessible to me? Or, is it for some other reason not the case that I ought to doubt that my inference is valid in such a case?

I shall address these questions with the assistance of some of the work of Paul Boghossian. He too argues, along the lines of our idea in chapter 5, that one can be entitled to an inference without actually *knowing* or even being able to know of the soundness of the principle of the inference, or of the validity of the inference. He argues that if we deny this then, firstly, we can have no coherent general view of how such a belief could be justified, and, secondly, Carrollian regress is inescapable. Boghossian contends, therefore, that it is enough that I am *not* epistemically irresponsible in *carrying out* an inference, and that we can see how I can be epistemically responsible (or at least not *irresponsible*), even though I can't know that the inference is valid, with the help of a conceptual role semantics. The upshot of this is a viable defence of rule-circular argument, as I'll explain and discuss in chapter 7.

As will become clear, my account diverges from Boghossian's in that I think that anyone who is entitled to an inference of *C* from *P* *does* know that *P* entails *C*, and we *can* explain how she is entitled to her belief that *P* entails *C* in the case in which her inference is of the form of modus ponens, say, by Boghossian's own notion of epistemic blamelessness. Then, I do *not* think that we cannot avoid regress if we think that such a subject infers *because* she believes that *P* entails *C*, and is entitled to that inference because she *knows* that *P* entails *C*. I agree with Boghossian, however, that a thinker can be entitled to, say, an inference of the form of modus ponens without knowing that *modus*

ponens is truth-preserving – that *all* inferences of that form are valid. This, as we'll see in chapter 7, will allow me to endorse Boghossian's own defence of rule-circular justifications.

In the two sub-sections to follow, respectively, I shall present Boghossian's arguments contending that if we think that a reasoner knows that *P* entails *C* whenever she is entitled to infer some *C* from some *P*, this knowledge can in the general case be neither, firstly, inferential nor, secondly, non-inferential.

Following that, I shall explain Boghossian's conceptual roles semantics and his reasons for thinking that a subject is epistemically *blameless* for *carrying out* an inference (not for *believing* that the inference is valid) when that inference is of the form of, say, modus ponens. In section III-iv I shall urge that this is indeed the right sort of move to make – it takes us a step closer to seeing propositional and practical knowledge as intimately connected – but I'll argue that Boghossian does not carry the project through to the extent that he ought to. Once we have appreciated fully that a close and inextricable connection exists between propositional and practical knowledge, we will not recoil from the thought that anyone who is entitled to an *act* of inference (because, say, she is blameless in carrying out the inference) has the propositional knowledge *that* the inference is valid (because, say, she is blameless in having the *belief* that it is). For it is quite simply *not* the case that '[h]owever many strata of knowledge-that are postulated, the same crux always recurs that a fool might have all that knowledge without knowing how to perform'.⁷ We will be able to see why this is not the case when we have a better understanding of what it is to understand and believe something. By keeping in mind that propositional knowledge cannot be separated from practical knowledge, I shall explain how we can see knowledge *that* an inference is valid as being brought to bear on the relevant *act* of inference without confusedly allowing ourselves to be steered into Carrollian regress.

⁷ Ryle (1946, p. 8). See chapter 4 (IV

Telling Myths to get into Regress?)

III-i Simple Inferential Internalism

In arguing against an externalist account of inferential justification, Boghossian agrees that such examples as Bonjour's Norman and the Fermat-example⁸ are decisive against it. He suggests that the intuition underlying our reluctance to see believers in such circumstances as *justified* believers is the intuition that '... someone [is not] justified in holding a given belief if they are being epistemically *irresponsible* in holding that belief. Being justified is, at least in part, a matter of being epistemically blameless.'⁹ He goes on to consider and reject the most natural alternative to such an externalist view – to be justified in believing some proposition, *P*, a subject 'must have to hand a *reflectively accessible* warrant for the proposition that *P*'.¹⁰ This is at least one way to ensure that justification does not come apart from epistemic blamelessness. So we have what Boghossian calls Simple Inferential Internalism:

A deductive inference performed by *S* is warrant-transferring just in case (a) *S* is justified in believing its premises (b) *S*'s justification for believing its premises is suitably independent of his justification for believing the conclusion, and (c) *S* is able to know by reflection alone that his premises provide him with a good reason for believing the conclusion.¹¹

Now, in the case of deductive inference, when I infer *C* from *P*, condition (c) above cannot be satisfied unless I know or am able to know that *P* entails *C*. So, Boghossian asks, how might I be in a position to have that knowledge? The knowledge would have to be either inferential or non-inferential, and 'an inferential route to knowledge of the validity would be completely useless. The very sort of reasoning whose justification is at issue would have been presupposed.'¹²

⁸ See chapter 3 and chapter 4 (II-i The Sceptic's Argument and the Most Natural Reply), respectively.

⁹ Boghossian (2003, p. 228)

¹⁰ Boghossian (2003, p. 228)

¹¹ Boghossian (2003, p. 229)

¹² Boghossian (2003, p. 230)

So if Simple Inferential Internalism is right I must be able to have *non-inferential* knowledge of the validity of an argument if I am to be entitled to infer according to it. And non-inferential knowledge, as Boghossian points out, is justified either by observation alone or by nothing at all.¹³ In ‘How are objective epistemic reasons possible’, Boghossian considers and rejects some arguments in favour of the latter option¹⁴, and we’ll see that my account does indeed proceed along the lines of this approach.¹⁵ For now, however, let’s consider his reasons for dismissing the idea that a quasi-perceptual capacity of ‘rational insight’ allows us to know of the truth of some propositions, such as some logically necessary truths.

III-ii Rational Insight

The idea that by an act of rational insight a reasoner can appreciate that a given proposition is true has been historically prevalent, and also has some contemporary supporters, such as Kyburg and Bonjour:

I think that in some sense...our justification of deductive rules must ultimately rest, in part, on an element of deductive intuition: we *see* that *modus ponens* is truth-preserving—that is simply the same as to reflect on it and fail to see how it can lead us astray.¹⁶

When I carefully and reflectively consider the ... inference ... in question, I am able simply to see or grasp or apprehend ... that the conclusion of the inference must be true if the premises are true. Such a rational insight, as I have chosen to call it, does not seem to depend on any particular sort of criterion or any further discursive or ratiocinative process, but instead is direct and immediate.¹⁷

The main difficulty with the idea, which Boghossian thinks to be almost decisive against it, is that it is mysterious what such a faculty of ‘rational intuition’ could consist in and

¹³ Boghossian (2003, p. 230)

¹⁴ Boghossian (2002, p. 21 – 23)

¹⁵ See also Wright (2002, section 4)

¹⁶ Kyburg (1965), cited by Boghossian (2002, p. 20)

¹⁷ Bonjour (1998, p. 106-107), cited by Boghossian (2003, p. 230)

how it could function.¹⁸ Crispin Wright has the following to say ‘in support of his impatience’:

...it may be reflected that someone who is inclined to believe in the objectivity of some contested region of thought *always* has the option of postulating a special capacity of direct sensitivity to the relevant putative region of special fact. This move does nothing to support objectivism unless the claim that we have such a faculty is rendered appraisable, and that demands some sort of account of how the proposed faculty—in the present case, a quasi-perceptual faculty of rational intuition—goes to work on the relevant subject matter and is indeed conducive to beliefs which keep track of it. While that account is missing, to invoke such a faculty, and the attendant conception of the range of facts which constitute its special province, is simply to pay ourselves an empty compliment. But in the present case no one seems to know how to deliver such an account.¹⁹

Boghossian does not, of course, want to deny ‘that we can *think about* ... logical properties and relations, and ... reason our way to general conclusions about them.’ But he finds highly questionable the claim that we ‘have some sort of *non-discursive, non-ratiocinative, insight* into their natures, an insight that would disclose immediately, and without the help of any reasoning whatsoever, that all instances of [modus ponens] are truth-preserving.’²⁰

Note, however, that Boghossian is arguing here that a faculty of rational insight surely cannot allow us to see that ‘*all* instances of modus ponens are truth preserving’. (My emphasis.) And it seems that a good argument for this can indeed be made. ‘One might (quasi-) perceive that a *particular* object of attention is thus-and-so—but how can one (quasi-) perceive that *all* objects of a certain kind are thus-and-so?’²¹ But now, part of what Simple Inferential Internalism requires of a reasoner, we said, was that she must be

¹⁸ See Boghossian (2003, p. 230-232) for his arguments against Bonjour’s account of the workings of the supposed faculty of rational insight.

¹⁹ Wright (2002, p. 52)

²⁰ Boghossian (2003, p. 231)

²¹ Wright (2002, p. 53)

able to know that her *particular* inference is valid if she is to be entitled to that inference. Later, Boghossian writes that

For obvious reasons, it's not plausible to think of [a] capacity for rational insight as operating on individual inferences one by one, generating for each of them the insight that if its premises are true, then so is its conclusion. Rather, we suppose that rational insight equips the thinker to arrive at the wholly general insight that MPP is valid, that is that:

For all p, q: Necessarily: If both p and 'p→q', then q.²²

I must admit that I'm not convinced that there are such obvious reasons at all²³, but I do accept the basic complaint as expressed by Wright against a rational insight account of how a reasoner knows that a basic inference is valid when she carries it out. As Haack expresses it, 'The claim that one can just *see* that the premisses justify the conclusion is implausible in the extreme in view of the fact that people can and do disagree about which arguments are valid.'²⁴ I have also urged that, while we can hold that a child who doesn't have the ability to work with logical *vocabulary* has all the same the *belief* and perhaps knowledge that an argument is valid when she carries it out, it is implausible to think that she is able sometimes to *see* that the argument is valid.²⁵

Having rejected, therefore, the option of appealing to a rational insight theory to assist us in dealing with the problems at hand, let's consider Boghossian's alternative proposal.

III-iii Boghossian's Way Forward — Conceptual Role Semantics and Epistemic Blamelessness

Because Boghossian rejects the idea of rational insight and thinks that this is the most plausible way to argue that whenever a subject infers with entitlement she knows or is

²² Boghossian (2003, p. 232)

²³ See Wright (2002), especially p. 52 and sections 11-13 where he explains how a rational insight theorist (which Wright is not) might defend herself against the arguments of Boghossian.

²⁴ Haack (1976, p. 118)

²⁵ See chapter 5 (

II My Beliefs are Not the Matter?).

able to know that her inference is valid, and because he furthermore finds the Carrollian regress argument convincing²⁶, he concludes that someone can be entitled to *carry out* an inference without being able to know that it is valid. But he also wants to avoid a purely externalist account of inferential justification, so his suggestion is that our internalist intuitions are satisfiable by requiring only that a reasoner must be *epistemically blameless* in carrying out an inference if her inference is to be warrant-transferring.²⁷ And an inferer can be epistemically blameless when she infers from some *P* to some *C*, even if she is unable to know that *P* entails *C*. Though in the most common counter-examples to an externalist account of justification, such as Bonjour's clairvoyant and Boghossian's Fermat example, the subjects involved have no reflectively appreciable warrant for coming to believe what they do in the way that they do and are subjects who are, seemingly *for that reason*, epistemically blameworthy in holding those beliefs, Boghossian argues that someone might lack such a warrant and yet be epistemically blameless in holding her given belief. In the case of someone who carries out an inference from *P* to *C* though she has no reflectively appreciable warrant for that inference, Boghossian contends, the subject is yet epistemically blameless as long as inferring according to the rule which licenses that inference is 'a precondition for having one of the concepts ingredient in [either *P* or *C*, or both]'

The thought is this. Suppose it's true that my taking *A* to be a warrant for believing *B* is constitutive of my being able to have *B*-thoughts (or *A*-thoughts, or both, it doesn't matter) in the first place. Then doesn't it follow that I could not have been epistemically blameworthy in taking *A* to be a reason for believing *B*, even in the absence of any reason for taking *A* to be a reason for believing *B*? For how could I have had *antecedent* information to the effect that *A* is a good reason for believing *B*, if I could not so much have had a *B*-thought without taking *A* to be a reason for believing *B* in the first place? If inferring from *A* to *B* is required, if I am to have the ingredient propositions, then it looks as though so inferring cannot be held against me, even if the inference is blind.²⁸

²⁶ Boghossian (2003, 232-236); see also Boghossian (2002) and Boghossian (1999).

²⁷ We will see, in chapter 7, how Wright puts pressure on this view, to which my own account is closely related.

²⁸ Boghossian (2003, p. 240)

This is Boghossian's way of beginning to fill in the details of an account which takes as 'one of Wittgenstein's fundamental insights' the thought that 'giving grounds ... comes to an end', and so, in the case of inference in particular, an inference can be 'blind' yet justified.

When I obey a rule, I do not choose.

I obey the rule *blindly*.²⁹

Giving grounds ... comes to an end – but the end is not certain propositions' striking us immediately as true, i.e., it is not a kind of *seeing* on our part; it is our *acting*, which lies at the bottom of the language game.³⁰

Boghossian thinks that a conceptual role semantics puts us 'in a position to mount an explanation of the blameless blindness of [modus ponens].'³¹

The idea that the content of our concepts is determined by introduction and elimination rules for (or circumstances for and consequences of) their correct use is now quite widely accepted and has been given, especially in the case of the logical constants, considerable support. On this view the meaning of the logical constant *and*, for example, is determined by the introduction rule

from A, B
infer A and B,

and the elimination rules

from A and B ; from A and B
infer A infer B,

²⁹ Wittgenstein (1953, remark 219); cited by Boghossian (2003, p. 237).

³⁰ Wittgenstein (*On Certainty*); cited by Boghossian (2003, p. 237).

³¹ Boghossian (2003, p. 241)

and someone has the concept of conjunction only if they are disposed to infer according to these rules.

Similarly, the thought would be, a subject has the concept of entailment only if she is disposed to infer, for any propositions P and C , from P together with P entails C , to C .³² And if I must be disposed to carry out such inferences in order to have the concept of entailment in the first place, which is plausible, then I cannot be held blameworthy for doing so in a particular case just because I don't know, or lack a reflectively appreciable warrant to believe, or am *unable* to know that P entails C . *Of course* I don't know and am unable to know that P entails C if I don't have the concept of entailment, for that concept is a component concept of P entails C , but I must first be disposed to carry out the inference I have just carried out if I am to be in possession of that concept.

III-iv Putting Theory into Practice

Boghossian has rightly, I think, seen the way forward to require a rejection of a sharp distinction between propositional and practical knowledge, and this is indeed quite widely taken to be the correct diagnosis of the problem. In 'Inference, Belief, and Understanding'³³, Stroud shows how we meet with severe difficulty if we think of belief in this way - of having a belief as analogous to having an entry of the relevant proposition in a notebook, as Stroud puts it. As we have seen, he approaches the question, "Just because I believe this, why *must* I believe that?" posed to us by the tortoise, by asking what it is for one belief to be *based on* another. If we neglect the fact that there is a practical aspect of belief, then the fact that I believe that P and that C , for some P and C which are related (even obviously) by *entailment*, would not seem to imply that the latter belief is *based on* the former. We will think that I might not have seen any connection between P and C , so I might believe C for some other reason than P entirely.

³² Provided she can understand these three propositions – I might have the concept of entailment but not some other concept ingredient in either P or C , say; or P entails C could be too long or complex for me to understand.

³³ Stroud (1979)

So this line of thought will easily lead us to make the mistake of thinking that to have based my belief that *C* on my belief that *P*, I must in addition have ‘seen’ or at least have the belief that, say, *P* entails *C*³⁴, as another notebook entry. Then we are on the road to regress because, again, there is still the question of whether or not I have seen that *C* is entailed by *P* together with *P* entails *C*, or whether I perhaps believe that *C* for some other reason not connected with *P* and *P* entails *C* at all.³⁵

But ‘believing something is more than having something written in a notebook.’³⁶ What we have forgotten, though we would not so easily forget it in everyday life where we treat someone who *behaves* inappropriately as disqualified from having a given belief, is that ‘belief requires understanding and understanding requires seeing connections and drawing conclusions in appropriate circumstances.’³⁷

This, as Stroud explains, might ‘be thought to provide a way around Lewis Carroll’s regress and to explain the source of the ‘must’ in “He believes this so he must believe that.”’³⁸ ‘A person’s understanding and believing things must ultimately be seen as a certain complicated disposition or competence or practical capacity.’³⁹

Thus, though someone such as the tortoise might genuinely *believe* that *P* and *accept* ‘*P* entails *C*’, her subsequent non-acceptance of *C* would show that she does not in fact *believe* that *P* entails *C*, that she hasn’t *understood* what she has accepted in accepting ‘*P* entails *C*’. Anyone who is convinced of the truth of *P* and also convinced that the truth of *P* *guarantees* the truth of *C*, quite simply *will* move to an acceptance of *C* if asked to when reflecting on, as she takes it, the truth of those two propositions, *P* and *P* entails *C*. Anyone who is convinced of the truth of *P*, accepts ‘*P* entails *C*’, but remains doubtful of the truth of *C* because she wishes first to check for the truth of (3), even though such

³⁴ Or that some other relationship holds between the propositions *P* and *C*, such as that *P* is good reason to believe *C*.

³⁵ See Stroud (1979, p. 185-187) for Stroud’s excellent discussion of this.

³⁶ Stroud (1979, p. 191)

³⁷ Stroud (1979, p. 189-90)

³⁸ Stroud (1979, p. 189-90); see also Black (1970, p. 21)

³⁹ Stroud (1979, p. 194)

behaviour seems at first blush to be reasonable, just does not really believe that *P* entails *C*.⁴⁰

Brown also expresses a thought along these lines as follows:

A man knows that if *p*, then *q* if, when he knows that *p*, he is able to see that, consequently, *q*. If a man knows that *p*, but cannot see that *q*, this is just what shows him not to know that if *p*, then *q*. In other words, knowledge of the hypothetical ‘If *p*, then *q*’ is such that there cannot be a man of the kind the Tortoise said might exist, namely [one] who accepted ‘*p*’ and ‘If *p*, then *q*’ but failed to see ‘*q*’.⁴¹

We must appreciate that there is, as Brown puts it, ‘an intimate relation between knowledge of hypotheticals and ability to see conclusions’.⁴² Thus, as I would put it, *pace* Boghossian, a man knows that *P* entails *C* when he infers from *P* to *C* with entitlement, and this idea does not steer us into Carrollian regress because if a man knows or believes that *P*, but is unwilling to infer *C*, this is just what shows him *not* to believe that *P* entails *C*.

In the last section of his essay, Brown attacks the argument presented by Ryle against the idea that ‘knowing how to reason ... [is] ... analysable into the knowledge ... of some propositions.’⁴³ Providing hypothetical examples purporting to demonstrate this, as we have seen⁴⁴, Ryle contends that someone might know any number of things (‘dutifully reciting’ them, as Ryle puts it) without knowing *how* to reason with them at all, just as the tortoise (supposedly) knows that *P* and that *P* entails *C* but does not infer *C*.

But Brown argues that

⁴⁰ See Stroud (1979, p. 190)

⁴¹ Brown (1954, p. 175). Brown’s last claim here does not seem quite right for I can surely ‘accept’ something without actually believing it. If that’s right then I can indeed ‘accept’ any number of things but be unwilling to accept others, or be unable to ‘see’ them. The question is whether or not I *understand* what I am accepting in such a case – I cannot *believe* that *P* and that *P* entails *C* and yet ‘fail to see’ that *C*.

⁴² Brown (1954, p. 177)

⁴³ Ryle (1946, p. 6); cited by Brown (1954, p. 177).

⁴⁴ Chapter 4 (IV Telling Myths to get into Regress?)

... none of this will do. The case described is not possible. There is no puzzle provoked by its possibility, and no call here for a distinction between 'knowing that' and 'knowing how'. ... inability to see the conclusion is sufficient to show that the pupil does not understand and accept the hypothetical, regardless of what he 'dutifully recites'. It is doubly so [if] the premises entail the conclusion.⁴⁵

Ryle and Achilles, then, commit what is essentially the same error. Achilles is led into it by the seemingly harmless act of adding as an entry in his notebook, on the instruction of the tortoise, proposition after proposition which the tortoise 'accepts'. But he will never show *in that way* that the tortoise must accept the conclusion, which, besides, was the case if and when the tortoise believed only that *P* and that *P* entails *C* in the first place, as he at least claimed to. Ryle makes the similar mistake of thinking that 'dutifully reciting' or 'knowing by heart'⁴⁶ a proposition is sufficient for believing it, which requires *understanding* it.⁴⁷ To the contrary, as Stroud points out⁴⁸, the fact that the tortoise 'dutifully' accepts whatever Achilles asks him to accept gives us further reason (apart from the fact that he doesn't go on to accept obvious consequences) to think that he *doesn't* believe what he is accepting at all. Before one accepts something it is important that one understands it, and has some sort of reason to believe it. But the tortoise asks Achilles neither for any such reasons nor for any explanation of the meaning of what he accepts before he does so. He also does not call for an explanation of the relevance of the truth of what he is accepting to the argument with which he and Achilles are at that moment concerned. Again, 'believing something is more than having something written in a notebook,'⁴⁹ and more than a dutiful recitation of it. As Ryle himself points out, 'knowing these principles exhibits itself not in the recitation of formulae but in the execution of valid inferences and in the avoidance, detection, and correction of fallacies, etc.'⁵⁰ One of the morals of Carroll's story is, not that we should distinguish between practical and theoretical knowledge with the focus falling more firmly on the former as

⁴⁵ Brown (1954, p. 178)

⁴⁶ See quote from Ryle in chapter 4 (IV Telling Myths to get into Regress?).

⁴⁷ See Brown (1954, p. 178) for his discussion of this.

⁴⁸ Stroud (1979, p. 190)

⁴⁹ Stroud (1979, p. 191)

⁵⁰ Quoted in chapter 4 (IV Telling Myths to get into Regress?).

irreducible to the latter, but, indeed, that we should precisely *not* separate out too cleanly these two forms of knowledge. There is ‘no call ... for a distinction between ‘knowing that’ and ‘knowing how’.

Making this mistake is, as should already be evident, and as I’ll explain more carefully below, what leads us to think that we are led into regress if we attribute to someone who infers *C* from *P* the *belief* that *P* entails *C*. It’s also therefore the mistake which lures us into proposing alternative accounts of inferential justification, such as my own only partially satisfying account of the previous chapter as well as Boghossian’s account, which run counter to our intuitions and actual practices of attributing *propositional* knowledge together with an *inferential* entitlement. With the correct view of the matter it might indeed be correct to say, with Ryle, that ‘knowing how to reason ... [is not] ... analysable into the knowledge ... of some propositions’, but we should add that knowing-*that* is not analysable into knowing-*how*, for the two are intimately and inextricably connected.

The mistake Achilles makes is not so much to be willing to represent the Tortoise’s beliefs in a notebook (though he is easily led astray by doing so), but to think that by doing that *alone* he will at some point convince the tortoise to accept the conclusion. It is not confused to represent a belief someone has as an entry in a notebook, but to think that to be *all there is* to having that belief. Or, to put it another way, Achilles ought not to add, say, *P entails C* in his notebook on the basis only of the tortoise having accepted ‘*P entails C*’, for there is more to believing something than the mere acceptance of some formulation of its content. The explanation of the fact that the tortoise *must* accept the conclusion will come when we have a list of what the tortoise genuinely believes *together with* an account of what it is to understand and believe something⁵¹, specifically, that thing or those things the tortoise does indeed believe. Such an account, therefore, will have to not only support the view that understanding and believing something requires correctly drawing *some* conclusions from it, seeing *some* of the obvious consequences of

⁵¹ See Stroud (1979, p. 194) - ‘What is needed is some explanation of what his having that alleged knowledge amounts to, as opposed to what it is knowledge of.’

it; we would hope it to explain that believing some given proposition, that $P \& (P \text{ entails } C)$, say, requires accepting another *given* proposition, that C .⁵² I shall turn now to discuss one developed form that such an account might take.

IV Brandom's Semantic Inferentialism and Logical Expressivism

What we are looking for is to salvage our inferential practices, to explain why someone who knows (and therefore believes) that P and that P entails C , *must* accept that C in appropriate circumstances, and also why she would be *entitled* to do so, by an account of the *meaning* of our expressions or of the *content* of our concepts. If to have the concept of entailment one must be willing to carry out any inference of the above form just as long as one can understand its premises and conclusion, then one cannot believe the premises and yet be unwilling to accept the conclusion - believing the premises involves having the concept of entailment because it is a component of one of those premises.

We have been given ample opportunity to appreciate the crucially important role that one's relevant *practical* knowledge has to play in the question of what *propositional* knowledge one has. We have also seen argument against the idea that knowing-*how* is analysable into knowing-*that*, and we were even led to consider the possibility of dismissing such a requirement as that I must know *that* P entails C if I am to be entitled to infer C from P , suggesting that what I need is to be *entitled* to believe that P entails C , whether or not I in fact believe it.

My aim is to dismiss such arguments and suggestions – as I said above, it may indeed be the case that practical knowledge is not 'analysable' into propositional knowledge, if the idea of such analysis is to do away with all talk of know-how as if, again, to know something is merely to have it as an entry in a notebook; but that is precisely because of the intimate connection between knowing-*how* and knowing-*that*, so that we must take equal care not to make the opposite mistake of thinking that all knowing-*that* is really a

⁵² See Stroud (1979, p. 190, footnote 2); and Stroud (1979, p. 191). See chapter 7 for why these hopes might be unsatisfiable.

case of knowing-*how*. There is no error we commit in writing ‘*P entails C*’ in a list of beliefs a certain subject has when she infers *C* from *P*, or listing it under ‘Things known’⁵³ when she does so with entitlement, just as long as we understand that her genuine possession of this belief, or knowledge, requires that she understands the listed proposition and that this in turn requires that she *acts* appropriately.

For a clear view of the intimate connection existing between practical and theoretical knowledge to come in the way of an account of meaning or content, we must take a first step, as we have already done, along the lines of Wittgenstein’s thought of meaning as use.⁵⁴ To understand an expression or to have a concept, such as one which would have to be had if it is a component of a proposition believed or known, one must *use* it in appropriate ways in appropriate circumstances. For instance, if one believes a particular proposition having that concept ingredient in it, one must be willing to carry out certain inferences involving that proposition in certain circumstances. Thus, Wittgenstein writes that

The grammar of the word “knows” is evidently closely related to that of “can”, “is able to”. But *also* closely related to that of “understands”. (‘Mastery’ of a technique)⁵⁵ (My emphasis.)

I say that we have already taken a first step in this direction because we have seen some such account as this, a conceptual role semantics, advocated by Boghossian. In the same vein, Robert Brandom develops an *inferentialist* theory of meaning⁵⁶ whose approach, he writes, ‘takes the form of a linguistic pragmatism that might take as its slogan Sellars’s principle that *grasping a concept is mastering the use of a word*.’⁵⁷

On Brandom’s view, the content of a proposition is given by its inferential role, that is, by its relation to other propositions from which it follows and by which it is followed.

⁵³ See Stroud (1979, p. 192-193)

⁵⁴ Wittgenstein (1953)

⁵⁵ Wittgenstein (1953, remark 150)

⁵⁶ In Brandom (1994) and Brandom (2000)

⁵⁷ Brandom (2000, p. 6)

And, in the order of explanation, propositional content comes prior to the content of concepts which can only figure as *components* of propositions. Brandom's 'inferentialist pragmatism' 'starts with a practical distinction between good and bad inferences, understood as a distinction between appropriate and inappropriate *doings*, and goes on to understand talk about truth as talk about what is preserved by the good moves.'⁵⁸

This inferentialism makes use of Dummett's idea, as Brandom puts it, of 'any linguistic expression hav[ing] two aspects: 'the *circumstances* under which it is correctly applied, and the appropriate *consequences* of its application, utterance or use.'⁵⁹ The idea is that it is the correct circumstances for and consequences of the application of a linguistic expression which determine its meaning. And for a given (theoretical)⁶⁰ proposition, corresponding to the appropriate circumstances for and consequences of applying it (characteristically, *asserting* it) we have the propositions from which it follows, and which follow from it, respectively. So the relationship between our given proposition and these other propositions - the given proposition's inferential role - is constitutive of that proposition's content.

We have a given concept or understand a given proposition, then, when we have a mastery of its *use*, when we have a certain know-*how*; when we do indeed infer in the appropriate circumstances from and to the proposition in accordance with its content-conferring inferential role. 'Understanding a propositional content ... is here presented ... as practical mastery of a certain kind of inferentially articulated doing ...'.⁶¹ And central to Brandom's account is a distinction between what he calls, following Sellars, *materially* good inference and *formally* good or valid inference. The latter category is derivable from the former - what is and what is not a *materially* good inference is set up by our actual practices of using our expressions, and this comes first; then, a formally valid inference, for example one which is valid by virtue of its *logical* form, is one which is materially valid given any grammatical substitution in premises and/or conclusion of non-logical for

⁵⁸ Brandom (2000, p. 12)

⁵⁹ Brandom (2000, p. 62)

⁶⁰ A non-theoretical proposition, such as *that's red*, might have other sorts of appropriate circumstances (or consequences) of application, such as the presence at hand of a red object.

⁶¹ Brandom (2000, p. 63)

non-logical vocabulary, once what is our 'logical' vocabulary has been specified. Thus, the inference from 'Cape Town is to the west of Hermanus' to 'Hermanus is to the east of Cape Town' is materially good and anyone who does not endorse this inference does not understand 'west' and 'east'. And the inference from 'Cape Town is beautiful and Johannesburg is ugly' to 'Cape Town is beautiful' is materially good as well as *logically* valid given that 'and' is one of our *logical* constants and the above inference remains materially good no matter what we substitute grammatically for 'Cape Town is beautiful' and/or 'Johannesburg is ugly' (or, for that matter, for any of the sub-sentential units component of these two sentences/propositions) throughout. '[A]ccording to this way of thinking, the *formal* goodness of inferences derives from and is explained in terms of the *material* goodness of inferences, and so ought not be appealed to in explaining it.'⁶²

Now, merely sentient creatures too behave in rule-governed or regular ways, responding reliably in certain ways in like circumstances (even communally, I mean: responding in the same or similar way community-wide). And even non-sentient *things* such as thermostats and bits of iron respond reliably to, in these cases, temperature and the presence of atmospheric water vapour, respectively. What is it, on the inferentialist account, that disqualifies these merely *sentient* beings from being *sapient* – from being *concept-users*? Brandom has us consider a parrot who is trained to respond reliably to the presence of red objects by saying, "That's red," as one of the most plausible candidates for being a concept-possessor of this kind. (Note the similarity between this parrot and the tortoise who 'dutifully recites' whatever *true* propositions he is asked to accept.)

On Brandom's inferentialist view of meaning, this parrot has neither the concept *red* nor the belief that *that's red* when she 'says' that it is, since she knows nothing (or certainly *almost* nothing) of that proposition's inferential role. She does not, for instance, endorse the inference from *that's red* to *that's coloured*, nor to *that's not a prime number*⁶³, and since the content of a proposition is given by its inferential role, and one understands a

⁶² Brandom (2000, p. 53)

⁶³ Brandom (2000, p. 65)

proposition only if one has a practical mastery *of* its inferential role, she cannot be said to understand the proposition at hand.

This, then, is the mark of sapience. To be a concept-user is, as Brandom (following Sellars) puts it, to be a player in the game of giving and asking of reasons. It is not only to be a reliable responder but to be willing and able to provide reasons for saying what one says, for believing what one believes, and to accept consequences thereof. Indeed, if one is not such a game-player, one cannot be said to *say* or *believe* anything. So on Brandom's view, according to which *concept* possession is grounded in language use - 'grasping a concept is mastering the use of a word.'⁶⁴ - the situation is one in which a practical mastery of the expressions of one's language allows one to *say* all that one says (and not merely make utterances), to believe what one accepts and to have all of the concepts one employs (rather than to be a merely reliable responder), all at once. One understands the expressions and has the concepts involved in one's activities *en masse*. And so 'inferentialist semantics is resolutely *holist* ... One cannot have *any* concepts unless one has *many* concepts. For the content of each concept is articulated by its inferential relation to *other* concepts.'⁶⁵

By his account of what it is to be a concept-user, Brandom accounts for what is distinctive and important about language-games and the players of such games – they are players of the game of *giving and asking for reasons*. 'By contrast to Wittgenstein...language has a *center* ... in linguistic practices. Claiming, producing and consuming claims are not on a par with other 'games' one can play, they are what in the first place make possible talking, and therefore thinking: sapience in general.'⁶⁶

I said at the close of the previous section that we would like a theory of meaning to do two things for us. First, we need it to explain why believing one proposition requires that I am willing to accept *some* other propositions in appropriate circumstances. And the account we have seen does indeed do this for us. On the inferentialist view at hand,

⁶⁴ Brandom (2000, p. 6)

⁶⁵ Brandom (2000, p. 15)

⁶⁶ Brandom (2000, p. 14-15)

believing one proposition (which requires that we understand it and have the concepts ingredient in it) requires that we endorse the materially good inferences whose very correctnesses are constitutive of the content of that proposition, and we genuinely endorse these inferences only if we are willing to carry them out. But this means that the account also satisfies our second requirement – that it should explain why believing some specified proposition requires that I am willing to accept another *given* proposition in appropriate circumstances. I do not understand a given proposition, such as that *Cape Town is to the West of Hermanus* unless I endorse the inferences whose material goodnesses determine that proposition's content, one of which, for our given proposition, is the inference from *Cape Town is to the West of Hermanus* to *Hermanus is to the East of Cape Town*. Thus, on this view, if I am to understand the former proposition, I must be willing to infer from it to this other *particular* proposition, the latter.⁶⁷

But this inferentialist account of meaning, when coupled with Brandom's expressivist view of logic, provides even more elucidation for the case in which the tortoise purportedly believes, for some *P* and *C*, both that *P* and that *P* entails *C*, while refusing to accept *C*; or claims that though Achilles knows that *P* and that *P* entails *C* he cannot on that basis, via an inference from those two propositions, know that *C*. We have seen that Brandom thinks that 'the inferences that matter for ... [conceptual] contents in general must be conceived to include those that are in some sense *materially correct*, not just those that are *formally valid*.'⁶⁸ But what is the role that logic has to play then? It is Brandom's answer to this question which will shed further light on the question of why the tortoise *must* accept *C* if he genuinely believes that *P* and that *P* entails *C*, as has been suggested above. Brandom writes:

According to the inferentialist account of concept use, in making a claim one is implicitly endorsing a set of inferences, which articulate its conceptual content. ... Making explicit that know-how ... is putting it in the form of a claim *that* things are thus-and-so. In this

⁶⁷ In chapter 7 I shall discuss a possible objection to this view.

⁶⁸ Brandom (2000, p. 56)

case, a central expressive resource for doing that is provided by basic *logical* vocabulary.⁶⁹

The idea here is that it is the conditional (for which I use ‘entails’) which allows us to make *explicit* the commitments which are *implicit* in any of our claimings. In claiming, for instance, that Cape town is to the west of Hermanus, I implicitly endorse the inference from that proposition to the proposition that Hermanus is to the east of Cape Town, and, as I’ll argue below, I have the implicit *belief* that *Cape town is to the west of Hermanus* entails *Hermanus is to the east of Cape Town*. The conditional allows me to make *explicit* what I am *implicitly* committed to when I claim that Cape town is to the west of Hermanus, by allowing me to *say* that *Cape town is to the west of Hermanus* entails *Hermanus is to the east of Cape Town*.

To say, then, for some *P* and some *C*, that *P* entails *C*, is to explicitly endorse the inference from *P* to *C* as materially, if not necessarily formally, good; and I in fact endorse that inference only if I am indeed willing to carry it out. This suggests again that the tortoise, who accepts ‘*P* entails *C*’ but is unwilling to infer *C* from *P*, cannot actually understand what he accepts when he accepts ‘*P* entails *C*’.⁷⁰ And since to believe that *P* entails *C* is to be committed to the goodness of the inference from *P* to *C*, we surely cannot simultaneously (without confusion) attribute to someone an entitlement to that belief, as we do if we attribute *knowledge* that *P* entails *C*, while denying that she is entitled to the very inference whose goodness she is committed to by virtue of having *that* belief.

Important to note, however, is that if someone knows that *P* entails *C* then the inference she is thereby directly licensed to carry out is the inference from *P* to *C*, and not the inference from *P* together with *P* entails *C*, to *C* (though she *is* also licensed to carry out this latter inference). That means that this knowledge, that *P* entails *C*, must entitle our

⁶⁹ Brandom (2000, p. 21)

⁷⁰ Also, by what has gone before, the goodness of the inference from *P* together with *P* entails *C*, to *C*, is partly constitutive of the content of *P* entails *C*, so anyone who is not willing to carry out that inference doesn’t understand the proposition *P* entails *C*.

subject to carry out a particular inference even though *P* entails *C* is not itself one of the premises of that inference. In section V-i I shall argue that we are right to attribute the belief that *P* entails *C* to anyone who infers from *P* to *C*, and right to think that this belief must be a case of knowledge if she is to be entitled to that act of inference. In V-ii, below, I'll argue that requiring that I must know that *P* entails *C* in order to be able to have *inferential* knowledge that *C* does not amount to requiring that *P* entails *C* must be one of the *premises* of the inference via which I can gain that knowledge, so that this requirement does not lead us into regress.

V Knowledge Needed, Though not as a Premise

V-i Knowledge Needed

I have explained Boghossian's account of how we can be entitled to a given act of inference without knowing that the rule according to which that inference proceeds is sound, or that the *given* inference is valid. His account separates out practical from propositional knowledge, by distinguishing between being entitled to *use* a given rule of inference and knowing *that* the rule is sound, but worse, between being entitled to *carry out* a given inference and knowing *that* the inference is valid. He argues that I can be entitled to *carry out* an inference, to *use* the rule according to which that inference proceeds, even when I do not know *that* the rule is truth-preserving, and do not know *that* the particular inference at hand is valid. I am entitled to carry out a given inference as long as I am epistemically blameless in doing so, which I am, Boghossian explains, if being willing to carry out that inference is a precondition for having some of the concepts ingredient in one or more of the propositions involved in the inference. And one of Boghossian's main reasons for rejecting the view that I must know that an inference is valid in order to be entitled to carry it out is that he thinks that this view suffers from the problem of Carrollian regress.

But I have been urging that propositional and practical knowledge be seen as intimately connected and *not* separated out. I think that knowing *that* an argument is valid goes

hand-in-hand with being entitled to *employ* that argument. I've argued at various points that to infer *C* from *P* is to have at least the implicit belief that *P* entails *C* – we at least naturally attribute that belief to anyone whom we take to be an inferer of *C* from *P*, and we would not say that someone understood '*P* entails *C*' if she did not accept its truth while still endorsing the inference from *P* to *C*. According to our actual practices, that is, the belief that *P* entails *C* is tied to a willingness to perform the act of inference from *P* to *C*, and knowledge that *P* entails *C* to an entitlement to that inference – *belief* is seen as inextricably linked to *action*; propositional knowledge to entitlement to act. So though in the previous chapter we withdrew from talking of the belief that *P* entails *C* that someone has when she infers *C* from *P* in considering what it is for her to be *epistemically entitled* to that act of inference, I wish now to return to the idea that the fact that she believes that *P* entails *C* has a role to play in our explanation of her inference as one to which she is *entitled*, and, perhaps less contentiously, as an act which is *rational*.

The view that anyone who infers *C* from *P* believes that *P* entails *C* is widely endorsed in one form or another. We have seen that on Brandom's view I can endorse the goodness of an inference from *P* to *C* in one of at least two ways: by the act of carrying out the inference or by that of *saying* that *P* entails *C*. (Also, even if I merely assert or believe that *P*, where the goodness of the inference from *P* to *C* is partly constitutive of the content of *P*, I thereby implicitly endorse that inference as good.) His inferentialist account, therefore, provides further support for the view that if I infer *C* from *P* I have at least the *implicit* belief that *P* entails *C*.

Boghossian himself, as we have seen above, expresses inferring *C* from *P* as 'taking [P] to be a warrant for believing [C]'.⁷¹ Similarly, J. J. Thomson⁷² has argued, as Stroud puts it, 'that someone who infers *C* from *P* at least takes it for granted that *P* is reason to believe *C*'.⁷³ I think that's right and, furthermore, by a natural extension, that it is correct to say of anyone who infers *deductively* from *P* to *C* (remembering that by this I mean

⁷¹ Boghossian (2003, p. 239); See (III-iii Boghossian's Way Forward — Conceptual Role Semantics and Epistemic Blamelessness) above.

⁷² Thomson (1965, p. 296-8)

⁷³ Stroud (1979, p. 186, footnote 1)

that the inferer is *convinced* of the truth of *C* on the basis of her knowledge that *P*) that she takes it that *P* entails *C*. But it seems hardly controversial to move from saying that someone *takes* it that *P* entails *C* to saying that she *believes* implicitly that *P* entails *C*. We might reject this idea in the face of the threat of Carrollian regress, but I see that move as defeatist - a turning away from the problem by a rejection of some of our perfectly natural practices of the attribution of belief and knowledge. I think we ought to be more positive in our approach, perhaps taking more to heart Wittgenstein's contention that 'Philosophy ... leaves everything as it is'.⁷⁴

So, as Brown writes, 'a man knows that if *p*, then *q* if, when he knows that *p*, he is able to see that, consequently, *q*,'⁷⁵ and 'if the-fact-that-*p* convinces [someone] that *q*, then he thinks that if *p*, then *q* *even though he never formulates it*.'⁷⁶ (My emphasis.)

Now if, along the lines of the proposals of Boghossian and Brandom, having the concept of entailment or understanding the proposition *P* entails *C* involves endorsing certain inferences in which *entails* figures as a component concept of one or more of the propositions involved in the inference, such as the inference from *P* together with *P* entails *C*, to *C*; and if to carry out an inference (and also to be *willing* to do so) is to believe that its premises jointly entail its conclusion, then the situation we have is as follows. To believe that *P* entails *C* it is required that I am willing to infer from *P* together with *P* entails *C*, to *C*, and therefore also required that I believe that (3), that is, that (*P*&(*P* entails *C*)) entails *C*. Believing something involves not only being willing to act in certain ways, but also actually believing other things. Stroud expresses a similar thought when he says of someone to whom we attribute the belief that *P* that '[h]is acceptance of the conditional that if *P* is true then *Q* is true would be part of what we attribute to him in attributing to him at the outset a belief in *P*, "⁷⁷if understanding *P* involves seeing that certain other propositions are true. And Black writes that

⁷⁴ Wittgenstein (1953, remark 124)

⁷⁵ Brown (1954, p. 175)

⁷⁶ Brown (1954, p. 177)

⁷⁷ Stroud (1979, p. 189-90)

... to understand a proposition of this sort [what he calls a 'self-evident' proposition] is already to be in possession of all that is relevant to its truth. Thomas Reid said self-evident propositions 'are no sooner understood than believed', but the correct formula is, rather, that understanding a self-evident proposition necessarily includes *knowing* the proposition to be true ... if we have any doubt we could not have understood what we thought we were doubting. There can be no suspension of belief about a self-evident proposition.⁷⁸

Thus, the close relationship between belief and practice, specifically, *inferential* practice, is maintained, whereas it is partially *lost* if we think, with Boghossian, that understanding something requires *acting* in certain ways but not also *believing* those things which we take to be true by those very acts. Boghossian attempts, as we did in chapter 5, to explain entitlement to an *act* of inference in the basic case, where the rule according to which the inference proceeds is meaning-constituting, while prescinding from talk about the *belief* that the inference is valid which a reasoner has when she carries out an inference. He has provided an explanation of why the reasoner is epistemically blameless in carrying out the inference in such a case. But if it is true that this reasoner has the belief that the inference is valid, then, by the same token, she is epistemically blameless in having that *belief*. For she must have this belief in order to have the concept of entailment in the first place. Thus we explain how she is *non-inferentially* entitled to the belief without recourse to the suspect notion of rational insight (and we might even say that the belief is justified by 'nothing at all'⁷⁹). Furthermore, an inference which proceeds according to a meaning-constituting rule is always valid because it is the very goodness of that inference which partially determines the contents of the concepts involved in the inference. So our subject at hand has the *true* belief that the inference is valid with entitlement – she *knows* that the inference is valid.⁸⁰ And this view of the matter will allow us to explain in a more satisfying way why she is *entitled* to her act of inference, and explain it as a rational act. I'll argue that we can do this without being led into Carrollian regress if we have the

⁷⁸ Black (1970, p. 18); the Reid (1986) quote comes from Woozley, ed. (1986, p.358).

⁷⁹ See above - (III-i Simple Inferential Internalism)

⁸⁰ See chapter 7 for more detailed discussion.

correct view of how *propositional* knowledge of validity can be brought to bear on an *act* of inference.⁸¹

I think that anyone who infers C from P believes that P entails C and that she is entitled to that inference *only if* she in fact *knows* that P entails C . At least, she must have an entitlement for that belief and P must also *in fact* entail C , as I've argued in previous chapters.⁸² I also think that she is entitled to the inference of C from P *if* she knows that P entails C . But now the objector complains: "But surely if she knows that P entails C and that fact is relevant to whether or not she comes to know that C as the result of an inference, then P entails C must be a premise of her argument." Or, she might say: "Surely if our inferer checks and justifies the belief that P entails C before she comes to believe that C as the result of an inference (which is one way she might have known that P entails C when she inferred) then P entails C is one of her *reasons* for believing C just as much, and in precisely the same way, as P is one of her reasons."

V-ii Reason to Infer

I disagree with the claim made in the first objection, and agree with only part of that made in the second. If I check that P entails C prior to coming to believe that C via an inference, then P entails C is indeed, just as much as P , one of my *reasons* for believing that C . But it is not a reason in 'precisely the same way' because it is not a *premise* of my argument for C . It is of course true that I do often check that P entails C before I infer from some P to some C , especially when that inference is not obviously valid, but then, as I have just said: I check that P entails C prior to inferring C from P , *alone* - I do not include P entails C as one of my premises. I do not, for instance, if operating within a system of logic with whose rules I am unfamiliar, verify that C follows from P according to one of the rules of inference of that system, and then go on to write P entails C below

⁸¹ It's been pointed out to me that such talk is reminiscent of some passages written by Bertrand Russell in *The Principles of Mathematics* (1903), but I have been unable to locate these passages, and understand that his views there, and mine here, are not quite alike.

⁸² See chapter 7 for further discussion.

P, before writing *C* (or perhaps (3), and then (4), and so on).⁸³ After I have *P*, I write down immediately whatever follows from it according to one of the rules, for all I am interested in is whether or not what I want to write down does indeed *follow from P*.

We have already seen, in chapter 2, that the hypothetical associated with an argument need not be, *per impossible*, a premise of that argument in order for the argument to be valid. Why then think that though this is the case, I must either employ *P entails C as a premise* of my argument for *C* nonetheless, or not bother checking that *P* entails *C* at all? If I check that *P* entails *C* I do so because I want to be sure that the inference I am about to carry out is valid, not because I am unsatisfied with that inference and wish to first supplement it by adding further premises and then carry out another inference entirely. It is precisely a verification of the truth of *P entails C* which *does* satisfy me that the inference of *C* from *P alone* is a good one; I then go ahead and infer.

J. F. Thomson writes of Achilles, who has just asked the tortoise to accept the next proposition in the regress, that '[i]n saying that *Z* follows from *A* and *B* and *C* Achilles implies that it does not follow from *A* and *B* alone, he implies that these premises are not by themselves sufficient.'⁸⁴ And that is quite right. Achilles' task was to get the tortoise to infer *Z* from *A* and *B*, and to do so he pointed out that *C* is true - that *A* and *B* together entail *Z*. He then makes the mistake of writing '*C*' under '*A*' and '*B*' in his notebook and asking the tortoise to carry out a *new* inference, from those three propositions together, *as premises*, to *Z*. In an attempt to do that, he asks the tortoise to accept another proposition, *D* (the one Thomson refers to above), and so on. But he should have explained to the tortoise that as long as he genuinely believes that *C*, he must carry out the *original* inference, from *A* and *B* to *Z*, and not have got started on this regress of premises.⁸⁵ Achilles is not

⁸³ See chapter 4 (III-ii Explicitly Accepting that *P* entails *C* in the Course of Inferring).

⁸⁴ Thomson (1960, p. 97)

⁸⁵ J.F. Thomson (1960) has a very interesting discussion of this. Another writer who makes a similar point is W. A. Wisdom (Lewis Carroll's infinite regress, p. 572).

... or should not regard himself as, asking [the tortoise] to accept another premise. For *ex hypothesis* he suppose[s] that [the tortoise] already [has] enough premises.⁸⁶

Achilles and the tortoise, are, indeed, both confused, just as I would be if I wrote down *P entails C* below *P* in my notebook after I had just in fact checked, by checking that *P* entailed *C*, that the argument from *P alone* to *C* is valid.

If I believe that *P* entails *C* then, for that reason, I must be (which is simply to say that I *am*) willing to infer *C* from *P* alone.⁸⁷ If I accept '*P entails C*' but remain unwilling to infer *C* from *P*, then either I do not understand what I have accepted or I do not understand what deductive inference is all about - that it is about believing whatever *follows* from whatever else I believe (in which case I surely couldn't be seen as an *inferer* at all⁸⁸).

I have said that if I believe that *P* entails *C* then, *for that reason*, I must be willing to infer *C* from *P*. But I want to add to this, riskily, one might think, because of the threat of regress, by saying that when I infer *C* from *P*, for some *P* and some *C*, I believe that *P* entails *C* and *my reason* for carrying out that inference is that *P* entails *C* (as I take it to). This, then, is the answer to our objector's question, "How can *P entails C* be part of your reason for believing that *C* except as one of the premises of your argument for *C*?"⁸⁹ As well as having more direct reasons for *believing* such-and-such, I can also have reasons for *acting* in such-and-such a way, for doing *this* or *that*, for instance, for *inferring C* from *P*. The idea is that I might first check that *P* entails *C* and then go ahead and *perform* the inference from *P* (alone) to *C* on that basis, so that if I were to verbalise what I was doing I would not say,

"*P*, and *P entails C*; *So C*,"

⁸⁶ Thomson (1960, p. 101)

⁸⁷ It is also true that if I believe that *P* entails *C* then I must be willing to infer *C* from *P together with P entails C* in appropriate circumstances, but this in no way undermines the claim I have just made.

⁸⁸ See (VI Entitlement to *Act*) below.

⁸⁹ This is in effect the challenge presented by the sceptic, recall, as claim (b) above. See note 3, in chapter 5 above.

which would be an inference from *P together with P entails C*, to *C*, but rather

“*P entails C; So, we have P; So C.*”

There is one ‘move’ (the first “so”) from the belief that *P entails C* to the *act* of inference, and another ‘move’ (the second “so”) which *is* the *act* of inference.

In arguing, you may need to point out that you are [arguing validly]. You then (as it were) step aside from what you are *doing* and comment on your own *performance*. But then the *performance* must be there, independently of the comment, to be commented on.⁹⁰ (My emphasis.)

Now if *P entails C* is my reason for *inferring C* from *P*, then it is perhaps in some sense one of my reasons for believing that *C*, for if I had not had the belief that *P entails C* and had not acted on it, then I would not have had the belief that *C*. But *P entails C* is at least not in such a case one of my reasons for believing that *C* in the same way as *P* is. If you ask for my reason(s) for believing that *C*, I will most likely tell you that it is *P*, for I know (or at least believe) that *P*’s truth is *by itself* sufficient for the truth of *C*, that is, I believe that *P entails C*.

If you ask me why I believe that today, Monday, is the 20th day of the month, I might offer as my reason that last Monday was the 13th. In such a case, I certainly *believe* that the latter proposition entails the former, but *that* belief has *not* here been deployed as a *premise* of my argument for the claim that today, Monday, is the 20th – the *only* premise of that argument is the proposition that last Monday was the 13th. Now, if you remain unconvinced that today is the 20th, I might respond by telling you that there are seven days in a week. In *that* case I have presented a *new* argument, with *two* premises, in an effort to convince *you* of the truth of the conclusion. But in any case in which I try to convince you of the truth of the conclusion by telling you that the premises we *already*

⁹⁰ Thomson (1960, p. 100)

have jointly entail the conclusion, what I give you patently was *not* a premise of the *original* argument, and is also not meant to be a premise of any *new* argument for the conclusion.

... to point out that someone has good or conclusive reasons for believing something is not to state another reason he has to believe it. The remark is a comment on the original premises; it says something about them, but it does not give a part of the original argument.⁹¹

I am not in total agreement with the claim that to point out that *P* entails *C* is not to give someone another *reason*, apart from *P*, to believe that *C*. Indeed, I think that *P* entails *C* is *always* part of my *reason*, in a sense, for believing that *C* when I infer *C* from *P* *alone*; but it is certainly not one of my *premises*. What we have in such a case is perhaps more correctly described as *P* being my reason for believing that *C*, with *P* entails *C* being precisely my reason *for* having *inferred* *C* from *P*.⁹²

VI Entitlement to Act

We've now reunited belief and practice, and we can see what mistakes we've made in previous chapters. In chapter 4, our mistake was not the very thought of attributing the implicit belief that *P* entails *C* to anyone who infers *C* from *P*, or the implicit knowledge that *P* entails *C* to anyone who does so with entitlement. Our mistake was the thought that doing so would inevitably lead to regress, that 'a fool might have all that knowledge without knowing how to perform'.⁹³ We are only led into regress by always attributing the belief that *P* entails *C* to someone who infers *C* from *P* if we mistakenly think that a belief can only be deployed so as to be useful for justifying another belief as a *premise* of an argument for that latter belief.⁹⁴ A belief, as well as possibly being deployed as a

⁹¹ Stroud (1979); see also chapter 2 (IV-i We Want What We Cannot Have).

⁹² See Wright (2002, p. 79-83) for discussion, along similar lines, of how the simple internalist would need to respond to Boghossian's Carrollian regress arguments.

⁹³ Ryle (1945-1946, p. 8); see chapter 4, section IV.

⁹⁴ I mean, of course, that the *proposition* believed should not be seen as a premise of the argument; I express it in this way here and below for ease of expression.

reason for *believing* something else in the sense of being a premise of the argument for that belief, can also be deployed as a reason for *acting* in a certain way, one of those ways being, specifically, *inferring*. And by seeing this, by remembering and maintaining the intimate connection between one's beliefs and one's inferential activities in this way, we can avoid the problem of regress. Thus, instead of being led to propose an only partially satisfying, insufficiently internalist account of inferential entitlement such as that of chapter 5, or being led to propose some even more far-fetched *externalist* account of inferential entitlement, my argument is that one need *only* know that an inference is valid, and have that knowledge deployed as one's reason for *carrying out* the inference, in order to be entitled to that act of inference. And the fact that one believes that the inference is valid in such a case also explains *why* one carries out the inference – so the view allows us to explain acts of inference as *rational* acts.

Now, I must provide further support - apart from bringing attention to our actual practices - for the claim that if I infer *C* from *P* I believe that *P* entails *C*, and *P* entails *C* is my *reason* for *inferring*, even if I cannot have the *explicit* belief that *P* entails *C*.⁹⁵ And I must provide further support for the claim that if I do not only believe that *P* entails *C*, but *know* it, then I am entitled to that inference.

Because I think that an inference is a kind of act, and because I see it as uncontroversial that anyone who knows that *C*, for some *C*, is entitled to the sincere *assertion* that *C*, which is also an *act*, I think it will help us to see clearly what is going on in the case of inference if we focus first on assertion, as we did in the previous chapter.

Just as we naturally attribute the (implicit) belief that *P* entails *C* to someone who infers *C* from *P*, we naturally attribute the (implicit⁹⁶) belief that *C* to someone who sincerely asserts *C*, thinking, of course, that she did not only *come to believe* that *C* once she had explicitly asserted it. We attribute the implicit belief that *C* to such an assertor in order to see her act of assertion as rational, and we can appreciate that we would be unable to do

⁹⁵ See my discussion of Ryle - chapter 4, section IV.

⁹⁶ This qualification often omitted below.

this if we rejected the idea of *implicit* belief as the idea of a 'myth'. Our understanding is that the subject has performed her act of sincere assertion *because* she believes that *C*; *C* is her *reason* for asserting that *C* - if she had not believed it she would of course *not* have asserted it.

Let us represent an assertion that *C* as follows.

$\Rightarrow C,$

where an arrow of this form indicates that an *act* has been performed.

My view is that when I sincerely assert that *C*, my reason for performing that *act* of assertion is that *C*. We can represent this in the following way.

$C \Rightarrow C$

But is this the correct view of the matter? I think that we can see that it is by taking a look at what the *point* of assertion is. The point of assertions may indeed vary from case to case, and I might have more than one aim in asserting something. I often assert something in order to communicate information to another. I sometimes assert something to make someone aware of the fact that I know it; sometimes (and possibly also) to give her a good reason to believe something else. I might also assert something just to remind *myself* of it. But in all these cases, no matter what the *ultimate* aim of it, an assertion is only successful, in a narrow sense, if the proposition asserted is *true*. In *any* case in which I believe *C* to be *false*, at least, I will most assuredly *not* sincerely assert that *C*.

To be able to assert at all I must understand that, generally, a proposition asserted is meant to be *true*. Although I can (*insincerely*) assert something though I know it to be false - if I wish to deceive someone - even this I cannot do unless I understand that the point of assertions generally is for them to be true; or else how could I imagine that I could deceive someone by asserting a falsehood? But if I must understand that the point

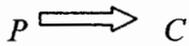
of assertions generally is for the propositions asserted to be true in order to be able to assert anything at all, then if I *sincerely* assert something I do so with the *intention* of the proposition I assert to be true. And if I *intend* for *C* to be true when I sincerely assert that *C*, then *whenever* I sincerely assert that *C*, I *believe* that *C* is true, and I have *C* as my *reason* for carrying out that act of assertion. There is an inextricable connection between assertion and truth.

That's why we are right to automatically attribute the belief that *C* to anyone who asserts that *C*, who we take to be sincere in her assertion. That's also why one of the most natural and most direct ways to challenge someone's entitlement to an assertion that *C*, as we've seen in chapter 5, is to claim that she doesn't have an entitlement to *believe* that *C*. *C* is her *reason* for performing the act of assertion, so by claiming that she is not entitled to her belief that *C* an attempt is made to undermine her reason for *acting* in a certain way, for *asserting* that *C*. It is uncontentious that anyone who does not have an entitlement to *believe* that *C* is not entitled to *assert* that *C*; and also that anyone who *is* entitled to *believe* that *C* is entitled to that act of assertion. A subject is entitled to assert that *C* if and only if she is entitled to *believe* that *C*. And it is most certainly the case that if she *knows* that *C*, then she is entitled to assert that *C*.⁹⁷

Now, similarly, we attribute the implicit belief that *P* entails *C* to anyone who carries out an inference from *P* to *C* in order to see her act of inference as rational, and we would be unable to see it as such if we rejected the idea of *implicit* belief - as the idea of a 'myth'. Our understanding is that the subject has performed her act of inference *because* she believes that *P* entails *C*; *P* entails *C* is her reason for *inferring* *C* from *P* - if she had not believed it she would not have carried out that inference.

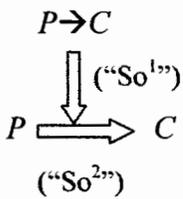
⁹⁷ Does an assertor need to know that she is entitled to her act of assertion *in order to be* entitled to that act? This might seem not to be the case given the truth of the last claim I have made in the text. But, perhaps to know that *C* I must know that I am entitled to *believe* that *C*, in which case we would worry less about the question of whether I know that I am entitled to an *act* of assertion that *C* just when I know that *C*.

Let us represent an inference of C from P as follows.



Again, the arrow indicates that an *act* has been performed – here, an act of *inference*.⁹⁸

My view is that when I infer C from P , my reason for performing that act of inference is that P entails C :



We have represented here the act of inference from P to C being performed with my reason for having performed it being P entails C ($P \rightarrow C$); we might say that I ‘moved’ (“So¹”) from my belief that P entails C to perform the act of inference (“So²”) from P to C .

But is this the correct view of the matter? Well, what is the *point* of (deductive) inference in general? Again, the point of inferences varies from case to case, and I might have more than one aim when I carry out an inference. Part of my aim in carrying out an inference might be to *further* my knowledge – maybe to come to have, firstly, a true belief I did not previously have, *and*, secondly, to have it *justified*. Or, my aim might be to merely *justify* a belief I already had, which may also have already counted as knowledge, as *justified* (because, say, I was *able* to justify it). Alternatively, I might want to *show* someone that a given belief of mine *is* justified. But in any of these cases, no matter what the *ultimate* aim of it, an inference is only successful, in a narrow sense again, if its premise entails its

⁹⁸ The context of discussion will make clear what sort of act – an assertion or inference, or whatever - I represent by one of these arrows.

conclusion; if it is *valid*.⁹⁹ In any case in which I believe *P* entails *C* to be *false*, at least, I will most definitely *not* deductively infer *C* from *P*, as I have already argued.

To be able to infer deductively at all I must understand that a deductive inference is meant to be *valid*. Now, we might think that I can carry out an inference while I know it to be invalid, for instance, if I am teaching logic and wishing to test for the vigilance of my class. But this will not do because inferring is more than writing one proposition after another on a board – to infer is to come to *believe* one thing on the basis of another belief, or to justify a belief one already has by noting the truth of something else (from which it *follows*). But in the case described I might not even believe the conclusion of my ‘inference’, and even if I do, I know full-well that I do not justify that belief by inferring it from the premise, from which it *does not* follow. There is no such thing, on my view, as ‘insincere’ inference.¹⁰⁰

Now if I understand that the point of inferences generally is for their premises to entail their conclusions, then if I infer some *C* from some *P* I do so with my intention being for *P* to entail *C*. And if I intend for *P* entails *C* to be true whenever I infer some *C* from some *P*, then I *believe* in some particular such case that *P* entails *C*, and I have *P* entails *C* as my *reason* for carrying out that act of inference. There is an ineradicable connection between deductive inference and deductive validity.

That’s why we are right to automatically attribute the belief that *P* entails *C* to anyone who infers *C* from *P*. That’s also why one of the most natural and most direct ways to challenge someone’s entitlement to an inference is to claim that she doesn’t have an entitlement to *believe* that her inference is valid. *P* entails *C* is her *reason* for carrying out the inference, so by claiming that she is not entitled to her belief that *P* entails *C*, an attempt is made to undermine her reason for *acting* by inferring *C* from *P*. Analogously to the case of assertion, then, a thinker is entitled to her act of inference from *P* to *C* only

⁹⁹ Recall that I argued in chapter 3 that for me to be entitled to an inference, for it to be ‘warrant-transferring’, the inference must be valid.

¹⁰⁰ Even if what I have said in this paragraph is questionable, I cannot test my students by carrying out an *invalid* inference unless I understand that deductive inferences are on the whole meant to be *valid*.

if she is entitled to her associated *belief* that *P* entails *C*. But being entitled to the belief that *P* entails *C* is not *sufficient* for being entitled to the inference from *P* to *C*, insofar as we want an inferer's entitlement to her inference to enable her to have *certain* knowledge of her conclusion as a result of her inference. As I argued in chapter 3, for *that*, it is required that *P* does in fact entail *C* – her belief must be true. Therefore, assuming¹⁰¹ that when a thinker carries out an inference with entitlement, the entitlement to her belief that it is valid, together with the belief's truth, is sufficient for her knowing that the inference is valid, any reasoner who is entitled to an inference in fact *knows* that inference to be valid. And I think that, again, it is most certainly the case that if she *knows* that her inference is valid, then she is entitled to it.

There is an ineradicable connection between truth and assertion, and an assertion is only successful, in a narrow sense, if the proposition asserted is *true*, as I argued above. Thus, if the tortoise witnessed me asserting that *C* without providing any justification for that act of assertion, his natural challenge would be to say that I am not entitled to believe that *C* is true.

If, on the other hand, I had the tortoise grant me first that I *know* that *C* is true (however I might do such a thing), and then went ahead to assert that *C*, how could the tortoise raise a sensible challenge to that act of assertion?

$C \implies C$

He could not.

¹⁰¹ See chapter 7.

If, however, I convinced the tortoise not that I knew that C , but that I knew that P , for some proposition P from which C followed, even easily, and then went ahead and asserted that C on the basis of my knowledge that P (which I take to be in no important way different from *inferring* C from P here):

$$P \implies C,$$

the tortoise's method would be to challenge me by saying that I did not know that P 's truth guaranteed the *truth of* C – for my assertion that C is successful only if C is *true*. There is an ineradicable connection between inference and validity, and an inference of mine is only successful if it is valid. To be entitled to an act of inference from P to C it must *at least* be the case that I am *entitled* to believe that P entails C .

Finally, imagine that the tortoise granted me not only an entitlement to believe, but *knowledge* that P entails C , and that I went ahead to carry out the inference from P to C on the basis of that knowledge:

$$\begin{array}{ccc}
 P \rightarrow C & & \\
 \Downarrow \text{("So"}^1\text{")}& & \\
 P \implies C & & \\
 \text{("So"}^2\text{")}& &
 \end{array}$$

How could the tortoise raise a sensible challenge to that act of inference - from P , *alone*, to C ? I am sure that he could not.

Chapter 7

Objections, Implications, and Conclusion

I Introduction

In chapter 3 I argued that, for any P and C , it must at least not be the case that I ought to doubt that P entails C if I am to be entitled to the inference from P to C . But because, for a simple argument of the form of modus ponens, say, it is enough for me to know of the truth of the premise in order to be entitled to infer the conclusion, we must maintain that, in such a case, when I know the premise I *do* have the required entitlement to believe that the inference is valid. In chapters 5 and 6 I have provided some indication of why, on my view, anyone in that position has the required entitlement to believe that her inference is valid – she is epistemically blameless in having that belief; we cannot sensibly maintain that she ought to have doubt about it. In fact, I think that she *knows* that her inference is valid, and is therefore entitled to carry it out. I shall now clarify and refine my position, showing how it stands up to some of the challenges we have already encountered, and considering how we might deal with some unseen objections to it. The main challenge will be to defend the view that being blameless in this way in having a belief, together with the belief being *true*, is sufficient for having knowledge. In section III I shall consider the implications of my account for the problem of the justification of deduction raised in chapter 2.

II *How do You Know?*

In chapters 5 and 6 I argued that we'd do better not to think that a reasoner has the required entitlement to believe that a given modus ponens inference is valid because of some mysterious faculty of 'rational insight' of hers. And that we cannot think that anyone who carries out an inference by modus ponens with entitlement has an *inferential* entitlement to believe that the inference is valid; having some item of propositional

knowledge which counts as a good reason to believe that it is valid. I have said that these two options are seemingly exhaustive of the possible ways in which a thinker could have the required entitlement which could be entirely satisfying to an internalist, and I now wish to return to this point, which raises a certain difficulty for my account.

For an argument of the form of modus ponens, with premise P and conclusion C , P entails C is true by virtue of the meaning of the conditional, since the inference from P to C is one whose very goodness in part *determines* the meaning of those two propositions, and of the conditional. (I shall explain this more carefully shortly.) And I have argued, with Boghossian, that anyone who has the concept of the conditional is willing to infer from such a P to such a C ¹, but, further, that she *believes* that P entails C . But it might be better to say here that she *understands* that P entails C – if she has the concept of entailment she must *understand* that this C follows from this P , that the inference from P to C is a good one, since this is true by virtue of the very inferential norms which *determine* the meaning of *that* concept - entailment - of which she has possession. The knowledge that P entails C is part of the knowledge she has by virtue of having an understanding of her language, having its concepts. And we might think that calling such knowledge ‘understanding’ brings attention to this fact though it is indeed all the same a kind of *knowledge* possessed. By this view, as I have urged, the knowledge which one has when one has an understanding of one’s language, or a possession of its concepts, has both a practical and a theoretical aspect.

The grammar of the word “knows” is evidently closely related to that of “can”, “is able to”. But *also* closely related to that of “understands”. (‘Mastery’ of a technique)² (My emphasis.)

Thus, we can see why it is not because of a faculty of rational insight or the possession of some item(s) of propositional knowledge that a reasoner has an entitlement to believe that our given inference from P to C is valid. To have the concept of entailment one must understand, for such a P and C , that P entails C , so anyone who believes that P entails C ,

¹ Again, as long as she can understand these two propositions.

² Wittgenstein (1953, remark150)

who therefore has the concept of entailment, already *knows* that *P* entails *C* (by virtue of having an understanding of her language). One cannot first *believe* that *P* entails *C* and then be able to ‘see’ that it is true or justify it in some other way, because, again, to believe that *P* entails *C* one must have the concept of entailment, and this requires already understanding that *P* does indeed entail *C*. We have already seen Black’s expression of a thought along these lines.

... to understand a proposition of this sort is already to be in possession of all that is relevant to its truth. Thomas Reid said self-evident propositions ‘are no sooner understood than believed’, but the correct formula is, rather, that understanding a self-evident proposition necessarily includes *knowing* the proposition to be true ... if we have any doubt we could not have understood what we thought we were doubting. There can be no suspension of belief about a self-evident proposition.³

In chapter 6 I indicated how we could have a deeper explanation, with the help of Boghossian’s notion of epistemic blamelessness, of why it is that anyone who has the concept of the conditional is entitled to a given belief she has that some modus ponens inference is valid, and in fact *knows* that it is valid without having an antecedent warrant for that belief. She must have the belief in order to have the given concept, and she is therefore *epistemically blameless* in having the belief, which is a *true* belief (see below). To clarify this now, and to lead into discussion which will be valuable to us in sections to follow shortly, we might consider whether it is plausible that anyone who has the concept of entailment in fact believes that *all* instances of modus ponens are valid. For if that were the case then we could argue, as above, that anyone who has the concept of the conditional in fact *knows* that *all* instances of modus ponens are valid – that modus ponens is truth-preserving.⁴

³ Black (1970, p. 18); Reid (1986, p.358)

⁴ Note that this would make a rule-circular argument (or *any*) for the validity of MP otiose – anyone who could construct, or even understand the propositions involved in, such an argument would have the concept of entailment and therefore already know MP to be truth-preserving. See below.

Boghossian is in fact relatively sympathetic to the view that ‘there may be beliefs that are such that, having those beliefs is a condition for having one of the concepts ingredient in them. Thus,’ Boghossian tells us, ‘Christopher Peacocke has written of a special case

in which it is written into the possession condition for one or more concepts in [a] given principle that to possess those concepts, the thinker must be willing to accept the principle, by reaching it in [a particular] way.’⁵

And ‘the special case that Peacocke has in mind concerns our belief in the validity of the basic truths of deduction.’⁶

But Boghossian contends that ‘it is not remotely plausible that anyone having the concept of [the] conditional would have to have the *belief* that [modus ponens] is valid.’ For, he claims, ‘one can have and reason with [the] conditional without so much as having the concept of logical implication.’⁷

I agree that it is too much to require that any subject must believe that *all* instances of modus ponens are valid in order to have the concept of the conditional. But it should be evident from what has gone before that I do *not* agree that someone can be a deductive *inferer* without having the concept of entailment; without having an understanding of the very idea of deductive inference – which is to believe one thing which *follows* from something else one believes. I have argued that *anyone* who carries out a deductive inference believes that inference to be deductively valid - if we deny this then we cannot explain her act of inference as a *rational* act - and this surely requires her to have the concept of entailment. But yes, of course, we have been given ample opportunity to appreciate the fact, stated in Brandom’s terms, that

[t]he grasp of logic that is attributed must be an *implicit* grasp, since it need be manifested only in distinguishing material inferences as good and bad, not in any further

⁵ Peacocke (2000), cited by Boghossian (2002, p. 23)

⁶ Boghossian (2002, p. 23)

⁷ Boghossian (2002, p. 23)

capacity to manipulate logical vocabulary or endorse tautologies involving them.⁸ (My emphasis.)

In chapter 4 we saw the contention from Ryle that the idea that anyone who carries out an inference has the *implicit* belief that the inference is valid, whether or not she has acknowledged that it is valid, is the idea, as he puts it, of a ‘myth’. I hope, however, that I have shown that this idea is, apart from being a quite natural one which is embedded in our very practices with respect to inferential activity, an idea from which we *must not* withdraw if we are to have a satisfyingly rounded account of inference. Indeed, the idea that a subject only believes something if she has explicitly accepted it, which would seem to be the suggestion of Ryle, is quite absurd. Beliefs are rightfully attributed to agents in order to explain their *acts*, and even to explain their behavioural *dispositions*. So,

if the-fact-that-p convinces [someone] that q, then he thinks that if p, then q even though he never formulates it.⁹

And these considerations can help to explain why Bonjour, with his allegiance to the notion of rational insight, is even willing to maintain - implausibly, I would say - that

...even an apparent rational insight ... must involve a genuine awareness by the person in question of the necessity or apparent necessity of the proposition in question in something like the strong or metaphysical sense ... thus requiring at least an approximate (though perhaps, in some cases, very implicit) grasp of that very demanding concept of necessity.¹⁰

I don’t, however, have the same reasons as Boghossian¹¹ for finding what Bonjour has written here disagreeable. I agree with Bonjour that *any* inferer has the ‘in some cases, very implicit ... grasp of [the] very demanding concept of necessity’. But, as I have explained in previous chapters, I think that the idea that *anyone* who carries out an

⁸ Brandom (2000, p. 53)

⁹ Brown (1954, p. 177)

¹⁰ Bonjour (1985, p. 114), cited by Boghossian (2001, p. 637)

¹¹ See Boghossian (2001)

inference is able to *see* that the inference is valid stretches our imagination beyond reasonable limits. We have very good reason for wishing to remain faithful to the view that a thinker infers *because* she believes her inference to be valid, but we have seen that it is *by far* more plausible to explain how she is *entitled* to that belief in basic cases in terms of Boghossian's notion of epistemic blamelessness, than to opt for a theory of rational insight.

Now, to return to the possibility of maintaining that anyone who has the concept of the conditional believes that *all* instances of modus ponens are valid, I can see only one plausible way of defending this view.¹² If being disposed to infer from some particular proposition to another involves believing that the former entails the latter, then being disposed to carry out *any* modus ponens inference one might be faced with surely involves believing that *all* instances of modus ponens are valid. So *if*, in order to have the concept of the conditional a subject must have this latter disposition, then she must have the belief that *modus ponens* is truth-preserving. But unfortunately, although I have followed Boghossian up to this point in saying that someone who has the concept of entailment must be willing to infer, for *any P* and *any C*, from *P* together with *P entails C*, to *C*, the reader can probably see that this view of the matter is more than a little naïve. We'll shortly have the opportunity to examine it more closely.

Before moving on to consider how we might deal with some further possible objections to my account, however, I shall explain more carefully my current position, according to which: where *P* is the premise and *C* the conclusion of some modus ponens argument, anyone who has the belief that *P* entails *C*, being disposed to infer *C* from *P*, in fact *knows* that *P* entails *C*. Assume for now¹³ that anyone who has the concept of entailment is disposed to carry out *any* inference of the form of modus ponens, and thus believes, for any given modus ponens inference, that it is valid. We already have a good understanding of why, on my view as it now stands, anyone who believes that *P* entails *C*, where *P* is the premise and *C* the conclusion of a modus ponens inference, is epistemically *blameless*

¹² See also (III-i Begging the Question) below.

¹³ See sections below where I discuss the question of whether a subject must be disposed to carry out *any* given modus ponens inference in order to have the concept of entailment.

in having that belief. She is epistemically blameless in having it because she *must* have it in order to have the concept of entailment in the first place, so she cannot be faulted for not having an antecedent warrant for the belief because having such a warrant requires that she *does* have the concept of entailment, and therefore that she already has the belief.

But, of course, being epistemically blameless in believing something is not sufficient for *knowing* it. For knowledge, we also need at least¹⁴ the truth of the belief. So I shall now explain with greater clarity why any belief that *P* entails *C*, where *P* is the premise and *C* the conclusion of a modus ponens inference, which is a belief that any possessor of the concept of entailment must have, is in fact *true*.

The belief is true since the inference from *P* to *C* is in such a case an inference whose very correctness partly determines the meaning of the conditional, and therefore partly determines the meanings of those very propositions, *P* and *C*. By a conceptual role semantics

A logical constant *C* expresses that logical object, if any, that makes valid its meaning-constituting inferences.¹⁵

If a rule is 'genuinely meaning-constituting then it has a semantic value such that [it] is truth-preserving.'¹⁶

The conditional means what it means in part by virtue of participating in all inferences which are of the form of modus ponens. So, since any modus ponens inference in part *constitutes* the meaning of the conditional, it is impossible that there should be some such inference whose premise was true and conclusion false. Any genuinely meaning-constituting rule is truth-preserving.

¹⁴ See below - (II-iv Epistemic Blamelessness is indeed *not* sufficient for Warrant).

¹⁵ Boghossian (1996, p. 25)

¹⁶ Boghossian (1999, p. 251)

For a slightly different but very closely related approach, let's look at Peacocke's account of the link between possession conditions and knowledge. In Peacocke's terms, a subject possesses a given concept only if she finds 'primitively compelling' certain inferences (or 'transitions') and certain contents. Loosely translated¹⁷ – a subject possesses a given concept only if she is disposed to carry out certain inferences and accept certain propositions. Peacocke sees this approach as 'consistent with acknowledging a major insight of the later Wittgenstein. This is the insight that an account of what is involved in employing one concept rather than another, following one rule rather than another, has at some point to mention what thinkers employing the concept find it natural to believe.'¹⁸

Now, 'at an intermediate level of generality', Peacocke formulates the link between possession conditions and knowledge as follows.

Link Between Possession Conditions and Knowledge Take any mental state of the thinker that a possession condition for a concept says is sufficient for the thinker finding primitively compelling a given content containing that concept. Then when a thinker judges that content and for the reason that he is in that state, his judgement constitutes knowledge.¹⁹

Let's see how this principle would be applied to the sort of scenario in which we are interested. For any P and C which are the premise and conclusion of a modus ponens inference, according to the possession condition for the conditional, anyone who believes (who is in the 'mental state' of believing) that P , finds primitively compelling the content of the proposition C – she is disposed to accept that C . But then if such a subject *does* accept and believe that C , and 'for the reason that [she] is in that state' (of believing that P), her belief is a case of knowledge.

¹⁷ See Peacocke (1992, p. 6): 'To say that the thinker finds such transitions primitively compelling is to say this: (1) he finds them compelling; (2) he does not find them compelling because he has inferred them from other premises and/or principles; and (3) for possession of the concept C in question ... he does not need to take the correctness of the transitions as answerable to anything else. (The thinker may in fact take them as answerable to something else. If he is a philosophical theorist himself...)'

¹⁸ Peacocke (1992, p. 13)

¹⁹ Peacocke (1992, p. 157)

Peacocke explains this link in a similar way to what we have already seen above. It is explained, firstly, as a ‘consequence of the approach ... under which semantic values are assigned so as to make the practices of belief formation mentioned in the possession conditions always *correct*.’ (My emphasis) Secondly, one’s *entitlement* to the belief is defended as follows: ‘If your reasons for a belief ensure its truth and do so as a consequence of the nature of the concepts it contains, you are in the best possible epistemic position with respect to the content of the belief.’²⁰

I shall turn now to consider how my account, and these we have seen of Boghossian and Peacocke, fare under pressure – my account, that is, of how someone who carries out an inference by a meaning-constituting rule is entitled to the associated belief she has that the inference is valid, in fact knows that it is valid, and is entitled to the inference.

II-i Reflectively Appreciable Warrant

I cannot be required, for *any* inference I carry out, to have an *inferential* justification for the belief that it is valid, and cannot be required to be able to *see* that it is valid. Taking my cue from *Boghossian’s* explanation of how a subject is entitled to the *acts* of inference in accordance with meaning-constituting rules she carries out, I have explained my entitlement to the attendant belief that such an inference is valid, when I carry it out, in terms of *epistemic blamelessness*. But it is obvious that to be entitled to a belief “because I cannot have justified or be in a position to justify the belief until I have its ingredient concepts, and cannot have *those* until I in fact have the belief”, is not to have an *internalist* warrant for the belief.

What precisely is it for the warrant I have for some belief to be properly internalist? In chapter 5 I suggested that for the warrant for my belief that a given inference is valid to be properly internalist it would have to at least be the case that I possessed either some item(s) of propositional knowledge which counted as good reasons for believing that it was valid, or a faculty of rational insight which enabled me to ‘see’ that it was. The idea

²⁰ Peacocke (1992, p. 157-158)

was that when I carry out some modus ponens inference, my belief that the inference is valid must be one to which I am entitled because of something ‘internal’ to me – the belief must be justified either because I *have* justified it by an inference or by an act of insight, or *can* do so by reflection alone, without any further gathering of ‘information’ not already in my possession.

In ‘On basic logical knowledge’²¹, Crispin Wright applies pressure on Boghossian’s notion of epistemic blamelessness, which, recall from chapter 6²², the latter introduced in an effort to make space for an account of inferential justification to occupy a middle-ground between the extremes of simple inferential externalism and simple inferential internalism. To gain clarity on what such an account of the *internalist* ilk would look like, Wright, speaking for the *externalist*, says that an *externalist* account of inferential entitlement ‘will seem unacceptable only if one takes possession of a warrant to require being in a position *reflectively* to *appreciate* that one has one.’²³

Now, even to be able to offer up a reason for believing something, or to be able to apprehend its truth by an act of rational insight, is not by itself sufficient for ‘being in a position *reflectively* to *appreciate* that one has’ a warrant for believing it. To have a good *reason* to believe something, for instance, is not to know or have reason to believe *that one does*. And it is not to know or have good reason to believe that one is justified in believing what one believes. But since we have already accepted that someone who carries out a modus ponens inference is not required to have justified or be able to justify *inferentially* her belief that the inference is valid (and not required to have ‘seen’ or be able to ‘see’ that it is valid), we need not consider any further whether, or how, having *reasons* for believing in, or insight into, the validity of an inference would be to have an *internalist* warrant for believing that it is.

The view at hand is that having certain dispositions to infer – in accordance with the rule of modus ponens, say - is a precondition for having certain *concepts* – the concept of the

²¹ Wright (2002)

²² See chapter 6, sections III-I to III-iii.

²³ Wright (2002, p. 59)

conditional. And that one is therefore epistemically blameless in carrying out such inferences albeit without having an antecedent warrant for doing so, because to have such a warrant one must first have the relevant concept. On my account, having the disposition to carry out those inferences involves having the *beliefs*, for each of them, that they are valid, and one is for similar reasons epistemically blameless in having those *beliefs*. But now Wright correctly points out that the idea here is clearly *not* ‘that exercise of those dispositions on warranted premisses puts one in position [*reflectively to appreciate* that one has a] warrant for the conclusion.’

If we had an explanation of how an unsophisticated, unreflective, reasoner had a reflectively appreciable warrant for the *conclusion* of a modus ponens inference she carried out, what would it look like? Given my view of how her belief that the inference is valid figures in her act of inference as her *reason* for carrying it out, to explain how she had a reflectively appreciable warrant for the belief that the *inference is valid* would be to get us at least some of the way. And presumably we would be required to provide this latter explanation if we were to have the complete explanation – of why she had a reflectively appreciable warrant for the *conclusion*. But on our account, of course, this reasoner has no such reflectively appreciable warrant for the belief that her inference is valid. The reasons for her being justified in that belief are not *her* reasons. That belief is warranted, the claim is, because the rule according to which she has inferred is meaning-constituting of the conditional – so she must be disposed to carry out that inference and have the relevant belief associated with it in order to have the concept of entailment, and is therefore epistemically *blameless* in having that belief without having an antecedent warrant for it. But she has of course not even an inkling of these reasons for her being worthy of praise (or at least unworthy of blame) for having the belief that her inference is valid. Indeed, the inferer might not even be able to cite the belief that the inference is valid as her reason for inferring.

I shall shortly return²⁴ to defend my view as satisfactory, even if partially externalist, and more satisfactory than an account such as Boghossian’s which makes no reference to the

²⁴ In (II-iv Epistemic Blamelessness is indeed *not* sufficient for Warrant).

belief that a given inference is valid which a thinker has when she carries it out, or is disposed to. I'll argue that a reasoner has *certain* knowledge of the truth of a conclusion when she has certain knowledge of premises from which she infers the conclusion in accordance with a meaning-constituting rule. Revisiting the discussion of chapter 3, where I argued that to have certain knowledge of the truth of some proposition reached by some inference it is required (a) to be the case that the inference is valid, and (b) to not be the case that I ought to doubt that it is; I shall argue that the satisfaction of these two constraints, *together with* the inference being one in accordance with a meaning-constituting rule, is *sufficient* for it being possible²⁵ that I should have knowledge of the truth of the conclusion, and knowledge that is certain. Before then, we must take a look at two problems which Wright²⁶ raises for Boghossian, which are, potentially just as much, problems for me too.

Instead of granting that his account is partially externalist in nature, Boghossian proceeds to make a case for the claim that it is in fact 'properly *internalist* (in spirit).'²⁷ As Wright puts it, Boghossian asks 'what *virtue* there is in seeking reflectively appreciable warrants and returns the answer that, by so doing, a thinker is able to form beliefs in a fully *epistemically responsible* fashion.'²⁸ Boghossian's idea is that an inference via a meaning-constituting principle is warranted, and in a way which is (or should be) acceptable to an internalist. Why? Because an act of inference which proceeds according to a meaning-constituting rule is one which is not epistemically irresponsible, and lack of epistemic irresponsibility is what the internalist constraint really boils down to. To have a reflectively appreciable warrant for an act (or belief, I would say) is one way to be epistemically blameless in that act (or belief), which is really all that an internalist requires of an act or belief for it to be warranted. But to have a reflectively appreciable warrant is, as we have seen the contention to be, not the *only* way of being epistemically blameless.

²⁵ We also need it to be the case that I have certain knowledge of the truth of the *premise*, and - if I am to acquire *new* knowledge by the inference - justification for believing the premise which is suitably independent of the justification I have for believing the conclusion.

²⁶ See Wright (2002, p. 62-69)

²⁷ Wright (2002, p. 60)

²⁸ Wright (2002, p. 60)

But now, *Wright's* contention is that, (i) an act of inference being in accordance with a meaning-constituting rule is *not* sufficient for the epistemic blamelessness of the act; and (ii) an act being epistemically blameless is *not* sufficient for it being warranted.²⁹

II-ii Is an act of inference being in accordance with a meaning-constituting rule sufficient for its epistemic blamelessness?

Let's start with a consideration of Wright's argument for (i)³⁰. He has us consider a pair of natural deduction rules (the 'N-rules') and informs us that he and others have provided argument which puts them in a 'position tantamount to [holding] the view that the two rules ... serve to constitute a suitable meaning for the [cardinality] operator 'N'. [But] critics have rejoined that the analogy in structure between ... the N-rules and those for the course-of-values operator [which are unsound³¹] casts doubt on their suitability to serve' this function. Now, Wright argues that even if the two rules at hand *are* properly meaning-constituting of the cardinality operator, as *he* believes them to be, a subject who inferred by them would surely be epistemically irresponsible in doing so, since, to put it one way, Wright's critics 'are at least owed an answer'.

This argument, though presented somewhat briskly (by me here *and* by Wright), is indeed quite persuasive. I must confess to not being in the best possible position to offer an adequate response to it, but I don't think that it's a crippling objection to the accounts of Boghossian and myself. I shall attempt to work out some of the details of one possible avenue of response. The discussion of the remainder of this sub-section shall be carried over to section II-iii.

Wright's idea, as I understand it, is that at least if there is some unresolved debate about whether or not some given rule is (perhaps partially) meaning-constituting of a given operator, then, even if it *is*, an act of inference via that rule is not one for which a thinker

²⁹ Wright (2002, p. 62)

³⁰ Wright (2002, p. 64)

³¹ See below - (II-iv Epistemic Blamelessness is indeed *not* sufficient for Warrant).

is (epistemically³²) blameless. But now, what could be a ‘critic’s’ reasons for opposing the view that the relevant rule is meaning-constituting? She might of course have more than one reason, and different critics might themselves have different reasons. But let’s look specifically at *Wright’s* critics, then. They point to the fact that the structure of the rules in which we have a special interest, the N-rules, and the structure of those rules for the course-of-values operator, are analogous. And Wright tells us earlier in his paper³³ that these latter rules ‘are—famously—unsound’. So, presumably, the critics of Wright worry that the N-rules are likewise *unsound*.

Now here, as I see it, is the rub. To change tack slightly so as to focus on the specificities of my account, my view is that if the N-rules are indeed constitutive of the meaning of the cardinality operator then, following Boghossian, anyone who has the concept of that operator is disposed to infer according to the rules and therefore, for each inference (by either one of the rules) which she is disposed to carry out, she believes that it is valid. And she is blameless in having these beliefs without having an antecedent warrant for them, the thought is, because she cannot coherently be expected to have such an antecedent warrant. But Wright’s suggestion seems to be that she is *not* blameless in having these beliefs even if the N-rules *are* meaning-constituting unless she *does* have an antecedent warrant for them – unless she has some acceptable response to Wright’s critics. And this, despite the fact that Wright is evidently quite sympathetic to a conceptual role semantics. What then, according to Wright, is the flaw in the reasoning I have just presented? There must be some mistake I (and Boghossian) have made if Wright’s position is to be coherent.³⁴

Since Wright’s critics have doubt about the soundness of the N-rules, they presumably will have doubt about the validity of some inferences of the form of either one of those rules with which they are presented. Then, if the N-rules do indeed serve to constitute a meaning for the cardinality operator, these doubters supposedly *do not*, on my account as it stands, have the concept of that operator! If they ‘have any doubt [they] could not have

³² I omit this qualification below.

³³ Wright (2002, p. 63)

³⁴ See the next sub-section for a fuller discussion.

understood what [they] thought [they] were doubting. There can be no suspension of belief about a self-evident proposition.³⁵ Or, more in line with my proposal - and truistically, perhaps: There can be no suspension of belief about a proposition which one is *required* to believe in order to have some concept which is a component of that very proposition.

The surprising situation we would then have is thus as follows. Either the N-rules constitute a suitable meaning for the cardinality operator, and Wright's critics do not have the concept of that operator, and do not understand what they are talking about (or are *not* talking about anything!); or, Wright's *critics* are right, and Wright is wrong! – the N-rules *are not* meaning-constituting of the cardinality operator. Which is it? Surely not the former—Wright's adversaries are in possession of the given concept, if *anyone* is. So, are we then plausibly in a position, on the basis of these reflections alone, to conclude that Wright is mistaken in thinking that the N-rules constitute a suitable meaning for the cardinality operator – even without having considered the arguments he has elsewhere presented in defence of this claim?

Surely not. As I have said, I am not particularly well-equipped at this point to offer a proposal of how we might come to a satisfactory resolution of the issues at hand. I *will* say, however, that since it is evident that Wright's opponents have the concept of the cardinality operator, but also that Wright might well be correct in his view that the N-rules can serve to constitute a suitable meaning for that operator, we must maintain that rules might *be able* to serve such a function for a given concept, and yet be ones that not all *possessors* of that concept are disposed to infer by. Does that not, however, undermine my (and Boghossian's) whole project? I don't think it's cut and dry.

To maintain that it does not we would have to hold a view, I think, something along the lines of the following. All of Wright and his critics certainly already share (inferential) practices in which the cardinality operator figures, and that is why they are, all of them, in possession of the relevant concept. The meaning of the cardinality operator is already

³⁵ Black (1970, p. 18)

determined by some set of principles for its use which does not include the principles given by the two N-rules. But it may yet be coherent to hold the view, and also coherent to oppose it, that the N-rules can also serve (by themselves, presumably) to constitute a suitable meaning for that operator. Perhaps, for instance, we can maintain that Wright's proposal reduces to the thought that the principles for the use of the cardinality operator currently adhered to by all concerned can be *replaced* by those given by the N-rules; perhaps in that way providing a foundation for them. At this point, however, I am probably already out of my depth, and am unsure about how to proceed. I must add, however, that this last idea of mine would have the somewhat unwelcome consequence that a concept might have more than one set of use-governing rules, where each set is by itself sufficient for constituting a suitable meaning for the concept. And, therefore, even if some particular principle is partially meaning-constituting of a concept, being a member of one set of principles which together fully determine the meaning of that concept, a thinker might have the given concept without endorsing that principle as good – for she might endorse the goodness of each member of a *distinct* set of principles which likewise, together, bestow a determinate meaning on the relevant concept; and, in some sense, the *same* meaning. I'll pursue these matters further in the next sub-section.

II-iii The set of inferences which are meaning-constituting of a concept *for a thinker*

Our considerations above make it seem highly implausible that, for any concept, there is a single set of inferences which each partially constitute the content of the concept, and jointly *suffice* to constitute it; so that for each of the inferences in *that* set, all thinkers who have the concept have the disposition to carry it out.

Even in the case of a basic logical constant, a thinker might not endorse the goodness of a particular inference which is an instance of one of the rules standardly taken to be meaning-constituting of that constant, and yet have the relevant concept. Take the conditional. Timothy Williamson, approaching the issues at hand by first arguing that one

has the concept *if* as long as one understands a word that means *if*³⁶, brings to our attention the following fact.

Vann McGee, who is a ‘distinguished logician’, ‘refuses to make some inferences by modus ponens’, as is evident from a reading of his ‘A counterexample to modus ponens’.³⁷ McGee has constructed some pretty uncomplicated arguments of the form of modus ponens, and argued that they are not valid. But

[d]oes McGee lack the concept *if*? Although I deny that his purported counterexamples to modus ponens are genuine, I admit that they have some initial plausibility on the ordinary understanding of ‘if’. In conversation with McGee, he appears to understand the word ‘if’ quite well by ordinary standards. He certainly seems to know what we other English speakers mean when we use the word ‘if’. Before he had theoretical doubts about modus ponens, he understood the word ‘if’ if anyone has ever understood it; surely his theoretical doubts did not make him cease to remember what it means. We may therefore assume that McGee has the concept *if*, just like everyone else.³⁸

It may be hard to believe that there are some modus ponens inferences which Vann McGee, the ‘distinguished logician’ that he is, refuses to carry out. And perhaps even more astonishing a fact that there are some inferences of the form of modus ponens which he believes to be invalid. But it is not important to take a look at McGee’s purported counterexamples to modus ponens, because it is undeniable that McGee believes what he sincerely tells us he believes, and is not willing to carry out a given inference when he says that he is not. Yet, it is equally undeniable that McGee has the concept of the conditional, ‘just like everyone else’.

What, then, to make of this fact? It is, now that we look at it, not all that surprising a fact at all. Any concept I have is ingredient in an indeterminate number of inferences which I am disposed to carry out. But, for *any* given one of these inferences, I might *not* have

³⁶ Williamson (2003, p. 3)

³⁷ McGee (1985)

³⁸ Williamson (2003, p. 4)

been disposed to carry it out and yet surely still have had the given concept. Here is the Stroud's expression of the point, which will also, I think, give us an idea of how best to proceed.

Understanding something is a complex matter. Certainly if someone saw *none* of the relations between a particular proposition *P* and others, and was not disposed to accept it even though he already believed what were obviously good reasons against it, we would at some point justifiably conclude that he did not understand it. But that shows at most that understanding something requires seeing *some* (perhaps even a great many) of its obvious connections with other things.³⁹

If 'understanding something requires seeing *some* (perhaps even a great many) of its obvious connections with other things', then having a *concept* requires seeing the truth of some (perhaps even a great many) of the propositions which are obvious consequences of other propositions one believes, which have that concept as a component. And since the conclusion of a modus ponens argument is an obvious consequence of its premises, having the concept of the conditional requires seeing the truth of some (in fact, a great many, I think) of the conclusions of modus ponens inferences whose premises one believes.⁴⁰

To move away from the analogy with sight, the idea now is that anyone who has the concept of the conditional is disposed to carry out a great many inferences by modus ponens. This is surely right – if we saw that someone was only willing to carry out a (relatively) small number of inferences by modus ponens, for whom there was a (relatively) great number of modus inferences she was *not* disposed to carry out, we would not say that she had the concept of the conditional; and rightly so. Of course, we might wrongly, on the basis of some subject's inferential behaviour, refuse to attribute to her the possession of the concept of the conditional. If it happened that most of the modus

³⁹ Stroud (1979, p. 191)

⁴⁰ Williamson (2003, p. 4) also considers such an option as a possible refinement of the possession-condition for the conditional – 'Is a willingness to make most inferences by modus ponens (no matter which ones) a precondition for having the concept *if*?' In a single paragraph he dismisses the idea on the basis of it suffering from the problems I am about to consider.

ponens inferences she came across were of the McGee problem case variety, she might refuse to infer by modus ponens most of the time – having been persuaded by McGee’s arguments – but still have the concept of the conditional.⁴¹ What is important, however, what is *necessary* for her to have that concept, is that she be *disposed* to carry out a great many inferences of the form of modus ponens.

Now, this might not help us to argue (in the way that I have above) that anyone who genuinely believes the premises of some particular modus ponens inference *must* accept its conclusion, since she could presumably have the concept of the conditional (and therefore *could* believe the premises of that inference) without being disposed to carry out *that* particular inference.⁴² Well, the case of Vann McGee should have already convinced us that it is *not* the case that anyone who genuinely believes the premises of a modus ponens inference *must* accept its conclusion. But I do think that our refined position can allow us to argue that a reasoner is blameless in many of the inferences she carries out by modus ponens; and blameless, for each of those inferences, in having the *belief* that it is valid.

Let us work with Boghossian’s idea ‘that there is a particular set of inferences involving ‘if, then’ that are meaning-constituting [of the conditional] *for a thinker*’.⁴³ (My emphasis.) Our thought might now be that, given the set of inferences which are meaning-constituting (I’ll call it her ‘M-C set’) of the conditional for some subject, for each and every one of those inferences, the subject needs to have the disposition to carry it out in order to have the given concept. All the same, another thinker’s M-C set for the conditional might (or probably *will*) be distinct from our first subject’s set, though these sets will have a substantial overlap and both serve to constitute a suitable meaning for ‘if, then’. Both could, the hope is, rightfully be said to have the concept of *the conditional*.⁴⁴

⁴¹ Bernhard Weiss brought this point to my attention.

⁴² See Stroud (1979, p. 191-192) We are even, as I have explained, reluctant to grant the tortoise an understanding of the conditional, though he might be disposed to carry out ‘a great many’ inferences by modus ponens, but is *unwilling* to carry out a ‘fair number’ of inferences by modus ponens (and not because there is something unique or peculiar about *those* inferences).

⁴³ Boghossian (1999, p. 248)

⁴⁴ See the end of the previous sub-section.

Now, the idea was, McGee is disposed to carry out a great many of the inferences of the form of *modus ponens*, and each of these inferences partially constitutes the meaning of the conditional *for him*. So he needs to have each of these dispositions (and each of the attendant beliefs in the inferences' validities) in order to have *that* concept – of the conditional. Therefore, he's blameless when he carries out any of these inferences without an antecedent warrant for it, since he needs to first have the concept of the conditional in order to have such a warrant.

Unfortunately, this is of course no good. We have already seen why. It is implausible in the extreme to hold that a McGee who did not endorse the goodness of just one inference the real McGee *does* endorse - perhaps restraining himself from sanctioning it in order to reflect first on the question of whether it is good or not (like the tortoise does), perhaps even being persuaded by theoretical considerations that it is *not* valid - would not have the concept of *the conditional*. We would have to hold that subtracting just one inference from the set of inferences which are meaning-constituting of the conditional for McGee, would turn that set from one whose members jointly sufficed to determine a suitable meaning for the conditional, into one whose members did *not*.

Now, another thought might occur to us: if the inference we 'subtract' from McGee's meaning-constituting set is indeed one which serves to partially constitute the meaning of the conditional for him, as we supposed it to be (being a member of his M-C set), then surely if he did *not* have the disposition to carry out that inference, he would not have *the same* concept. But what could we mean here by 'the same' concept. On one reading, the claim made here is just false, and, again, for the reasons we have just seen. McGee would still have the concept of *the conditional* no matter what single inference we subtracted from his M-C set – many of us have *that* concept without each having the same M-C set for it. On another reading the claim made above may be true, but is not to the point in so far as we are interested in whether or not McGee is automatically blameless, for each of the inferences in his M-C set, in carrying them out when he does. As long as he would have the concept of *the conditional* even without having some particular inference in his M-C set, he could in principle have a *warrant* for that inference without having it in that

set; so we could not argue, *in the way that we have* (see below), for the claim that he is automatically blameless in carrying out that inference when he does. Now, perhaps it is true that though he would have the concept of *the conditional* even without having this inference in his M-C set, he would not have, in some sense, ‘precisely the same’ concept in such a case, or not have precisely the same understanding of ‘if, then’. But even if we could make sense of this claim and see that it was true (in which case, note, there would be some, hopefully not insurmountable, difficulties we would have to face up to⁴⁵), all of that would be of no assistance to us here, as I have explained.

The position of the critic of the account at hand is presently as follows. For any given inference in a subject’s M-C set for the conditional, it is not true that she must be willing to carry it out in order to have the concept of the conditional. So when she carries it out, she can in principle have an antecedent warrant for doing so – she is not automatically blameless in that act of inference in the way that Boghossian and I would like to maintain that she is. Presumably, then, she is only really blameless in carrying out that inference if she *does* have an antecedent warrant for it – she ought not to carry it out unless she has such a warrant for doing so.

And yet, if our thinker was not first disposed to carry out a great many inferences by modus ponens, she would not so much as be *able* to have an antecedent warrant for *any* of her acts of inference. I think that there is something incoherent in this position, and I’ll try now to provide some support for this. I certainly don’t think that the critic we have on our hands is yet home and dry.

We will be required to effect some change on our account of why, for any inference which is in a subject’s M-C set for the conditional and which she carries out, the subject is blameless in carrying it out (and believing that it’s valid) when she does, and without

⁴⁵ This is a potential problem for Brandom’s inferentialist semantics. On his view *all* of the inferences involving a given expression which a thinker is committed to (or disposed to carry out), have a role to play in bestowing a meaning on that expression for that thinker. See Brandom (2000, p. 167) for his discussion of the problem and proposal for dealing with it. To paraphrase his formulation of the challenge: The inferential significance of a belief depends on what else one believes. Thus, we have holism. So how do we communicate (if communication is sharing meanings)?

having an antecedent warrant for it. For we have seen that it's not the case that for any inference in a subject's M-C set for some concept, that subject must be disposed to carry out that inference if she's so much as to *have* the given concept. My rejoinder is that if, for any inference a subject carries out, she's required to have an antecedent warrant for it in order to be blameless in carrying it out, then she *can't* have an antecedent warrant for *any* of the inferences she carries out. So it is incoherent to hold that the antecedent of the conditional statement just expressed is true – to hold that is to hold that I can be expected to have an antecedent warrant for any arbitrary inference I carry out, where it would follow from the truth of that claim that I *cannot* have an antecedent warrant for *any* of the inferences I carry out whatsoever.

Take some inference which I carry out and which is in my M-C set for the conditional. The claim that I am irresponsible in carrying it out unless I have an antecedent warrant for it is only coherent if I *can* in principle have an antecedent warrant for it. The initial idea was that I *cannot* have an antecedent warrant for this inference because I need to be disposed to carry this inference out in order to have the concept of the conditional. But we've seen that this is false.

But now, similarly, the claim that, for any *arbitrary* inference I carry out, I am irresponsible in that act of inference unless I have an antecedent warrant for it, is only coherent if, for any *arbitrary* inference I carry out, I *can* have an antecedent warrant for that act of inference. We've seen a defence of the claim that I *can* in principle have an antecedent warrant for any arbitrary inference I carry out. I agree, but think that I can in principle have said warrant for any arbitrary inference I carry out only because I don't *need* an antecedent warrant for any arbitrary inference I carry out in order to be blameless in doing so.

Take, again, some arbitrary inference in my M-C set for the conditional, and imagine that I have carried it out. I could indeed, in principle, have had an antecedent warrant for this act of inference. How? The idea is that I could have held my dispositions to carry out all the other inferences in my M-C set fixed, as it were, and therefore still have had the

concept of the conditional, and in some way have acquired a warrant for the given inference; then carried it out.

How could I acquire such a warrant? The idea must be that I must either justify it by an act of rational insight or justify it *inferentially* (the thought is that I must acquire an *internalist* warrant for the act of inference before I perform it). But we have dismissed the notion of rational insight for several reasons. So I must have an *inferential* justification for the given inference before I carry it out - I must carry out *another* inference in order to justify this one. But, on my opponents view, if I don't have an inferential justification for *this new* inference, then I am not blameless in carrying it out, and therefore certainly not entitled to carry it out. And so on. I can indeed have an antecedent warrant for any arbitrary inference I carry out but only because there are *some* inferences which I am blameless in carrying out *without* having an antecedent warrant for them.

We could argue for this from a slightly different angle. Since I don't have the concept of the conditional until I am disposed to carry out a great many inferences involving 'if, then', we can assume that when I am learning the meaning of the conditional, there is some point (even if it's vague) at which I have that concept and at which I have not justified any of the inferences involving 'if, then' which I am disposed to carry out - because until then I didn't have the concept of the conditional. Fine, my opponent now claims that at this point I can 'start to acquire' warrants for those inferences. I can acquire them, say, one by one in turn when I am faced with an inference, and before I carry it out; because *now* I have the concept of the conditional and for any single one of those inferences, I don't need to be disposed to carry it out in order to have that concept.

But now imagine that I have before me an inference involving 'if, then'. How I will get a warrant for it? Again, I must *inferentially* justify a belief that the inference at hand is valid. That means I must carry out one of the other inferences I am disposed to,

presumably one of those involving 'if, then'⁴⁶, all of which I have at present *not* justified. And the only way I could justify this new act of inference would be to carry out, yet again, another inference which remains to be justified by me. And so on. If it is true that I must *always* have an antecedent warrant for an inference if I am to be blameless in carrying it out, then I can *never* have an antecedent warrant for any inference I carry out. So if we hold the antecedent of the conditional statement just expressed to be true, then we hold that I really ought to accomplish something in particular prior to carrying out any inference, where it would follow from the truth of this claim that I quite simply could *not* accomplish that something prior to carrying out any of my inferences whatsoever.

My idea is that it is incoherent to maintain that I am not entitled to *any* of the inferences I carry out by modus ponens without antecedent warrant. Though for any given inference I carry out a sensible epistemic challenge to my entitlement to that inference can be made out, that is only because such a challenge is presented against a background of *unchallenged* inferences.⁴⁷ I must be disposed to carry out a great many inferences by modus ponens if I am to have the concept of the conditional; so if I am to be able to justify any inference I am inclined to carry out whatsoever. If I had not a great many dispositions, if you like, to carry out modus ponens inferences (at first, unreflectively), then I wouldn't be able to justify any inference at all. So one cannot coherently maintain that I am not entitled to have - that I ought *not* to have - *any* of those dispositions, if you like, on the basis of them being dispositions to infer *unreflectively*, for if I did not have any of those dispositions, I would not so much as be *able* to reflect at all.

Though a sensible epistemic challenge can be leveled against any given modus ponens inference I carry out without antecedent warrant, it is incoherent to maintain that I am not entitled to *any* of the inferences I carry out by modus ponens without antecedent warrant. And I think it follows from this that an opponent cannot indiscriminately challenge my entitlement to any inference whatsoever I carry out by modus ponens without antecedent

⁴⁶ If the inference proceeds by some rule other than modus ponens, we are still left with the question of whether I have an antecedent warrant for that inference; and if I *do*, how I might have acquired a warrant for it without being blameless in employing some argument *without* having an antecedent warrant for it.

⁴⁷ Thanks to Bernhard Weiss for helping me to see the matter in this way.

warrant, since this would commit her to the view that I am not entitled to *any* of the inferences I carry out by modus ponens without antecedent warrant. But of course, she could coherently challenge *some* of the modus ponens inferences I carry out unreflectively. She could challenge such inferences as McGee would deem invalid, since McGee is at least ‘owed an answer’. But she could only coherently challenge my modus ponens inferences which are peculiar in some way – like these of the type McGee would find objectionable – *not* indiscriminately. So, it seems plausible that we could define my M-C set for the conditional as the set of all inferences involving ‘if, then’ I am disposed to carry out, excluding those that are peculiar in some such way as makes it reasonable for me to be expected to have an antecedent warrant for them. This cannot leave me with an M-C set which is ‘insufficiently large’ to bestow a suitable meaning on the conditional for me, because it is not only *unreasonable* to expect me to have an antecedent warrant for so many inferences that if I wasn’t disposed to carry out those inferences I wouldn’t have the concept of the conditional – it is *incoherent*.

I shall turn now to Wright’s argument for (ii) - that an act of inference being epistemically blameless is not sufficient for it being warranted.

II-iv Epistemic Blamelessness is indeed *not* sufficient for Warrant

Wright has us ‘imagine that Frege had invented the *Begriffsschrift* as a system of natural deduction, and that rather than as an axiom, he had gone on to formulate the notorious Basic law V of *Grundgesetze* as a pair of rules respectively for the inferential introduction and elimination of contexts containing a course of values operator’.⁴⁸ We now know that if Frege had indeed formulated such rules, they would have been unsound.

But it took a clever man to find Russell’s paradox. It might not have been noticed for many years, and generations of students might have been trained in what was erroneously take to be the elegant and definitive foundational system that Frege had invented. If that

⁴⁸ Wright (2002, p. 63)

had happened, there is the same clear intuitive sense in which their practices would not have been epistemically irresponsible.

The immediate conclusion to be drawn from [such] examples is that there is a kind of irresponsibility—or rather, lack of irresponsibility—which does not *per se* suffice for warrant. Someone who in this sense responsibly but fallaciously reasons to a true conclusion is not entitled to take it—does not know—that it is true.

Wright's argument here is not that an act of inference being in accordance with a *meaning-constituting* rule does not suffice for it being warranted - the rules for the course of values operator are, being unsound, *not* meaning-constituting of that operator. And he grants for the sake of argument that an act of inference by a meaning-constituting rule is one which is *responsible*, in Boghossian's sense. His point, however, is that even if any act of inference by a meaning-constituting rule *is* a *warranted* act of inference, the explanation of that fact cannot be given 'purely in terms of this notion of responsibility'.⁴⁹ And this is a problem for *Boghossian* because he attempts to provide just such an explanation. Now, we can note that the hypothetical scenarios Wright comes up with⁵⁰ are scenarios in which an *invalid* inference is blamelessly carried out, so the problem is that the blamelessness of an act of inference does not guarantee that inference's *validity*, and therefore does not guarantee that it is *warranted*.

Wright's two objections to Boghossian's account connect with our considerations of previous chapters in the following way. In chapter 3 I said that I wanted to defend the view that we are able to have deductive inferential knowledge that is *certain*. I argued for two constraints on the attainability of such knowledge. For me to be able to have certain knowledge that *C*, on the basis of some item of propositional knowledge of mine, that *P*, it must at least be the case that *P* entails *C*, and also not the case that I ought to doubt that *P* entails *C*. In my strong sense of being entitled to an *act*⁵¹ of inference, which connects nicely with Wright's talk of an act of inference being *warranted*: for entitlement to an

⁴⁹ Wright (2002, p. 63)

⁵⁰ Wright also has us imagine ('if we can') a community having the unreflective communal practice to infer by the 'tonk'-rules.

⁵¹ I say that someone is entitled to a *belief* as long as it's not the case that she ought to doubt it. But, therefore, being entitled, in this sense, to the belief that *P* entails *C* is not sufficient for being entitled to the act of inference from *P* to *C*, because, for *that*, the *truth* of *P* entails *C* is also required.

inference we need both the *validity* of the inference and the *blamelessness* of the act of inference (which amounts also, on my view, to the blamelessness of the *belief* that the inference is valid).

Now, Boghossian's view is that if an inference carried out is a meaning-constituting inference, then it is a warranted inference that can serve to confer warrant on its conclusion so that the inferer can come to *know* the conclusion as a result of carrying out the inference. Why? Because the inferer is blameless in her act of inference in such a case, so at least our second constraint is met. Wright's first objection, (i), was that a meaning-constituting inference is *not* necessarily blameless; and his second, (ii), is that even if it were, it would not be a *warranted* act of inference *just because* it was blameless - the blamelessness of an act of inference is not sufficient for the *validity* of that inference, which is also required if the act of inference is to be warranted (our first constraint). So,

a satisfactory account of the acquisition-condition⁵² needs clauses of *both kinds*: it needs a clause to ensure the soundness of the rules of inference being followed ... and it needs a clause to ensure that practice in accordance with those rules is open to no complaint of irresponsibility ...⁵³

I have argued, in chapter 6, that for me to be entitled to an inference it is necessary and sufficient that I *know* that it is valid. And I think that if a reasoner carries out one of the inferences in her meaning-constituting set for the conditional, say, then she *does* know that inference to be valid. The inference *is* valid because it is partly constitutive of the very meaning of some of the propositions involved in it, as I have explained; and, also, she *believes* that it is valid. Then, against Wright's first objection, I have provided some defence of the view that she is *blameless* in committing such an act of inference, and likewise blameless in holding the belief that it is valid. So what we have in such a case is a true and blamelessly held belief. And while I don't think that *any* true and blamelessly

⁵² 'The minimum condition that must be met by the rules which mediate a particular inference if it is to subserve a thinker's *coming to* knowledge of the truth of its conclusion.' (Wright, 2002, p. 56)

⁵³ Wright (2002, p. 65)

held belief constitutes knowledge – the blamelessness of a belief, even with its truth (as an external fact), is not sufficient for knowledge⁵⁴ – I think that when it is further the case that our subject's holding of the belief is in part what constitutes the meaning of one of its ingredient concepts for her, then we do have a case of knowledge.

Consider the following. Trusting her teacher (which, we might argue, she can hardly be blamed for), a young student follows that teacher in carrying out a *non*-meaning-constituting but *valid* inference. She thus *believes* that the inference is valid, she is plausibly blameless in that belief, and the inference *is* valid. Even given all that, however, this inference of hers is not *warranted*, at least not to the extent required for her to be able to acquire certain knowledge of its conclusion. I would also maintain that she does not know that the inference is valid. But if this same student carries out another inference (by modus ponens, say), and one which is meaning-constituting (of the conditional) for her, then, again, she is blameless in her belief that the inference is valid; and this belief of hers is *true*. But here the truth of the belief is ensured by 'the nature of the concepts it contains', so she is surely 'in the best possible epistemic position with respect to the content of the belief.'⁵⁵ If a subject is to have the concept of the conditional, she must believe a great many propositions involving 'if, then'. And for someone who has that concept, the truths of some of *those* believed propositions are *constitutive* of the meaning of the conditional for her.⁵⁶ She *knows* of the truth of these propositions, or *understands* them to be true, by virtue of having the concept of the conditional; as part of her understanding of 'if, then'.

This account is indeed partly externalist, at least if we are right to characterize the distinction between externalist and internalist accounts in terms of the notion of 'reflectively appreciable' warrant. On my view, if an inference is an element of the

⁵⁴ So the inference's validity together with the fact that it's not the case that you ought to doubt that it is valid, is not sufficient for you being *entitled* to the inference, though both of these are *necessary* for you being entitled to the inference, as we said in chapter 3.

⁵⁵ Peacocke (1992, p. 157-158); See above – (II *How do You Know?*).

⁵⁶ For instance, the believed propositions involving 'if, then' which are hypothetical statements associated with modus ponens inferences in her M-C set for the conditional

meaning-constituting set of some concept for a subject, then she has knowledge of the inference's validity – but she has no *reflectively appreciable* warrant for that belief.

But, in defence of my account, I have argued that it would be *incoherent* (or just highly implausible, if we opt to endorse some rational insight theory) to expect of a thinker that she should have, for *each* of the inferences she carries out, a reflectively appreciable warrant doing so. I argued in the previous section that she can have an antecedent warrant for any arbitrary inference she carries out, but that this is only because there are many inferences she is entitled to carry out *without* an antecedent warrant, maybe including the one she hypothetically has on her hands. So a reasoner can have an antecedent warrant for any given inference she carries out - and by this it is meant that she can in principle have a justified belief which counts as a good reason for carrying out her inference; so, a reflectively appreciable warrant –but cannot coherently be required to have one for every inference she carries out.

Then, I have an account of how a subject who carries out an inference which is meaning-constituting for her *knows* that inference to be valid, and an account which should help to appease the discomfort felt by an internalist at the fact that she often doesn't have a strictly internalist warrant for that belief. For she is blameless in her belief and 'the nature of the concepts it contains' guarantees its truth – that the inference is valid is part of what she believes by virtue of having some of those very concepts which are ingredient in that belief.

Furthermore, the account is in large part properly internalist. Unlike on Boghossian's view, the blamelessness of an act of inference which is a meaning-constituting inference for a thinker amounts also to the blamelessness of her associated *belief* that the inference is valid. And the blamelessness serves her, if you like, as part of what makes that belief, which she *has*, warranted. Also, her warranted and true belief that the inference is valid amounts to *knowledge*, so the fact that the inference is valid is not a 'merely external' fact for her (even if the fact that her *act* of inference is warranted is). Then, her knowledge that the inference is valid is brought to bear on her *act* of inference, for it is on the basis

of that knowledge that she carries out the inference, as I argued in chapter 6. Thus, we have an account of inference which allows us to see how a meaning-constituting inference can allow a reasoner to have certain knowledge of the conclusion of the inference, and to *explain* a thinker's acts of inference – enabling us to see acts of inference as *rational*.

III The Justification of Deduction

In chapter 2 we saw the following rule-circular justification of modus ponens.⁵⁷

Assume that the premises, *if P then C* and *P*, of an MP argument are both true. Now, by the truth-table for the conditional, if the two propositions *if P then C* and *P* are both true, then *C* is true. So *C* must be true.

I shall now consider what position we are in with respect to the question of the acceptability of such an argument. On the one hand, it may be felt that offering such an argument in an attempt to justify modus ponens seems ridiculous, especially given all that has come before. For we have seen that to have a given concept one must be disposed to carry out a great many inferences in which that concept figures. In particular, it is plausible that to have the concept of the conditional one must be disposed to carry out a great many inferences by *modus ponens* – modus ponens is meaning-constituting of the conditional, or such is the hope. Furthermore, we have an explanation of the fact that a meaning-constituting rule *is* truth-preserving. So we would be inclined to seek now a justification for the claim that modus ponens is indeed meaning-constituting of the conditional, so that we could draw the conclusion that modus ponens is indeed truth-preserving. By this, we would have an explanation of the fact, and justification for the belief, that modus ponens is truth-preserving, and one which was surely by far more

⁵⁷ Chapter 2 (

satisfactory than the justification we *might* be able to have as a result of the argument given above.

But, having said this, it is important to note that to show that *if* a rule is meaning-constituting *then* it is truth-preserving, and then that modus ponens is indeed meaning-constituting and infer that it is truth-preserving, is of course to provide an argument which is, again, rule-circular. If we think that the argument given at the start of this section is unacceptable simply because it is rule-circular, then, since *any* explanation of the fact that modus ponens is truth-preserving will surely make use of that very rule⁵⁸, we cannot hold out hope for any such satisfactory explanation. As Boghossian⁵⁹ explains, for instance, it would be confused to think that, because of the (supposed) problems of rule-circularity, we need an *empirical* justification for modus ponens - *any* justification of our 'core logical principles' would have to make use of those very principles. So, if we accept 'the ban on the use of a logical principle in reconstructing our a priori warrant for that very principle, we would have to conclude that there can be no such reconstruction'.⁶⁰

It seems to me to be true that by a conceptual role semantics we can have a far deeper and more satisfying explanation of the soundness of modus ponens, but this does not take away from the fact that *any* such explanation must face up to the problems we have seen raised for rule-circular justifications. In light of this, it is by no means a case of "two steps forward, one step back" to ask whether an argument such as the one given above (at the start of this section) *could* grant a subject knowledge of the soundness of modus ponens, and knowledge which she previously lacked. For instance, as Wright points out⁶¹, it would come very naturally to offer some such rule-circular justification to a student who queried the soundness of modus ponens. So, could an argument of that kind indeed allow such a student to come to know that modus ponens *is* truth-preserving?

⁵⁸ Or, at least, some other basic rule of inference, which would not make any real difference.

⁵⁹ Boghossian (1999, p.232-234)

⁶⁰ Boghossian (1999, p. 253)

⁶¹ Wright (2002, p. 55)

III-i Begging the Question

As we saw in chapter 2, a *grossly* circular argument presupposes its conclusion by including it as (or as ‘part of’) one of its premises. The employment of such an argument, therefore, cannot serve to *confer* warrant on its conclusion because, for *that*, a subject needs to already *have* warrant for believing the premises of an argument – and in such a case this involves already having a warrant for believing the *conclusion*.

But, as Boghossian points out, an argument does not seem to rely on ‘its implicated rule of inference in the same way ... [as] it relies on its premises.’⁶² Also, and relatedly, an argument asks it to be granted that *one* application of modus ponens is warranted and ‘promises to deliver the conclusion that [modus ponens] is *necessarily* truth-preserving’.⁶³

In chapter 2 we saw Wittgenstein’s claim in the *Tractatus* that

[t]he nature of the inference [from one proposition to another] can be gathered only from the two propositions. They themselves are the only possible justification of the inference. ‘Laws of inference’, which are supposed to justify inferences, as in the works of Frege and Russell, have no sense, and would be superfluous.⁶⁴

To counter the charge that a *rule*-circular argument begs the question, we would need to show only that a reasoner does not need to have a warrant to believe that the *rule* according to which an inference proceeds is truth-preserving in order to be entitled to that inference.⁶⁵ If she did need to know or have warrant to believe that, then the employment of a rule-circular argument would not allow her to acquire a warrant for the belief in the conclusion that she did not already have; but if she did not, then there is no way in which

⁶² Boghossian (1999, p. 245)

⁶³ Boghossian (1999, p. 252)

⁶⁴ Wittgenstein (1922, 5.132)

⁶⁵ See chapter 2 (II-ii Begging the Question, p. 12).

she has begged the question by her provision of the rule-circular argument.⁶⁶ Thus, Boghossian distinguishes between being disposed to reason according to modus ponens and having the ‘full-blown’ belief that modus ponens is necessarily truth-preserving.⁶⁷ As we have seen, Boghossian holds that anyone who has the concept of the conditional is indeed disposed to carry out any given inference of the form of modus ponens (whose premises and conclusion she can understand), but he maintains that having this disposition does not involve having the ‘full-blown’ belief that modus ponens is necessarily truth-preserving.

Indeed, Boghossian maintains that having a disposition to infer and having a *belief* are ‘distinct kinds of states’, and that even if we think, as I do, that having the disposition to carry out some particular inference involves believing *that* inference to be valid, we are led into Carrollian regress.

We have seen that I differ from Boghossian, firstly, in that I think that having the disposition to carry out some particular inference *does* involve believing *that* inference to be valid, and that, as I argued in chapter 6, we are not led into regress by remaining faithful to this pre-theoretic intuition. Secondly, we have seen that a reasoner might have the concept of the conditional while there are *some* modus ponens inferences she is *not* disposed to carry out. This observation helps us to argue that someone might carry out an inference by modus ponens without having the belief that *all* instances of modus ponens are valid. This is important, firstly, because if it were false then it would be difficult to explain how a subject’s belief that modus ponens is truth-preserving figures in a *given* modus ponens inference of hers, except with the proposition that modus ponens is truth-preserving as a premise from which she infers that the particular inference at hand is valid – which would indeed lead to regress; and, secondly, because we can now follow Boghossian in providing a defence of rule-circular argument against the charge of begging the question.

⁶⁶ See Wright (2002, p. 56)

⁶⁷ Boghossian (1999, p. 230)

But, though it helps me to observe that someone might have the concept of the conditional without being disposed to carry out *any* given inference of the form of modus ponens, I must meet a possible objection to the effect that believing that some particular modus ponens inference is valid involves believing that the *rule* of modus ponens is truth-preserving. Susan Haack, in arguing against the view, distinct from my own, that we can ‘justify [a given] form of argument by appeal to the validity of its instances’, writes that

there is an *implicit generality* in the claim that a *particular* argument is valid. For to say that an argument is valid is not just to say that its premisses and its conclusion are true—for that is neither necessary nor sufficient for (semantic) validity. Rather, it is to say that its premisses *could not* be true without its conclusion being true also, i.e. that *there is no argument of that form with true premisses and false conclusion*.⁶⁸

If Haack is right in her claim that to say that a particular argument is valid is to say that ‘there is no argument of that form with true premisses and false conclusion’, then if someone who carries out an argument (of the form of modus ponens, say) necessarily believes that it is valid, then she also necessarily believes that *modus ponens* is truth-preserving.

But I just can’t see why to say, or believe, of some particular inference that its premisses *could not* be true without its conclusion being true also is to say that *there is no argument of that form with true premisses and false conclusion*. Indeed, the fact that someone such as McGee might believe that some given modus ponens inference is valid, and yet believe that another is *not* valid, and that modus ponens is *not* necessarily truth-preserving, seems clearly to refute this suggestion.

Furthermore, I’ve argued that anyone who carries out an inference believes that it is valid, but that it’s certainly not generally the case that she *infers* that it is valid from something else which she believes. I’ve made a case for the view that even a child with no logical ability who carries out an inference, by modus ponens, say, believes that the

⁶⁸ Haack (1976, p. 118)

inference is valid. But that such a child surely has no *reason* or *reasons* - such as that the argument is of the form of modus ponens, and all inferences of that form are valid - for believing that her inference is valid.

Now, Boghossian explains a subject's entitlement to an *act* of inference (where the inference is meaning-constituting for her) without reference to her belief that it is valid. She is blameless in the act of inference because, on his view, she must be disposed to carry it out in order to have one or some of the concepts involved in that inference.

We have refined this account to some extent, but the basic point is still the same. On my account, to believe that a give inference is valid it is sufficient that I carry it out or am disposed to do so – for belief and action are intimately connected. Then, if the inference is one of those which is meaning-constituting of, say, the conditional for me, then my belief that it is valid is warranted in the way I've explained, and true. In fact, I know that the inference is valid, and this is *sufficient* for me being entitled to that act of inference.

Since none of this involves me in having the belief that *all* instances of modus ponens are valid, or a warrant for such a belief, there is no reason for a rule-circular argument employed in an attempt to *acquire* such a warrant to be unacceptable as far as the charge that it is question-begging is concerned.

III-ii Bad Company

In chapter 2 we saw two rule-circular arguments purporting to justify, respectively, the introduction rule of Arthur Prior's connective, 'tonk', and Susan Haack's rule, modus morons – two patently *unjustifiable* rules of inference.⁶⁹ These arguments are analogous to the rule-circular argument we have above for the justification of modus ponens, so the charge is that this argument of ours 'keeps bad company' – if we allow rule-circular

⁶⁹ Chapter 2 (

II-i Bad Company, p. 10)

argument then we can justify such obviously unsound rules as the tonk rules and modus morons, by rule-circular arguments. But the matter is not quite so straightforward as that. Of course, 'we cannot simply assert that rule-circular justifications are acceptable and leave it at that,' Boghossian points out. '[T]he question is whether there is some intuitively plausible constraint that they should be made to satisfy that will repel the bad company they would otherwise keep while leaving in place the justificatory arguments that interest us.'⁷⁰

The plausible constraint that we have on rule-circular justifications is of course that they must proceed by meaning-constituting rules, or by inferences which are meaning-constituting of some concept ingredient in them for the thinker who employs them. And the problem with rule-circular justifications for the tonk rules and for modus morons, for instance, is not that they are rule-circular but that the inferences which a subject carries out when she employs them are not meaning-constituting of any concept for her.

As far as the tonk-rules go, if we thought that they constituted a meaning for 'tonk' then

... we would have to hold that there is a thinkable proposition expressed by sentences of the form 'A tonk B'. If there were such a thinkable proposition, then there would have to be a way the world is when the proposition is true. How, though, must the world be if 'A tonk B' is to be true? Since the sentence is compound, its truth value will depend on the truth values of its ingredient sentences A and B. But we can readily see that there can be no consistent assignment of truth value to sentences of the form 'A tonk B' given the introduction and elimination rules for 'tonk' ... So there can be no determinate way the world has to be, if 'A tonk B' is to come out true.⁷¹

We can 'readily see that there can be no consistent assignment of truth value to sentences of the form 'A tonk B' by attempting to construct a truth-table for 'tonk'. For when 'A' is true and 'B' is false, 'A tonk B' must be *true* because of the connective's introduction rule, and also *false* because of its elimination rule.

⁷⁰ Boghossian (1999, p. 248)

⁷¹ Boghossian (2002, p. 43)

Tonk-introduction:

For any propositions P and C ,

from	P
infer	$P \text{ tonk } C$.

Tonk-elimination:

For any propositions P and C ,

from	$P \text{ tonk } C$
infer	C .

Similarly, the adoption of modus morons, *together with* the introduction and elimination rules already in place for the conditional, would result in the impossibility of a consistent assignment of truth value to conditional statements.

We have seen plausible argument for the claim that an unsound rule is *not* meaning-constituting⁷², and it follows from the truth of this claim that a rule-circular justification of an unsound rule will not serve to confer warrant on its conclusion. For when a subject provides a rule-circular justification of an invalid rule, she infers by a *non*-meaning-constituting rule (the rule for which she seeks a justification) – she carries out an inference which is *not* meaning-constituting of any concept for her. She is therefore *not* blameless in that act of inference in the way that she is when she carries out some inference which *is* meaning-constituting of some concept for her, and not blameless in

⁷² I do not have space to consider the objection that some rules are unsound but meaning-constituting. See Boghossian (2003) for relevant discussion. See also Brandom (1994) and Brandom (2000).

believing that inference to be valid.⁷³ Indeed, her belief that the inference is valid is *false*⁷⁴ – she certainly does not know that it is valid, and is not entitled to carry it out.

IV Conclusion

Though the moral of Lewis Carroll's tale of what the tortoise said to Achilles has more to do with inference and belief than logical validity, a coherent but ultimately unsound argument for the impossibility of logically valid argument can be made out in the story. In its most persuasive (or at least most coherent form), such an argument demonstrates that the hypothetical associated with an argument is never a premise of that same argument; and in reply we can show that it need not *be* a premise of that same argument in order for the argument to be valid.

If we want to counter Descartes' sceptic by maintaining that we are entitled to have no doubt about some of our beliefs, we must hold that some of our beliefs are based on *conclusive* evidence – evidence which *guarantees* their truth. Further, we must contend that some of our beliefs are at least such that it is not the case that we ought to *doubt* that we have conclusive evidence for them. Indeed, it is in keeping with our actual practices of the attribution of entitlement to deductive inference always together with knowledge of the validity of deductive inferences to hold that anyone who is entitled to a deductive inference *knows* that inference to be valid. We have seen, however, by the very examination of our practices relating to inferential activity, that the knowledge an inferer has of the validity of some inference she carries out with entitlement is often *implicit* knowledge. Indeed, such reflections facilitate our appreciation of the fact that there is an intimate tie between propositional and practical knowledge.

For any given belief of some thinker, there is an indefinite number of propositions whose falsity would cast doubt on the truth of that belief, but which she does not need to *know*

⁷³ We might do better to say that she does not in fact believe that it is valid, because some of the (purported) concepts involved in her (purported) belief that her premise entails her conclusion have no content.

⁷⁴ So neither of our constraints of chapter 3 is met – she is blameworthy in her belief that the inference is valid, and the inference is in fact *not* valid.

(or even *believe*) to be true in order to be entitled to said belief. But, we rightfully attribute the belief that a given inference is valid to anyone who carries it out, in order to *explain* her act of inference, which is a *rational* act. Then, for her to be able to have certain knowledge of the truth of her conclusion as a result of that inference, again, the inference must be valid and she must be in the best possible epistemic position with respect to her belief that it is valid.

I have argued that anyone who carries out an inference which is meaning-constituting of some concept for her, can rightfully be said to *know* that inference to be valid. For she is epistemically blameless in holding the belief that it is valid, and the belief is true by the very nature of some of the concepts ingredient in it. Knowing that the inference is valid is part of what she knows by virtue of having one or more of the concepts involved in that inference. It is not that she is able to appreciate that it is valid by some mysterious act of 'rational insight'. Furthermore, when a reasoner has knowledge of the validity of some inference that knowledge has a bearing on her relevant *act* of inference when she carries it out, because *that it is valid* (as she believes it to be) is her *reason* for carrying it out. The proposition that a given inference is valid need not be a premise of any argument for the conclusion of that inference in order for the fact that a thinker knows that it is valid to have a bearing on the question of whether she is entitled to her belief in the conclusion. There is an intimate connection between propositional and practical knowledge, as, once more, is evidenced by our practices in relation to inferential activity, and the upshot of a full appreciation of this fact is that we do not fall into the trap of Carrollian regress.

Though an inferer needs to know that a given inference she carries out is valid in order to be entitled to that act of inference, she need not know that the *rule* by which the inference proceeds is sound. Thus, a rule-circular argument for one of our basic principles of logic is not objectionably question-begging. And we can acquire knowledge of the soundness of a basic rule of logic by a rule-circular argument even though we cannot justify an *unsound* rule of inference by such an argument. That is because an inference carried out by the employment of the first type of argument can be meaning-constituting of some

concept for the relevant inferer, whereas an inference carried out by the employment of the second type cannot be.

If we thought, with Putnam, that our basic principles of logic are “so basic that the notion of *explanation* collapses when we try to ‘explain’” their soundness – that there is “no room for an explanation of what is presupposed by every explanatory activity”⁷⁵ – I hope this essay to have contributed something to the continued attempts of philosophers to demonstrate otherwise.

⁷⁵ Putnam (1979, p. 110)

Reference List

- Alston, W. P. (1989) *Epistemic justification: essays in the theory of knowledge*. Itacha: Cornell University Press.
- Bach, K. (2003) *Accidental truth and would-be knowledge*. Posted by the San Francisco State University at URL [Http://online.sfsu.edu/~kbach/accidtruth.html](http://online.sfsu.edu/~kbach/accidtruth.html). Accessed 25/5/2004.
- Bermúdez, J.L. and Millar, A. (ed.) (2002) *Reason and nature: essays in the theory of rationality*. Oxford: Oxford University Press.
- Black, M. (1970) The justification of logical axioms, in: *Margins of precision* (ed. M. Black). Cornell: Cornell University Press. p.14-22.
- Boghossian, P. (1996) Analyticity reconsidered *Nous*, **30** (3), 360–391.
- Boghossian, P. and Peacocke, C. (ed.) (1999) *New essays on the a priori*. Oxford: Oxford University Press.
- Boghossian, P. (1999) Knowledge of logic, in: *New essays on the a priori*. (ed. P. Boghossian and C. Peacocke). Oxford: Oxford University Press. p. 229-254.
- Boghossian, P. (2002) How are objective reasons possible? in: *Reason and nature: essays in the theory of rationality*. (ed. J. L. Bermudez. A. Millar). Oxford: Oxford University Press. p. 15-48.
- Boghossian, P. (2003) Section 1 – I Boghossian, a section in the article: *Blind reasoning Proceedings of the Aristotelian Society*, **77**, 225-248.
- Boghossian, P. and Williamson, T. (2003) *Blind reasoning Proceedings of the Aristotelian Society*, **77**, 225-248.
- Boghossian, P. (2001) Inference and insight *Philosophy and Phenomenological Research* **63**, p. 633-641.
- Bonjour, L. (1985) *The structure of empirical knowledge*. Cambridge, Mass: Harvard University Press.
- Bonjour, L. (1998) *In defense of pure reason*. Cambridge: Cambridge University Press.
- Brandom, R. B. (1994) *Making it explicit*. Cambridge, Mass: Harvard University Press.

- Brandom, R. B. (2000) *Articulating reasons: an introduction to inferentialism*. Cambridge, Mass: Harvard University Press.
- Brown, D. G. (1954) What the tortoise taught us *Mind*, **63**, 250, 170-179.
- Carroll, L. (1895) What the Tortoise said to Achilles *Mind*, **4**, 278-280.
- Descartes, R. (1971) First meditation p. 61-65. Second meditation p 66-75. Fourth meditation, p. 92-100, in: *Descartes philosophical writings*. (ed. E. V. Anscombe. P. T. Geach). London: Dalro.
- Dilman, I. (1973) *Induction and deduction: a study in Wittgenstein*. Oxford: Basil Blackwell.
- Dummett, M. A. E. (1973) The justification of deduction, in: *Truth and other enigmas* (ed. M. A. E. Dummett). London: Duckworth. p. 215-247.
- Dummett, M. A. E. (1966) Wittgenstein's philosophy of mathematics, in: *Wittgenstein, The philosophical investigations* (ed. G. Pitcher): London: MacMillan.
- Harman, D. (1996). Analyticity regained *Nous*, **30** (3), p. 392-420
- Haack, S. (1976) The justification of deduction *Mind*, **85**, p. 112 – 119
- Haack, S. (1982) Dummett's justification of deduction *Mind*, **91**
- Haack, S. (1974) *Deviant logic*. Cambridge: Cambridge University Press.
- Kyburg, H. (1965) Comments on Salmon's "inductive evidence" *American Philosophical Quarterly*, **2**, p. 274-6.
- Lewis, D. (1996) Elusive Knowledge *Australasian Journal of Philosophy*, **74**, p. 549-567
- McGee, V. (1985) A counterexample to modus ponens *The Journal of Philosophy*, **82** (9), p. 462-471
- Peacocke, C. (2000) Explaining the a priori: the programme of moderate rationalism, in: *New essays on the a priori*. (ed. P. Boghossian. C. Peacocke) Oxford: Oxford University Press. p. 255-285
- Peacocke, C. (1992) *A study of concepts*. Cambridge, Mass: MIT Press.
- Putnam, H. (1979) Analyticity and Apriority: Beyond Wittgenstein and Quine, in: *A priori knowledge* (ed. P. K. Moser): New York: Oxford University Press. 1987, p. 85–111

- Prior, A. (1967). The runabout inference ticket, in: *Philosophical logic*. (ed. P. Strawson). Oxford: Oxford University Press.
- Quine, W. (1996) (ed.) *The ways of paradox and other essays*. End Edition. Mass, Cambridge: Harvard University Press
- Russell, B. (1903). *The principles of mathematics*. Cambridge: Cambridge University Press.
- Ryle, G. (1945-1946) Knowing how and knowing that *Meeting of the Aristotelian Society at the university of London Club on November 5th 1945*, p. 1-16.
- Ryle, G. (1949) *The concept of mind*. Hutchingson: London.
- Ryle, G. (1954) If, so and because, in: *Dilemmas*. (ed. G. Ryle). Oxford, Oxford University Press. p. 301-318
- Reid, T. (1986) *Essays on the intellectual powers of man*, in *Essays*. (ed. A. D. Woozley). New York: Lincoln-Rembrandt Pub.
- Stroud, B. (1979) Inference, belief and understanding *Mind*, **88**, p.178-196.
- Stroud, B. (1984) *The significance of philosophical scepticism*. Oxford: Oxford University Press.
- Stroud, B. (2000) Meaning, understanding and translation, in: *Meaning, understanding and practice*. (ed. B. Stroud). Oxford: Oxford University Press.
- Thomson, J.F. (1960) What Achilles should have said to the Tortoise *Ratio*, **3**, p. 95–105.
- Thomson J.J. (1965) Reasons and reasoning, in: *Philosophy in America*. (ed. Max Black). Cornell: Cornell University Press. p. 282 – 303.
- Weiss, B. (2002) *Michael Dummett*. Princeton: Princeton University Press.
- Williams, M. (1991) *Unnatural doubts*. Oxford: Blackwell Publishers.
- Williamson, T. (2003) Understanding and inference, a section in the article: Blind reasoning *Proceedings of the Aristotelian Society*, **77**, p. 249-93.
- Wisdom, W. A. (1974) Lewis Carroll's infinite regress *Mind*, New Series, **83**, p. 571-573.

Wittgenstein, L. W. (1922) *Tractatus logico-philosophicus*. (translated by C. K. Ogden). London: Routledge and Kegan Paul

Wittgenstein, L. W. (1953) *Philosophical investigations*. (translated by G. E. M. Anscombe). Oxford: Basil Blackwell.

Wittgenstein, L. W. (1956) *Remarks on the foundations of mathematics*. (translated by G. E. M. Anscombe, ed. G. H. von Wright G. E. M. Anscombe). Oxford: Basil Blackwell.

Wittgenstein, L. W. (1974) *On certainty*. Oxford: Basil Blackwell.

Wright, C. (1986) *Inventing logical necessity*, in: *Language, mind and logic* (ed. J. Butterfield). Cambridge University Press. p.187-210.

Wright, C. (2002) *On basic logical knowledge*, in: *Reason and nature: essays in the theory of rationality*. (ed. J. L. Bermúdez. A. Millar). Oxford: Oxford University Press. p. 48-76.