

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Please complete the following:

Full name:

University ID number:

2. THESIS DEPOSIT

2.1 Under your registration at the University, you are required to deposit your thesis with the University in BOTH hard copy and in digital format. The digital copy should normally be saved as a single pdf file.

2.2 The hard copy will be housed in the University Library. The digital copy will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.6 below), this will be made immediately openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service.
[At present, theses submitted for a Master's degree by Research (MA, MSc, LL.M, MS or MMedSci) are not being deposited in WRAP and not being made available via EthOS. This may change in future.]

2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the thesis **in excess of two years**. This must be applied for when submitting the thesis for examination (further information is available in the *Guide to Examinations for Higher Degrees by Research*.)

2.4 If you are depositing a thesis for a Master's degree by Research, the options below only relate to the hard copy thesis.

2.5 If your thesis contains material protected by third party copyright, you should consult with your department, and if appropriate, deposit an abridged hard and/or digital copy thesis.

2.6 Please tick one of the following options for the availability of your thesis (guidance is available in the *Guide to Examinations for Higher Degrees by Research*):

- Both the hard and digital copy thesis can be made publicly available immediately
- The hard copy thesis can be made publicly available immediately and the digital copy thesis can be made publicly available after a period of two years (*should you subsequently wish to reduce the embargo period please inform the Library*)
- Both the hard and digital copy thesis can be made publicly available after a period of two years (*should you subsequently wish to reduce the embargo period please inform the Library*)
- Both the hard copy and digital copy thesis can be made publicly available after _____ (insert time period in excess of two years). **This option requires the prior approval of the Chair of the Board of Graduate Studies (see 2.3 above)**

The University encourages users of the Library to utilise theses as much as possible, and unless indicated below users will be able to photocopy your thesis.

I **do not** wish for my thesis to be photocopied

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following:

- Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. **DECLARATIONS**

I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is either the same version as the final, hard-bound copy submitted in completion of my degree once any minor corrections have been completed, or is an abridged version (see 2.5 above).
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.
- At such time that my thesis will be made publically available digitally (see 2.6 above), I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.
- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.

5. **LEGAL INFRINGEMENTS**

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and ensure it is bound into the final hard bound copy of your thesis, which should be submitted to Student Reception, Senate House.

Student's signature: Date:

Enquiry and the Value of Knowledge

by

Barnaby Walker

A thesis submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy in
Philosophy

University of Warwick, Department of Philosophy

October 2015

Contents

Acknowledgements and Declaration	iii
Abstract	iv
Introduction	1
1. Williams on Knowledge and the Enquirer	2
2. Enquiry and the Value of Knowledge	16
3. Chapter Summaries	22
4. Summary	27
1. The Activity Question	29
1.1 Williams' Challenge.....	30
1.2 Williams' Solution	34
1.3 The Problematic Reasoning	42
1.4 Some Inadequate Responses to the Problematic Reasoning	47
1.4.1 Is True Belief Just as Good as Knowledge?.....	47
1.4.2 Is True Belief Necessary but Insufficient for Knowledge?.....	54
1.4.3 Does (3) Follow From (2)?	62
1.5 Does (4) Apply in the Case of True Belief and Knowledge?	66
1.6 Summary	70
2. Believing, Assuming and Judging	72
2.1 Enquiry and Belief	75
2.2 Belief and Judgement.....	79
2.3 Judging, Supposing, and the Acquisition/Manifestation of Belief	84
2.4 The Teleological Conception of Cognitive Mental Acts	89
2.5 Supposition, Judgement and Self-Imposed Constraints.....	95
2.6 Believing that p and Assuming that p for Practical Purposes	98
2.7 Summary	108
3. Guessing	110
3.1 Uses of 'Guess'	111
3.2 Sorensen's Account of Guessing	113
3.3 Guessing and its Standard of Success	115
3.4 Guessing, Believing and Knowing.....	122
3.5 Summary	133
4. How to Answer the Activity Question	135
4.1 The Problematic Reasoning Again.....	137
4.2 Thesis (T)	139

4.3 Why Principle (4) Doesn't Apply in the Case of True Belief and Knowledge.	146
4.4 Deciding to Believe.....	157
4.5 Concluding Remarks.....	159
 Bibliography	 160

Acknowledgements

First and foremost I would like to thank my supervisor, Matthew Soteriou, for his support throughout the writing of this thesis. I first encountered Matt as a first-year undergraduate in philosophy, and ever since then he has played a formative role in my intellectual development. When I was still fishing around for a topic for my PhD, he counselled me to think seriously about Williams' discussion, in *Descartes*, of the aim of enquiry, and how that discussion might bear on issues to do with the value of knowledge. This started me down the path which ultimately led to the writing of this thesis. Matt provided me with advice, assistance and encouragement at every step along the way. It has been an absolute privilege to be supervised by him.

I would also like to express special thanks to Naomi Eilan. Without Naomi's intervention (it is not too strong a word), I would never have become a postgraduate philosophy student at Warwick. She offered me a funded place on the MPhil programme at Warwick following a failed (and misguided) attempt to begin an MA at Bristol University, and was subsequently instrumental in securing funding for me to do a PhD. I would like to take this opportunity to express my gratitude for her ongoing personal, philosophical and professional support, and for her friendship.

Quassim Cassam commented on an earlier version of Chapter 3 at MindGrad 2013, and raised objections that caused me to substantially rewrite it. Joe Cunningham, Hemdat Lerman, Guy Longworth and Johannes Roessler gave me helpful comments at various stages in the writing of this thesis.

Wesley Chai and Elly Warner-Kennedy kept me going in the early stages of my PhD studies, during a difficult period in my personal life. Without them, I doubt this thesis would exist. My mum and dad were, as always, an essential source of support; my debt to them is literal as well as metaphorical. Amy's love and care during the final 18 months of this thesis brought me much happiness during a stressful time. Becky and Fergus made 11 Victoria Road into a real home. Michael Van Wijk served as a constant reminder not to take life (or philosophy) too seriously.

I dedicate this thesis to my sister, Molly. Molly studied philosophy at Warwick before me, and was the person who initially got me interested in the subject. Financial hardship meant she was unable to complete her postgraduate studies. Without her example, I would never have completed mine.

Declaration

This thesis is my own work and has not been submitted, in whole or in part, for a degree at another university.

Abstract

Philosophical discussion of the value of knowledge, inspired by Plato's seminal discussion in the *Meno*, typically focuses on the question why it is better to know that p than to have a mere true belief that p . This question is notoriously difficult to answer in a satisfactory way. I argue that the difficulty we experience in trying to solve this problem is a symptom of the fact that we are approaching issues about the value of knowledge in the wrong way. Beneath the traditional problem there lurks a more fundamental issue about the aim of enquiry, namely, why should an enquirer who wants the truth about whether p aim to find out (i.e. acquire knowledge of) whether p , and not merely aim to arrive at a true belief about whether p ? Identifying respects in which knowledge is superior to mere true belief is only one way of trying to answer this question, and, I argue, it is difficult to see how this approach to the question can succeed. An alternative is called for. Central to my alternative proposal is the idea that an enquirer will not have arrived at so much as a belief about whether p until he takes himself to have acquired knowledge of whether p . It is because this is so that an enquirer cannot make life easier for himself by merely aiming to arrive at true beliefs instead of knowledge. I justify this proposal by developing an account of belief according to which outright belief involves a disposition to judge that p , where judging that p is distinct from merely supposing that p for the sake of argument or guessing that p .

Introduction

The basic claim of this thesis is that to resolve philosophical issues about the value of knowledge, we need to pay attention to the position of an enquirer, i.e. someone seeking the truth about the answer to a question. This claim stands in marked contrast to the traditional approach to these issues. According to that approach, rather than thinking about the position of an enquirer, we should instead compare and contrast two subjects, one of whom knows and the other of whom is in some ‘lesser’ cognitive state (e.g. that of having a mere true belief), and try to identify respects in which the former subject is in a superior position to the latter. It turns out to be very difficult to do this in a satisfactory way. But is this the right approach to adopt? What do we presuppose in setting up issues about the value of knowledge in this way? I will be contending that the traditional approach to these issues is misguided, and that the correct approach to the value of knowledge is one that brings questions about the enquirer to the fore.

In this introduction, I want to set the stage for the arguments to come by doing three things. Firstly, I will introduce the idea that attention to the position of the enquirer is required to resolve philosophical puzzles about knowledge in more detail, by tracing its origin in the work of Bernard Williams. Aspects of Williams’ work will receive further discussion in Chapters 1 and 4. Secondly, I will argue that, when thinking about the value of knowledge, it is crucial to clearly distinguish two questions. The first, which corresponds to the traditional approach, is the question of why it is better to know that p than it is to merely be in some ‘lesser’ cognitive state with respect to the proposition that p —e.g. that of having a mere true belief that p . The second is the question of why an enquirer should want to know the answer to a question, and not merely to arrive at some ‘lesser’ cognitive state with respect to the answer to that

question—e.g. that of having a true belief about the answer. Concerning the relation between these questions, I wish to make two major claims. The first, as I have already indicated, is that we really do have two distinct questions here, and not merely two different ways of expressing the same question. My second claim is that the activity question is the more basic question in the sense that, in advance of identifying specific respects in which knowledge is superior to mere true belief, the only general reason we have for thinking that knowledge must be superior to mere true belief is that this claim is required to explain why knowledge, and not merely true belief, is the aim of enquiry. Thus, if it is possible to answer the activity question in a satisfactory way without appealing to the idea that knowledge is superior to mere true belief, we are relieved of the pressure of having to identify respects in which knowledge is better than mere true belief. Thirdly and finally, I will conclude the chapter with an overview of the argument of the thesis by providing brief synopses of the central claims of the remaining chapters. This should prepare the reader for what is to come.

1. Williams on Knowledge and the Enquirer

In two papers from the early 1970s, ‘Deciding to Believe’ and ‘Knowledge and Reasons’, Bernard Williams draws attention to what he calls ‘the examiner situation’, and make a series of bold claims about the role of this situation in philosophical reflection about knowledge. In ‘Deciding to Believe’ the introduction of the examiner situation is prompted by discussion of a machine which, Williams claims, could be attributed knowledge but not beliefs. Having made this claim about the machine, Williams provides the following in the way of comment:

This [i.e. the claim that the machine knows without believing] goes against what is a rather deep prejudice in philosophy, that knowledge must be at least as grand as belief, that what knowledge is, is belief plus quite a lot; in particular, belief together with truth and good reasons. This approach seems to me largely mistaken. It is encouraged by concentrating on a very particular situation which academic writings about knowledge are notably fond of, that which might be called the *examiner* situation: the situation in which I know that p is true, this other man has asserted that p is true, and I ask the question whether this other man really knows it, or merely believes it. I am represented as checking on someone else's credentials for something about which I know already. That of course encourages the idea that knowledge is belief plus reasons and so forth. But this is far from our standard situation with regard to knowledge; our standard situation with regard to knowledge (in relation to other persons) is rather that of trying to find somebody who knows what we don't know; that is, to find somebody who is a source of reliable information about something. In this sense the machine could certainly know something. Our standard question is not 'Does Jones know that p ?' Our standard question is rather 'Who knows whether p ?' (1973: 146)

When we try to find somebody who knows what we don't know so that we can find it out from them, we are seeking the truth about the answer to that question, and so are engaged in enquiry. Thus, although Williams doesn't put the point in quite these terms, the proposal he is making is that reflection on the position of the enquirer acts as a corrective to the mistaken approach to knowledge, encouraged by focussing on the examiner situation, according to which 'knowledge must be at least as grand as belief, that what knowledge is, is belief plus quite a lot; in particular, belief together

with truth and good reasons'. There is a parallel between the claims Williams makes here and the claims I wish to make about the value of knowledge. According to Williams, issues about the nature of knowledge are traditionally approached in a way that neglects the position of the enquirer, and this has led to mistaken claims about the nature of knowledge. In order to correct these errors, we need to reorientate our thinking towards the enquirer. I wish to claim that something similar, but more radical, holds of philosophical thinking about the value of knowledge. Issues about the value of knowledge have been approached in a way that ignores the perspective of an enquirer, and this has distracted philosophers from the fundamental problem in this area.

My argument does not depend on any of the specific claims that Williams makes in the quoted passage, but it is independently interesting to consider the plausibility of Williams' proposals. To begin with, it should be noted that in his description of the mistaken approach Williams runs together two different claims. The claim that 'knowledge must be at least as grand as belief' is naturally understood to express the widely accepted thesis that believing that p is necessary for knowing that p . The claim that 'what knowledge is, is belief plus quite a lot; in particular, belief together with truth and good reasons' is much stronger: it says not only that believing that p is necessary for knowing that p , but that it is possible to give a non-circular analysis of knowing that p as a conjunction of believing that p with further factors—in particular, truth and good reasons. It is possible to accept the first claim without accepting the second; this is the position of Williamson (2000).¹ Once the two claims are distinguished, we can raise two questions about Williams' account of the role of the examiner situation in philosophical thinking about the nature of knowledge.

¹ There is some further discussion of this point in Chapter 1, § 3.2.

Firstly, has a tendency to concentrate on the examiner situation encouraged the view that believing that p is necessary for knowing that p ? Secondly, has a tendency to focus on the examiner situation encouraged the stronger view that it is possible to give a non-circular analysis of knowing that p as a conjunction of believing that p with further factors?

I shall consider only the first of these questions here. The claim that we can give a non-circular analysis of knowing that p as a conjunction of believing that p with additional factors presupposes that believing that p is necessary for knowing that p : if believing that p is not even required for knowing that p , then we obviously can't analyse knowing that p as a conjunction of believing that p with further elements. Thus, if the former view has not been encouraged by focussing on the examiner situation, neither has the latter. Furthermore, in both 'Deciding to Believe' and 'Knowledge and Reasons', Williams' discussion of the role of the examiner situation in philosophical thinking about knowledge is intended to bolster his rejection of certain widely accepted necessary conditions on knowing that p , by providing a diagnosis of why these allegedly spurious conditions on knowing may nevertheless appear to be genuine. Whilst 'Deciding to Believe' rejects the view that knowing requires believing, 'Knowledge and Reasons' is concerned to show that 'it is possible for A to know that q without its being the case that A can rehearse reasons, or at least adequate reasons, for q ' (2006: 50). Thus, in considering whether preoccupation with the examiner situation has encouraged mistaken views about the conditions that are necessary for knowing that p , we also fasten onto the issue that is Williams' target.

Is it the case, then, that a tendency to focus on the examiner situation has encouraged the view that believing that p is necessary for knowing that p , as Williams alleges?

To answer this question, attention is needed to how Williams specifies the examiner situation. In ‘Deciding to Believe’ the description of the examiner situation contains a peculiar emphasis on assertion: the situation is the one in which ‘I know that p is true, this other man has asserted that p is true, and I ask the question whether this other man really knows it, or merely believes it’. This emphasis jars with the claim that academic writings on knowledge are ‘notably fond’ of the examiner situation: if memory serves, most articles on the nature of knowledge, at least in the post-Gettier literature, do not ask us to imagine that Smith, Fred Jones (or whoever else) has asserted that p . Why, then, does Williams build this into his specification of the examiner situation? Presumably, he is motivated here by the view, expressed earlier in ‘Deciding to Believe’, that ‘our very concept of assertion is tied to the notion of deciding to say something which does or does not mirror what you believe’ (1973: 146). If the examinee is described as having asserted that p , then it is being taken for granted that the examinee is a believer; and given that the sincerity of the examinee’s assertion is not in question, it is further being taken for granted that the examinee believes that p in particular. Thus, focussing on the examiner situation, so understood, does introduce into our thinking about knowledge an extraneous element—namely, the idea that the person who potentially knows that p has sincerely asserted that p —that would encourage the view that knowing that p requires believing that p even if that view were, in fact, false. As I have pointed out, however, the evidence simply does not bear out the idea that the examiner situation, so understood, is one that ‘academic writings about knowledge are notably fond of’. Williams’ diagnosis of the source of the view that knowledge requires belief is accordingly unconvincing.

We cannot get around this problem by omitting from the description of the examiner situation the idea that the subject has asserted that p , as Williams does in ‘Knowledge and Reasons’. There, the ‘examiner’s situation’ is introduced as ‘the case where A is admittedly convinced of the truth that p , and the question is whether that conviction is adequately based’ (2006: 48).² It is true, I think, that philosophical writings on knowledge have been preoccupied with the examiner situation, so understood. But this is just to say that epistemologists have tended to take it for granted that knowledge requires true belief, and have then enquired into what further conditions must be satisfied for the subject not only to truly believe, but to know. This can hardly be offered as a *diagnosis* of why philosophers have been attracted to the view that ‘knowledge must be at least as grand as belief’: it is because they are *already convinced* that knowledge requires true belief that epistemologists have focussed on the examiner’s situation, so understood. Williams is caught in a dilemma here. If he includes the idea that the subject has asserted that p in his specification of the examiner situation, then it is false that writings on knowledge have tended to focus on the examiner situation, so this cannot be what explains the currency of the allegedly mistaken view that knowledge requires belief. If, on the other hand, he omits assertion from the specification of the examiner situation, the situation reduces to the one in which the subject has a true belief, but then there is no extraneous element present in the situation that could misleadingly encourage the view that knowledge requires belief. Philosophers will have focused on that situation only because they are independently persuaded that knowledge requires true belief.

² I do not mean to imply that Williams omits assertion from his specification of the examiner’s situation in ‘Knowledge and Reasons’ in a vain attempt to avoid this problem. On the contrary, the difference is presumably due to the fact that, in that paper, Williams is not concerned to dispute the view that knowledge requires belief, but only the view that knowledge requires the subject to be able to rehearse reasons in favour of the truth of his belief.

What of the claim that knowledge requires the subject to be able to rehearse reasons in favour of the truth of his belief? Has this view been encouraged by a preoccupation with the examiner situation? The problem here is that the ability of an examinee to rehearse reasons in favour of the truth of his belief seems to be at most one way in which we might check that he really knows, and is not merely guessing. Suppose, for example, that we ask our examinee to state the product of 12 and 23; after a moment, he responds '276'. Does he know that the product of 12 and 23 is 276, or was his answer just a lucky guess? One way to settle this question would be to request reasons, by asking (e.g.) 'how do you know that $12 \times 23 = 276$?' If the examinee was able to explain how he worked out the answer, this might convince us that he really knows. But this is not the only way of checking whether he knows. Alternatively, we might ask the examinee to solve a series of similar arithmetical problems. Once he has got a sufficient number of these further problems right, this might convince us that he knows the answers to the questions—the odds of him getting so many of them right by sheer chance are so small that the possibility can safely be ignored.³ Here it might be objected that the ability of the examinee to answer related questions correctly convinces us that he knows only because it convinces us that he must be able to rehearse reasons in favour of his arithmetical beliefs. This, however, is dubious. Imagine it turned out that, although he was able to answer all of our questions correctly, the examinee was apparently unable to rehearse reasons in favour of his arithmetical beliefs; when asked 'how do you know that $12 \times 23 = 276$?', he responds by saying 'It just is', or 'Isn't it obvious?'. Would we now be inclined to judge that he doesn't really know that $12 \times 23 = 276$, despite his evident ability to answer a range of arithmetical questions correctly? I don't think

³ A similar case has been taken to impugn the claim that knowledge requires belief; see Radford (1966).

that we would; rather, I think that we would be inclined to judge that this person has a special way of knowing the answers to such questions, that he can ‘just see’ the answers to relatively complex arithmetical problems in something like the way ordinary people can see, without conscious calculation, that $5 - 3 = 2$, or that $3 + 1 = 4$. Even if one disagrees with this verdict, it must at least be admitted that reflection on the examiner situation does not specially encourage the view that knowledge requires the subject to be able to rehearse reasons for the truth of what he believes: the pull of the view that knowledge requires only the satisfaction of what Williams calls ‘external conditions’ can also become apparent in the examiner situation.

So far, I have considered only Williams’ diagnostic claim that preoccupation with the examiner situation has encouraged certain mistaken views about what is required for knowledge—in particular, the view that knowledge requires belief, and the ability to rehearse reasons in favour of the truth of one’s belief. I have been arguing that a tendency to focus on the examiner situation does not provide a plausible diagnosis of either view. The other half of Williams’ proposal is that these views cease to be plausible once we turn our attention to the position of an enquirer who wants to find out whether p , and who is looking for someone who knows whether p . This second part of his proposal is, on the face of it, independent of the first part; it remains to be seen, then, whether we should endorse what he has to say here.

Regarding the claim that attention to the enquirer discredits the view that knowledge requires the subject to be able to rehearse reasons in favour of the truth of what he believes, the following passage, from Edward Craig’s book *Knowledge and the State of Nature*, is apposite:

Why should we want more of a potential informant than that his views on the point at issue should be true, and at least confident enough for him to be prepared to come out with them? Then we come to hear the truth, which is what we wanted. But this overlooks a crucial point. It is not just that we are looking for an informant who will tell us the truth about *p*; we also have to be able to pick him out, distinguish him from others to whom we would be less well advised to listen. How is that to be done? Well, it will be easy enough to find out what he believes about *p*; and if we ourselves knew whether *p* that would suffice to tell us whether he has a true belief. But *ex hypothesi* we do not know whether *p*—we are in the position of inquirers, not of examiners (to borrow Bernard Williams's way of putting it); the informant is to be our means of access to that knowledge, and if we already had it, we would not be inquiring. Obviously, we have to detect the right informant without benefit of prior knowledge. So we need some detectable property—which means detectable to persons to whom it is not yet detectable whether *p*—which correlates well with being right about *p*; a property, in other words, such that if the informant possesses it he is (at least) very likely to have a true belief on that matter. (1990: 18–19)

What could this detectable property be? One answer is that the informant must be able to provide adequate reasons in support of his belief. This, one might think, is a property that correlates well with being right about whether *p*, and which the enquirer can detect without being able to detect whether *p*. Thus, reflection on the conditions under which someone will be a good informant about whether *p* for an enquirer who doesn't know whether *p* can encourage the thought that knowledge requires the subject to be able to rehearse reasons in favour of the truth of his belief.

That said, it certainly isn't the case that this is the only detectable property that correlates well with having a true belief about some matter. As Craig goes on to point out, '[i]f you want to know the way it will always be a good idea to ask a taxi-driver' (1990: 26), because being a taxi-driver correlates well with having accurate beliefs about directions. Furthermore, for some classes of beliefs (e.g. beliefs about salient features of one's autobiography, such as where one was born), 'the mere fact of being willing to offer an opinion [is] correlated excellently with being right' (ibid.). It would therefore appear that more careful reflection on the conditions under which someone will be a good informant about whether p should lead us to reject the view that knowledge requires the subject to be capable of rehearsing reasons in favour of the truth of his belief; the ability to rehearse reasons in favour of the truth of one's belief is at most *one* way of fulfilling the more general requirement of having some property, detectable to an enquirer who does not know whether p , that correlates well with having a true belief about whether p .

Although Craig is clearly influenced by Williams' general proposal that reflection on the position of an enquirer is important when it comes to gaining a philosophical understanding of the nature of knowledge, his argument relies on the more specific assumption that, to a first approximation, the conditions under which a person knows whether p are the conditions under the person would be a good informant about whether p . It is only if this more specific assumption holds that one can infer conclusions about what is required for S to know that p from premises about the conditions under which S would be a good informant about whether p . The basis on which Craig advocates this claim is that the concept of knowledge serves a certain function, and this function is to flag good informants. One issue raised by Craig's argument, then, is whether we should regard the concept of knowledge as having a

function, and, if so, whether its function is the one identified by Craig. Further discussion of this issue lies beyond the scope of the present work. However, there is another aspect of Craig's assumption I wish to comment on. Our conception of the conditions under which someone will be a good informant about whether p for an enquirer obviously depends on our conception of what the enquirer is after. If we think that the enquirer is after knowledge of whether p , then a person will not be a good informant about whether p for the enquirer unless the enquirer can come to know whether p on the basis of the person's testimony about whether p . It is significant that Craig's conception of what the enquirer is after falls short of knowledge: in particular, Craig assumes that the aim of an enquirer who wants the truth about whether p is merely to arrive at a true belief about whether p .⁴ Thus, put more explicitly, the assumption his argument relies on is that the conditions under which a person knows whether p approximate to the conditions under the person would be a good informant about whether p for an enquirer seeking a true belief about whether p . It is not at all clear why we should accept this more minimal conception of the enquirer's aim. As Williamson points out, '[i]t is no reply that believing truly is as useful as knowing, for it is agreed that the starting point [i.e. our conception of the enquirer's aim] should be more specific than 'useful mental state'; why should it be more specific in the manner of 'believing truly' than in that of 'knowing'?' (2000: 31, note 3).⁵

If we assume that the enquirer's aim is not merely to arrive at a true belief about whether p , but to acquire knowledge of whether p , the worry is that appealing to the

⁴ Attentive readers will notice that the passage I have quoted from Craig's book therefore contains a slip. He writes that 'we are in the position of inquirers, not of examiners...the informant is to be our means of access to that knowledge, and if we already had it, we would not be inquiring'. This implies that the enquirer's aim is the acquisition of knowledge. However, it is clear both from his explicit earlier statement of the enquirer's aim, and the logic of his overall argument, that his real view is that the enquirer's aim is the acquisition of true beliefs.

⁵ I will make a similar objection to an argument of Williams' in the next chapter.

conditions under which a person would be a good informant about whether p in order to refute the claim that knowledge requires the subject to be able to rehearse reasons in favour of the truth of his belief threatens to be question-begging. The following principle about the acquisition of knowledge via testimony is *prima facie* plausible: for hearer H to acquire knowledge that p via speaker S's testimony that p , S must know that p .⁶ Bearing this principle in mind, imagine now that someone tries to argue in a way analogous to Craig: 'S is a good informant about whether p , but S is unable to rehearse reasons in favour of the truth of his belief that p ; therefore, knowing that p does not require S to be able to rehearse reasons in favour of the truth of his belief that p '. Someone who is convinced that knowledge requires the ability to rehearse reasons in favour of the truth of one's belief can now respond as follows: 'S is a good informant about whether p only if an enquirer can come to know whether p on the basis of S's testimony about whether p , but given the aforementioned principle, this condition will be satisfied only if S himself knows whether p . Since S cannot rehearse reasons in favour of the truth of his belief, by my lights he does not know whether p . Hence, contrary to your argument, S is not a good informant about whether p '. No doubt the person who initially made the argument would now wish to further contest the idea that knowledge requires the ability to rehearse reasons in favour of the truth of one's belief, but the present point is merely that, once we think of the enquirer's aim as being the acquisition of knowledge, to argue that someone might be a good informant about whether p without being able to rehearse reasons in favour of the truth of their belief is tantamount to assuming that knowing whether p does not require the subject to be able to rehearse reasons in favour of the truth of their belief. Reflection on what

⁶ For further discussion of this principle, including purported counterexamples, see Lackey (1999).

makes a good informant is not neutral ground from which we can mount a non-question-begging challenge to internalism.

What about Williams' further claim that considering the position of the enquirer discredits the idea that knowledge requires belief?⁷ His argument for this claim occurs in the final part of the passage quoted earlier from 'Deciding to Believe'. To repeat: 'our standard situation with regard to knowledge (in relation to other persons) is...that of trying to find somebody who knows what we don't know; that is, to find somebody who is a source of reliable information about something. In this sense the machine could certainly know something'. The problem with the argument in this passage is that knowing something is not equivalent to being a source of reliable information about it; a book can be a source of reliable information about a certain subject matter, but a book cannot literally know anything about that subject matter. To assume that these two things are equivalent is to smudge the distinction between knowers and mere sources of information. It seems to me that the plausibility of Williams' argument depends on us overlooking this distinction. It is undeniable that a machine of the sort he describes might be a source of reliable information about something, and, provided we are not mindful of the distinction between being a source of reliable information and being a knower, this might persuade us that the machine knows about that thing. Once this distinction is drawn, however, it is not clear that such a machine can be anything more than a source of information. Consequently, it is not clear that the machine knows, despite being incapable of having beliefs.

In summary, Williams' proposal that a tendency to focus on the examiner situation has led to mistaken claims about the nature of knowledge, and that these views cease

⁷ For further discussion of this point see Craig (1990), especially Chapter II: 12–16.

to be plausible once we turn our attention to the enquirer, seems to me largely without merit. Such a tendency is not the source of the view that knowledge requires belief, and neither does it specially recommend the view that knowledge requires the subject to be able to rehearse reasons in favour of the truth of what he believes. Craig's argument that reflection on the position of an enquirer should lead us to reject the view that knowledge requires the subject to be able to rehearse reasons in favour of his truth of his belief relies on the assumption that the aim of an enquirer who wants the truth about whether p is merely to arrive at a true belief about whether p . It is not clear why we should accept this assumption, and I will argue in the next chapter that we should reject it on the grounds that the assumption generates an insoluble difficulty when it comes to explaining why enquirers engage in the pursuit of knowledge. Williams' suggestion that reflection on the position of an enquirer undermines the view that knowledge requires belief is implausible once we mark the distinction between knowers and mere sources of information. Fortunately, these difficulties for Williams' proposal do not undermine the central claim of this thesis, which is that we need to pay attention to the position of an enquirer to resolve philosophical issues about the *value* of knowledge. Although puzzles about the value of knowledge are closely related to puzzles about the nature of knowledge, issues about why knowledge matters to us are clearly distinct from issues about what knowledge is. Even if one is sceptical about the suggestion that reflection on the enquirer can help to illuminate issues about the nature of knowledge, then, it is a further question whether it is useful in the context of thinking about the value of knowledge. But how might thinking about the position of an enquirer contribute towards the resolution of the philosophical issues about the value of knowledge? This is the question I shall be taking up in the next section.

2. Enquiry and the Value of Knowledge

In thinking about whether reflection on an enquirer can help to resolve issues about the value of knowledge, it is helpful to return to Plato's original presentation of a problem about the value of knowledge in the *Meno*. In that dialogue, Socrates famously compares a guide who knows the way to Larissa with a guide who merely has a true belief about the way. He claims that the latter guide is just as good as the former, and, more generally, that true belief 'is just as good a guide as knowledge, when it comes to guaranteeing correctness of action' (97b). Furthermore, true belief is, of course, just as true as knowledge. These observations provoke puzzlement in Meno:

All this is making me wonder, Socrates, why, if this is so, knowledge is so much more highly valued than true belief and on what grounds one can distinguish between them. (97c—97d)

What does Meno mean in asking why 'knowledge is so much more highly valued than true belief'? There are at least two different ways of understanding this part of Meno's question. On one interpretation, Meno is here alluding to the fact that human beings desire and seek knowledge, and not merely true beliefs. (As Aristotle later put it, 'All men by nature desire to know' (*Metaphysics*: I.980a20).) Following this line, we can summarize Meno's question by asking, why do we desire and seek knowledge, and not merely true beliefs, given that true belief is just as true, and just as useful, as knowledge? Call this question *the activity question*. Notice that the activity question is one that concerns us in the role of enquirers, i.e. as subjects engaged in the pursuit of truth. However, the *Meno* problem is often formulated in a way that does not bring in the idea of a subject engaged in enquiry. According to this

second understanding of Meno's question, the problem is to explain why it is better to know that p than to have a mere true belief that p . Call this question *the comparative state question*.⁸ When the problem is expressed in this way, the predicament of an enquirer who neither knows nor truly believes that p , but who wants to know whether p , disappears from view. The exercise becomes one of comparing and contrasting two subjects, one of whom knows that p and the other of whom merely truly believes that p , and trying to identify respects in which the former subject is in a better position than the latter. This second interpretation of the question is encouraged by the subsequent direction of Plato's original discussion, which proceeds in precisely this way. (Plato's solution to the problem is discussed in more detail in Chapter 1 § 4.1.)

Here someone might object that the activity question and the comparative state question express essentially the same problem—the fact that the second question no longer mentions an enquirer, someone who is engaged in the pursuit of truth, is merely a cosmetic difference. After all, they may go on to say, if we can explain why the person who knows that p is in a better position than someone who merely truly believes that p , then we have explained why ordinary human beings desire and seek knowledge, and not merely true belief: there is no mystery in the fact that people prefer what is better to what is worse. Even if this is right, however, it misses the crucial point that this is only *one* way of responding to the activity question. Suppose, for example, that it could be shown that merely desiring or seeking true belief, without desiring or seeking knowledge, was impossible. Then there would be a straightforward answer to the question of why we desire and seek knowledge, and not only true belief, even though true belief is just as true, and just as useful, as

⁸ These names for the two questions were suggested to me by a conversation with Hemdat Lerman.

knowledge—namely, that we cannot merely desire or seek true belief. In answering the activity question in this way, though, we wouldn't have answered the question of why someone who knows that p is in a superior state, other things being equal, to someone who merely truly believes that p . Indeed, answering the activity question in this way is consistent with the admission that knowledge is not in any way superior to true belief. If this is right, then we really do have two questions here, and not merely two different ways of expressing the same question.

Here it may be objected that the supposition that it is impossible to merely desire or seek true belief is very implausible. Can't I now form the intention of acquiring a true belief about where Jane Austen was born, and do a Google search to arrive at one? What could possibly prevent me from doing so? The point that the supposition is implausible is right, but beside the point. My purpose in introducing the supposition was not that it is a serious candidate for truth, but rather to show that the activity question really is distinct from the comparative state question, by giving an example of a claim which, if true, would provide an answer to the activity question without providing an answer to the comparative state question. Once it is clear that we have two questions here, we might wonder what else we take for granted about enquiry, beyond the bare possibility of seeking true beliefs without seeking knowledge, in assuming that the correct way to tackle the activity question is by tackling the comparative state question.

Having distinguished the activity question from the comparative state question, a further issue we can consider is to which question, if either, we should assign priority in a philosophical investigation of the value of knowledge. I have already pointed out that a solution to the activity question needn't issue in a solution to the comparative state question. But it should also be noted that if the activity question is

answered in a way that doesn't issue in a solution to the comparative state question, it is far from clear that we have good independent reasons to think that it will be possible to identify respects in which knowing that p is superior to merely truly believing that p . Suppose, for example, that it turned out that it really was impossible to merely desire or seek true belief without desiring or seeking knowledge. This would explain why we desire and seek knowledge and not merely true beliefs, provided it is granted that we are rightly interested at least in arriving at true beliefs. It would be odd at this point if someone continued to press the question of why someone who knows that p in a superior state, other things being equal, to someone who merely truly believes that p . We already have a perfectly good explanation of why we desire and seek knowledge of the answers to the questions that interest us, and not merely true beliefs about the answers; why think it will be possible, in addition, to identify features of knowledge in virtue of which the state of someone who knows is superior to the state of someone who merely truly believes? What is it that is meant to be justified or explained by reference to such features? There is no obvious answer to these questions. Note that the claim here is not that there are no respects in which knowledge is superior to mere true belief. The claim is only that, in advance of identifying specific respects in which knowledge is superior to mere true belief, the only general reason we have for thinking that it must be possible to do so is that this is required to answer the activity question. It seems to me, therefore, that the activity question should not only be distinguished from the comparative state question, but that it should also be regarded as the more basic question, in the following sense: insofar as we think that the comparative state question is one that we need to answer, this will be because we think that we need to answer that question in order to provide a satisfying response to the activity question.

The points I have just made are absolutely fundamental to the argument of this thesis, so I want to reiterate them in a slightly different way, elaborating in places.

Another way into the issues here is to ask what is presupposed by the activity question and the comparative state question. The former question asks why we desire and seek knowledge, and not merely true beliefs, given that true belief is just as true, and apparently just as true, as knowledge. This question clearly presupposes that when we are interested in having the truth on some matter we are at least sometimes interested in having the truth, more specifically, in the manner of knowing the truth. This presupposition is difficult to deny. We frequently profess that we want to know things, from the trivial (e.g. ‘I want to know when the next train departs’) to the momentous (e.g. ‘I want to know who my real father is’). We also speak of wanting to work out/find out/discover/determine etc., whether..., or who..., or when..., or what..., or why..., or where..., or how..., and working out/finding out/discovering/determining the answer to a question of any of these forms is plausibly a matter of acquiring knowledge of the answer to that question.⁹ Given that we often express our desires and aims in more or less explicitly epistemic terms, it seems undeniable that we do in fact desire and seek knowledge of the answers to many questions. By contrast, the comparative state question presupposes something quite different. That question asks why a subject who knows that *p* is in a superior position, *ceteris paribus*, to a subject who merely truly believes that *p*. Obviously, this question presupposes that a subject who knows that *p* is in a superior position, *ceteris paribus*, to a subject who merely truly believes that *p*. Is this presupposition correct? It is not clear that we have any pretheoretical commitment to the truth of

⁹ I owe my formulation of this point to Soteriou (2013: 351). See that passage for further considerations that support the idea that, when a subject’s aim can be expressed by saying that he wants to work out/find out/discover/determine etc. the answer to some question, his aim is the acquisition of knowledge of the answer to that question.

this claim. Neither is there any obvious respect in which knowledge is superior to mere true belief, as the Platonic Socrates makes clear. Why, then, have philosophers tended to accept this presupposition without argument? In general, one way of explaining why an agent has *X* as his aim, rather than *Y*, is to identify respects in which *X* is superior to *Y*, given the agent's broader objectives. I suspect that philosophers have assumed that the explanation of why we are interested in knowledge, and not merely in true belief, must conform to this pattern. Once this assumption is in place, the question of why we seek knowledge and not just true belief reduces to the question of why knowledge is superior to mere true belief. Hence, the tendency to accept the presupposition of the comparative state question, and also the failure to distinguish the comparative state question from the activity question. In principle, however, there are ways of explaining why an agent seeks *X* rather than *Y* that do not appeal to the idea that *X* is superior to *Y*. One way to do this, as I explained earlier, would be to demonstrate that it is impossible for an agent to merely seek *Y* without seeking *X*. This raises the possibility that we might solve the problem of why enquirers are interested in knowledge, and not just in true belief, without solving the 'problem' of why knowledge is superior to mere true belief. And this possibility invites a further, more radical, proposal. If, as I suggested, we have no pretheoretical commitment to the truth of the claim that knowledge is superior to mere true belief, and the only theoretical purpose for which this claim is required is to explain why knowledge rather than true belief is the aim of enquiry, then the fact that we can explain why knowledge is the aim of enquiry without appealing to its relative superiority leaves us free to reject the claim that it is superior without violating any of our intuitive beliefs. In fact, I do not think that we should reject the claim that knowledge is superior to mere true belief, because philosophers have

succeeded in identifying respects, albeit rather limited ones, in which knowing is better than merely truly believing. But I shall be arguing that the rather limited respects in which knowledge is superior to mere true belief do not provide a satisfactory explanation of why knowledge, rather than true belief, is the aim of enquiry. Thus, the traditional approach to the value of knowledge fails to identify the basic problem philosophers should be concerned with in this area, and has indeed distracted us from the real issues.

To avoid misunderstanding, I must emphasize that I have not yet attempted to vindicate the claim that attention to the position of the enquirer is required to resolve philosophical puzzles about the value of knowledge. I have argued only that, beneath the traditional approach to the value of knowledge, embodied in the comparative state question, there lies a more basic problem about why, at least sometimes, knowledge, and not just true belief, is the aim of enquiry. The traditional approach, according to which this issue is to be resolved in a way that does not require attention to the position of the enquirer, remains on the table as one possible response to this problem; it therefore remains to be shown that this response is inadequate, and that the correct response is one that brings issues about the enquirer to the fore. To make further progress here, we need a deeper understanding of the thinking that lies behind the activity question. Why should the (alleged) facts that true belief is just as true, and just as useful, as knowledge, make our interest in knowledge at all puzzling? This question will be the focus of Chapter 1.

3. Chapter Summaries

It may be helpful, at this point, to provide the reader with a brief overview of the argument of the rest of this thesis.

Chapter 1 considers the activity question in more detail. The question challenges the rationality of seeking knowledge on the basis of considerations that purport to show that we ought to be interested only in arriving at true beliefs. In particular, true belief is claimed to be just as true, and just as useful, as knowledge. But why might these alleged facts about true belief be taken to impugn the rationality of seeking knowledge? The answer to this question is not obvious, and in Chapter 1 I identify two different ways in which they might be held to do so. The first, which can be traced to Williams' discussion of the method of doubt in his book *Descartes*, is as follows. The enquirer is someone who is seeking the truth about the answer to a question. Trivially, true belief is just as true as knowledge. Therefore, the enquirer will have got what he wants if he merely arrives at a true belief about the answer to his question; why, then, should he ever want any more than that? Seductive as it might seem, I shall argue that this reasoning is badly flawed, and therefore does not pose a serious challenge to the rationality of pursuing knowledge. There is, however, another way in which the possibility of merely pursuing true beliefs might be taken to undermine the rationality of seeking knowledge. Knowledge is harder to come by than true belief. If true belief is just as good as knowledge, why should enquirers bother themselves with the more demanding objective? I present this challenge in the form of a paradox, which I refer to as 'the problematic reasoning'. I spend the bulk of Chapter 1 arguing that various responses to the problematic reasoning are in fact inadequate, and that, by process of elimination, the right response is to argue that a certain principle of rationality on which it depends does not apply in the case of true belief and knowledge. To see why the principle does not apply in this case, however, it is necessary to first acquire a better understanding of the nature of belief itself.

I take up this task in Chapter 2. My investigation of the nature of belief occurs in the context of concerns about the aim of enquiry. This point is significant. An enquirer is a cognitively sophisticated subject. He is capable of recognizing that he does not have the truth about the answer to a certain question, and of adopting purposive means in order to obtain it. It follows from this, I shall argue, that the enquirer is capable of engaging in practical reasoning and theoretical deliberation (capacities which in turn imply that he is able to make conscious judgements and decisions), and also that the enquirer is capable of engaging in forms of acceptance besides belief. Once it is clear that we are dealing with a sophisticated kind of believer, certain issues about the nature of belief come to the fore. In particular: how are the beliefs of an enquirer related to his conscious thinking, and how are we to distinguish belief from other forms of acceptance that the enquirer is capable of engaging in? In relation to the first issue, I argue that having a disposition to consciously judge that p is necessary for an enquirer to have the outright (as opposed to the repressed) belief that p . This is relevant to the activity question, because the kind of belief that is relevant to that question is precisely outright belief. In considering the conditions under which someone who is enquiring into the question of whether p will acquire an outright belief about whether p , then, one thing it is relevant to consider is the conditions under which they will acquire a state that involves a disposition to consciously judge whether p . Relevant to this issue, in turn, is the question of what it is to judge that p . The nature of conscious judgement is far from obvious, but an apparently secure starting point is the thought that judging that p contrasts with other ways of consciously affirming the proposition that p , such as supposing that p for the sake of argument. How, then, should we distinguish judging that p in one's reasoning from supposing that p for the sake of argument? One simple proposal is

that judgement, unlike supposition, involves the acquisition/manifestation of the corresponding belief, but I argue that this proposal cannot explain why it is hard to make sense of the idea that a subject might be in error about the nature of his own conscious thinking. I consider and reject Shah and Velleman's (2005) claim that different ways of consciously affirming a proposition are to be distinguished by appealing to the intention with which the subject affirms a proposition. I go on to argue that, in order to explain why it is hard to make sense of the idea that a subject might be in error about the nature of his own conscious thinking, we should appeal to Soteriou's (2013) proposal that when the subject supposes that p for the sake of argument, the constraint of treating p as true in his reasoning is self-imposed. This is not the case when the subject judges that p in his reasoning. This is also relevant to distinguishing the state of believing that p from the state a subject is in when he is merely assuming that p for practical purposes. Although a subject who is merely assuming that p for practical purposes is disposed to act as though p is the case, and to plan on the assumption that p , when he plans on the assumption that p the constraint of treating p as true in his reasoning is self-imposed. This is not the case when the subject believes that p . A subject who believes that p is disposed to plan on the assumption that p by *judging* that p in his reasoning, where this involves treating the constraint of treating p as true as a non-self-imposed constraint.

Chapter 3 considers the nature of guessing. Guessing has not received much attention in the philosophical literature. However, Roy Sorensen (1984) proposes an analysis, and more recently David Owens (2003) has criticized the claim that belief aims at the truth by appealing to the idea that guessing also aims at the truth. One objective of Chapter 3 is to uncover inadequacies in these ideas. Thus, I argue that satisfying the conditions identified by Sorensen is insufficient for a subject to guess

that p , and also that satisfying his two main conditions—the ‘nondomination condition’ and the condition that believing is incompatible with guessing—is unnecessary for guessing. Furthermore, it is possible, contrary to Owens, for a subject to guess that p without the purpose of guessing truly. Another objective of Chapter 3 is to suggest how we might improve upon these ideas. One proposal I make here that we should conceive of the relationship between guessing and knowing at least partly in higher-order terms: when a subject guesses that p , he believes himself not to be answering that p knowledgeably, or at least suspends judgement about whether he is answering that p knowledgeably. These higher-order attitudes are absent when the subject judges that p . Thus, one difference between guessing that p and judging that p consists in the subject’s take on the epistemic status of his affirmation that p : when he judges that p , there is at least a purely negative respect in which he takes himself to know that p : he neither believes himself not to be affirming that p knowledgeably, nor suspends judgement about whether he is affirming that p knowledgeably.

Chapter 4 returns to the problematic reasoning set out in Chapter 1, with the aim of showing how the claims about belief and judgement made in Chapters 2 and 3 can be used to explain why that form of argument cannot be used to show that an enquirer ought at most to adopt the aim of arriving at true beliefs. Two claims are particularly important here. The first is that having the outright belief that p involves being disposed to judge that p . The second is that when a subject judges that p , he neither believes himself not to be affirming that p knowledgeably, nor suspends judgement about whether he is affirming that p knowledgeably. I argue that a subject who is capable of engaging in enquiry is a subject who has the capacity to adopt higher-order attitudes about his epistemic standing—e.g. to hold beliefs about what he does

and doesn't know—and that in a subject of this kind, the absence of such attitudes when the subject affirms that p suffices for the subject to take himself to know that p . In other words, the absence of certain *negative* higher-order epistemic attitudes is sufficient for an enquirer to manifest a *positive* higher-order attitude concerning his epistemic standing with respect to a proposition he has affirmed. Since outright belief involves a disposition to judge that p , it follows, I argue, that an enquirer will not have arrived at so much as an outright belief about the answer to his question until he takes himself to have found out the answer. This, I suggest, is the real reason why the enquirer cannot make life easier for himself merely by going after true beliefs. Even if he were only to adopt the aim of arriving at true beliefs, he would not have achieved his aim until he took himself to have achieved the more ambitious aim of knowledge. The correct response to the activity question, then, does not lie in the idea that knowledge is superior to mere true belief. The correct response, rather, is that, contrary to Williams, it is an error to think that an enquirer's interest in truth is basically only an interest in arriving at true beliefs, and, contrary to the problematic reasoning, it is an error to think that the enquirer can avail himself of more efficient methods by adopting only the aim of true belief.

4. Summary

In this chapter I have introduced the central idea of this thesis, which is that we need to pay attention to the position of an enquirer to resolve philosophical puzzles about the value of knowledge. I approached this idea by outlining a somewhat similar proposal that has been made by Bernard Williams. According to Williams, philosophical theorizing about the nature of knowledge has been led into error by a tendency to focus on the examiner situation, and the way to remedy this is to draw attention instead to the position of an enquirer. Although there appeared to be serious

objections to Williams' proposal, these objections do not undermine the suggestion that attention to the position of an enquirer is required to resolve philosophical puzzles about the value of knowledge, because issues about the value of knowledge are clearly distinct from (though related to) issues about the nature of knowledge. But how might considerations about enquiry be relevant to questions about the value of knowledge? I argued that it is important to distinguish the comparative state question that usually receives attention in philosophical discussions from what I called the activity question. The comparative state question asks why knowing that p is better than merely truly believing that p , whereas the activity question asks why enquirers desire and seek knowledge instead of merely caring about true beliefs. That these two questions are distinct is shown by the fact that an answer to the latter question needn't amount to an answer to the former. Furthermore, I argued that the activity question is the more basic question, in the sense that the only good general reason to accept the presupposition of the comparative state question, that knowledge *is* superior to mere true belief, is that this presupposition provides the only plausible response to the activity question. The basic problem in this area, then, is one that concerns enquirers, and even if tackling the comparative state question is one way of tackling the activity question, it is not clear that it is the only way of tackling that question. The next chapter undertakes a more detailed analysis of the activity question, with a view to clarifying what further options we have in responding to it.

Chapter 1

The Activity Question

In the Introduction I argued that philosophical reflection about the value of knowledge should begin with a question about enquiry: why do enquirers want to know the answers to questions, and not merely to arrive at some ‘lesser’ type of cognitive state with respect to the answers, e.g. that of having true beliefs about the answers? I contrasted this question, which I called ‘the activity question’, with the question of why it is better to know that p than to be in some ‘lesser’ cognitive state with respect to the proposition that p , e.g. that of having a mere true belief that p . I called this the ‘comparative state question’. The traditional approach to issues about the value of knowledge focuses on the comparative state question, but this is a mistake: the activity question identifies the more fundamental issue, and tackling the comparative state question is only one way of tackling the activity question. How else, then, might we respond to the activity question? To answer this question it is necessary first to know exactly why the assumption that true belief is just as true and just as useful as knowledge might be taken as a reason to think that enquirers should at most be in the business of seeking true beliefs. There is more than one line of thought that might be at work here. The main aim of this chapter is to present two different strands of thinking that might be operative at this point, and to indicate how we should respond to them.

Here is the plan. § 1 identifies and rejects an argument, found in Williams’ discussion of the method of doubt in his book *Descartes*, that might lead someone to press the activity question. § 2 underlines the verdict that this argument is problematic by showing that there are apparently insuperable obstacles in the way of

an adequate solution to the puzzle it generates. § 3 identifies an alternative motivation for pressing the activity question. I express this motivation in the form of a paradoxical argument that purports to establish the conclusion that it is rationally sub-optimal for an enquirer to seek knowledge of whether p , given that he could merely seek a true belief about whether p . § 4 argues that three responses to the problematic reasoning are inadequate; in particular, I argue that it is hard to see how we can come up with an answer to the comparative state question capable of defusing the problematic reasoning. § 5 identifies a more serious difficulty with the problematic reasoning. In general terms, the difficulty is that the problematic reasoning invokes a principle of rationality whose application depends on the satisfaction of conditions that may not be met in the case of true belief and knowledge. I conclude that the right response to the reasoning is indeed to press this kind of challenge.

1. Williams' Challenge

The version of the activity question I wish to focus on is the following: why do enquirers sometimes want to know the answers to questions, given that true belief is just as true, and apparently just as useful, as knowledge? So baldly put, the question invites the retort, 'Well, why shouldn't they?' This response, though flippant, brings out the point that, if someone is to be doing something philosophically interesting in pressing the activity question, he must have in mind some consideration, or set of considerations, which suggest that an enquirer should at most be interested in arriving at true beliefs. To answer the activity question in a satisfactory way, therefore, it is necessary for us first to know what these considerations are. I will not try to anticipate every possible consideration that might motivate someone here. I

shall merely identify, and respond to, what seem to me to be the two most obvious kinds of consideration that might be at work.

The first kind of consideration is raised by Bernard Williams in his discussion of the method of doubt in his book *Descartes*. There, he considers the question of ‘what reason Descartes has for regarding this unobvious strategy as straightforwardly the rational course’ (2005: 22). After all, Williams reasons, ‘we constantly want the truth about various matters, but hardly ever demand the indubitable’ (ibid.). In pursuing this question, Williams suggests that we should begin by considering the question of whether a subject who wants the truth should want so much as knowledge; if this can be established, ‘then anything peculiar in Descartes’s strategy will lie in a second step, from the search for knowledge to the search for certainty’ (ibid.: 23). Williams proposes that we should approach this question in the following way:

Let us take a person, call him ‘A’, who is in the most primitive situation of wanting the truth. He has no elaborate or reflective demands—it is not, for instance, that he wants to acquire or found a science (as Descartes does, or at least will want to do). He merely wants the truth on certain questions. Such questions can of course take many forms, ‘when...?’, ‘who...?’, etc.; we shall simplify, and take A as in each case wanting the true answer to a question of the form *whether p*. What exactly is it that A wants? What state does he want to arrive at? He wants, at the very least, to have a belief on the question whether *p*, and that belief to be true. That is to say, he wants at least to be in this state:

(i) if *p*, A believes that *p*, and if not *p*, A believes that not-*p*.

He wants *at least* to be in that state; why should he want any more? (2005: 23)

In this passage Williams apparently argues as follows. The enquirer is someone who wants the truth about the answer to some question. Trivially, true belief is just as true as knowledge. Insofar as all the enquirer wants is to have the truth, then, he will have got what he wants even if he arrives at no more than a true belief about the answer to his question. Why, then, should an enquirer ever want to *know* the answer to a question?

The correct response to this question is to challenge the assumption that an enquirer who wants the truth about whether p will have got what he wants if he arrives at a mere true belief about whether p . To see what is objectionable about this assumption, we need to begin by considering the question of what the enquirer wants in wanting the truth about whether p . An initial answer to this question is that what he wants is to *have* or *possess* the truth. However, talk of ‘having’ or ‘being in possession of’ the truth about whether p is metaphorical. How are we to specify what the enquirer wants in literal terms? The most natural way of understanding such talk is, I think, epistemic. Relevant here is Williams’ remark that ‘[o]rdinary speech...effortlessly expresses the thought that A wants the truth on the question ‘is p true or not?’ in the form of saying that A wants to know whether p ’ (2005: 23). But this interpretation clearly won’t do for the purposes of the present argument for the claim that the enquirer will have got what he wants if he arrives at a mere true belief about whether p . The argument now reads as follows: ‘The enquirer is someone who wants to know the answer to a question. Trivially, true belief is just as true as knowledge. Insofar as all the enquirer wants is to know the answer to a question, then, he will have got what he wants even if he arrives at no more than a true belief about the answer to his

question'. On this interpretation, far from establishing the conclusion that the enquirer will have got what he wants if he arrives at a mere true belief about whether p , the premise actually *contradicts* that conclusion, provided we assume that true belief is insufficient for knowledge. An alternative specification of what the enquirer wants in wanting the truth about whether p is clearly required.

It is not difficult to find one. The point that true belief is just as true as knowledge invites us to understand talk of 'wanting the truth' in a more minimal way: the enquirer is someone who merely wants to arrive at a true belief about the answer to his question. When this interpretation of 'wanting the truth' is plugged into Williams' argument, it reads as follows: 'The enquirer wants a true belief about the answer to some question. Trivially, true belief is just as true as knowledge. Therefore, the enquirer will have got what he wants even if he arrives at no more than a true belief about the answer to his question.' It seems impossible to disagree with this argument, although, heroically, Williams tries to refute it. The problem now, however, is that the argument fails to generate a puzzle about why an enquirer should sometimes want to know the answer to a question. If all the enquirer wants is to arrive at a true belief about whether p , why on earth should he want to know whether p ? Haven't we just built it into the case that he would be satisfied by a mere true belief?

There is no need to explain why someone who wants the truth about whether p wants to know whether p . The correct response to this challenge is ask what the enquirer wants in wanting the truth about whether p . If what the enquirer wants is to know whether p , then it is false that the enquirer will have got what he wants if he merely arrives at some state that falls short of knowledge, hence no puzzle arises about why he wants to know. If, on the other hand, what the enquirer wants is merely to arrive

at a true belief about whether p , then there is no reason to think that he does want to know, hence no puzzle can arise about why he does. Williams' challenge should be rejected.

Having raised the question of why an enquirer who would be satisfied, on the face of it, with a mere true belief about whether p would not in fact be satisfied with less than knowledge of whether p , Williams goes on to provide an answer. Although I take this to be a bad question, for the reasons I have indicated, I nevertheless wish to indicate why I think that Williams' answer is inadequate, for two main reasons. The first is that considering Williams' answer raises a number of points of philosophical interest. The second is that it is hard to see what other kind of strategy, aside from the one Williams pursues, someone might employ in trying to answer this question. The fact that there are apparently decisive objections to Williams' answer therefore suggests that no explanation of why an enquirer wants to know, given that he would be satisfied with mere true belief, can be given. This underlines the point that the argument that leads Williams to attempt to provide an explanation of this must be flawed.

2. Williams' Solution

As we saw, Williams begins by introducing 'A' as the name for an enquirer 'who is in the most primitive situation of wanting the truth' (2005: 23) on certain questions. For simplicity, Williams assumes that A's question is of the form 'Is p true or not?', and he goes on to suggest that what A wants, at the very least, is to have a true belief about whether p —that is, to be in the following state of affairs:

- (i) if p , A believes that p , and if not p , A believes that not- p .

The puzzle is therefore to explain why someone who is initially specified only as wanting a true belief about whether p will in fact want to know whether p .

In Williams' words, the 'basic point here...concerns the methods available to A to get even into state (i)' (ibid.: 24). As an enquirer, A wants, and not merely wishes, to be in state (i); given that he has no reason to believe that he will get into state (i) just by waiting and hoping, 'he must adopt purposive means to get into (i)' (ibid.). What, in general outline, do such means involve? Here Williams contrasts the situation of an enquirer seeking a true belief about whether p with the situation of a flint collector who wants to acquire a collection of prehistoric flints, and only prehistoric flints. A procedure the flint collector might employ is to collect flints indiscriminately and then determine which, if any, of the flints he has collected are prehistoric. But 'the analogous process with acquiring true beliefs would be not just inefficient, but incomprehensible', because 'to acquire a belief is already to assume an answer to the question of whether it is true' (ibid.: 25). Thus, 'a method which A uses as an enquirer to get into state (i) must be a method of acquiring beliefs which itself makes it likely that the beliefs A acquires by it will be true ones; or, equivalently, is such that he is unlikely to acquire beliefs by that method unless they are true' (ibid.). As shorthand for this, we can say that A must try to use a *reliable* method of acquiring beliefs. The requirement that A 's method of acquiring a true belief about whether p should be reliable is itself equivalent to the following requirement on A 's belief about whether p : that it should be 'produced in such a way that one is unlikely to acquire beliefs in that way unless they are true' (ibid.). Williams dubs this property of beliefs, conceived of as states of particular individuals, ' E '. If A succeeds in acquiring a true belief about whether p by means of the kind he must try to use, A will be in the following state:

(iv) if p , A believes that p , and if not- p , A believes that not- p , and in either case A 's belief has the property E .

Furthermore, Williams claims, 'we can put his [A 's] wants together' (ibid.: 26)—that is, his desire to be in (i) and his desire to use a reliable method to acquire a true belief about whether p —to arrive at the conclusion that (iv), and not merely (i), is the state of affairs the enquirer wants to end up in. Suppose that p is in fact the true answer to A 's question. It follows that if A is successful he will be in the following state of affairs:

(v) A truly believes that p , and his belief has the property E .

Is being in (v) sufficient for A to know that p ? 'If so', Williams says, 'we will have shown how it is that in wanting the truth, A wants to know' (ibid.: 27). Williams stresses that the notion of knowledge in play, unlike Descartes', is 'an absolutely minimal everyday conception of knowledge' (ibid.), and plausibly does not require A to know that he knows that p , to feel specially certain that p is true, or even to have conscious reasons for believing that p . Nevertheless, Williams thinks that the insufficiency of (v) for knowing that p , even in an absolutely minimal everyday sense of the term, is demonstrated by Gettier cases.¹⁰ Thus, the present argument has not yet succeeded in explaining why, in wanting the truth about whether p , A wants to know whether p .

The issue we are now confronted with resolves itself into two questions: (a) what, in addition to being in (v), is required for A to know that p , and (b), assuming that we can specify the additional condition required for knowledge, is it the case that an enquirer who wants the truth about whether p will want to have a true belief about

¹⁰ See Gettier (1963).

whether p , possessing the property E , that satisfies this condition? Williams does not attempt to provide a detailed answer to (a). He says only that, in broad terms, the additional factor required for knowledge absent in Gettier cases is that ‘the truth of the belief be non-accidental relative to the method or way by which it [was] produced’ (ibid.: 30). A more detailed specification of the requirement would of course be desirable, but for the purposes of his argument Williams takes it to be unnecessary: however exactly it is spelt out, A will want his true belief about whether p to satisfy this further condition:

While what he wants are true beliefs, as a conscious enquirer in a non-magical world [A] has to commit himself to a policy of acquiring them in reliable ways. Accidentally true beliefs, though they might seem welcome merely as true beliefs, are in fact only a sub-class of beliefs to which his methods are irrelevant; relative to his strategies of enquiry, they might as well have been false, and this state of affairs he cannot want. (ibid.)

If this is right, we can now say that A wants to arrive at state (v), but where ‘ E ’ is interpreted as the property of a belief’s having been ‘*appropriately* produced in a way such that beliefs produced in that way are generally true’ (ibid.: 31), i.e. produced in such a way that the truth of the belief is non-accidental relative to the use of the method. Williams suggests that, so understood, ‘(v) surely is sufficient for knowledge’. Thus, ‘starting merely from the idea of pursuing truth in a non-magical world, and so of the truth-seeker’s using methods of enquiry, we do arrive at the conclusion that the search for truth is the search for knowledge’ (ibid.).

Although Williams’ argument raises a number of issues, its central contention is clear enough: a rational enquirer who wants a true belief about whether p must want

to have a true belief about whether p that satisfies certain further conditions, such that a true belief about whether p that satisfies those conditions necessarily constitutes knowledge of whether p . Before I consider this contention, however, I want to make a more basic point about the argument which is connected with the concern I raised, in the previous section, about whether the problem Williams sets himself is well motivated.

Following his initial set up—i.e. the introduction of A as an enquirer who wants the truth on a question of the form ‘Is p true or not?’—Williams’ first move is to claim that, in wanting the truth about whether p , A wants at the very least to arrive at a true belief about whether p . I have already objected to this idea on the grounds that it is not clear why we should specify what A wants in a way that falls short of him wanting to know whether p . It is worth noting, however, that one can also object to this idea from the opposite direction: why should we specify what A wants as *so much* as a true belief about whether p ? It is not difficult to imagine circumstances in which a subject who ‘wants the truth about whether p ’ would be satisfied with less than a true belief. For example, a contestant in a quiz show might want the truth about whether p in the more specific sense of wanting to give the true answer to the question of whether p . But in order to give the true answer she doesn’t need to have a true belief about the answer, for she can give the true answer by correctly guessing whether p . Correctly guessing whether p does not require the acquisition of knowledge of, or a true belief about, whether p . (I will be discussing this kind of guessing in more detail in Chapter 3.) If this is right, then it is false that a subject who is described only as wanting the truth about whether p must be specified as wanting at least a true belief about whether p .

The effect of this objection is to reinforce the dilemma, raised in the previous section, about how we should specify *A*'s desire to have the truth about whether *p*. If we wish to say, with Williams, that what *A* wants is at least to have a true belief about whether *p*, we now need to explain not only why we should not specify *A*'s want as a desire to know whether *p*, but also why we should not specify *A*'s want as a desire to correctly guess whether *p*. And if we should not specify *A*'s desire merely as a desire to correctly guess whether *p*, why should we stop short of specifying it as a desire to know whether *p*?

Even if we are convinced of the arbitrary nature of the task Williams sets himself, we can still ask whether he manages to complete that task successfully. Does he succeed in showing that a rational enquirer who wants a true belief about whether *p* must want to have a true belief about whether *p* that satisfies certain further conditions, such that a true belief about whether *p* that satisfies those conditions necessarily constitutes knowledge of whether *p*? Williams' argument for this claim relies on the assumption that, since *A* is an enquirer and thus genuinely wants (and not merely wishes) to arrive at a true belief about whether *p*, *A* 'must' adopt purposive means to acquire a true belief about whether *p*. However, there does not appear to be any significant sense in which a subject who wants a true belief about whether *p*, and who has no reason to believe that he will arrive at one just by waiting and hoping, 'must' adopt purposive means in order to arrive at one. This is so even if we suppose, as seems plausible, that desiring that *p* disposes the subject to behave in ways that he believes are conducive to making *p* true. We often do adopt purposive means to obtain the things we desire, but this is not a matter of binding necessity. For example, consistently with my wanting it to be the case that *p* I may recognize that were I to try to make it the case that *p* then other aims I have would be

jeopardised. This might be because I recognize that devoting time and resources to making it the case that p would compromise the time and resources I can devote to other objectives (think of the choice a sportswoman might make to focus on just one of the sports she excels at), or because I recognize that, were it to be the case that p , other objectives of mine would be undermined or put at risk (if I eat the ice cream then it is less likely that I will succeed in losing weight). The recognition that adopting purposive means to realize one of my desires would jeopardise other aims I have in either of these ways might lead me to suppress that desire, so that I don't act on it, but suppressing a desire in this sense does not make it go away. Thus, it seems perfectly consistent with my having a certain desire that, although I am disposed to act on it, I suppress this disposition and do not adopt purposive means to bring about the desired state of affairs. This needn't involve any irrationality on my part: on the contrary, the kind of self-control I exhibit when I suppress one of my desires in order to focus my efforts on other ends that I value more highly is often an expression of rationality.

We can avoid this objection by supposing that A not only wants to have a true belief about whether p , but has decided to arrive at a true belief about whether p —or at least to try to do so—and has set about adopting means of arriving at one. Must A now want to arrive at a true belief about whether p in such a way that it constitutes knowledge of whether p ? It is here, I think, that we encounter the fatal flaw in Williams' argumentative strategy. Suppose that A must indeed try to use a reliable method, and, moreover, one that is sufficiently reliable for a true belief about whether p produced by that method to constitute knowledge of whether p , provided that the anti-Gettier condition is also satisfied. Does it follow that we can 'put his two wants together' and say that his objective is to arrive at true belief about

whether p that was produced by a reliable method—i.e. that his objective is to be in state of affairs (iv)? The problem with this inference is that A is interested in using a reliable method only as a means of arriving at a true belief about whether p . If I desire something merely as a means to some further end, it does not follow that what I desire is to arrive at the end *by* those means. Intuitively, if I fail to use the means but still get the end I will be perfectly satisfied with the outcome, for I was interested in the means *only* as a means to the end.

An example reinforces the general point that arriving at an end by reliable means confers no additional value on the end:

Imagine two great cups of coffee identical in every relevant respect—they look the same, taste the same, smell the same, are of the same quality, and so on. Clearly, we value great cups of coffee. Moreover, given that we value great cups of coffee, it follows that we also value reliable coffee-making machines—i.e. machines which regularly produce good coffee. Notice, however, that once we've got the great coffee, then we don't care whether it was produced by a reliable coffee-making machine... In order to see this, note that if one were told that only one of the great identical cups of coffee before one had been produced by a reliable coffee-making machine, this would have no bearing at all on the issue of which cup one preferred; one would still be indifferent on this score. (Pritchard 2010: 9)¹¹

Precisely the same point appears to apply to A . Even if he must aim to use a reliable method to arrive at a true belief about whether p , it does not follow that his aim is to arrive at a true belief about whether p *by* a reliable method. If this is correct, then no

¹¹ Pritchard's discussion at this point follows Zagzebski (2003), who also uses a coffee-based example—specifically, an espresso.

argument of the kind Williams advocates can possibly succeed. Considerations about the kind of means a rational agent would have to use to achieve a certain end cannot be used to justify an augmented conception of his end.

In this section I have made three main critical points about Williams' argument. The first was that it is not clear why we should specify *A*'s desire to have the truth about whether *p* as a desire for so *much* as a true belief about whether *p*. This point is connected with the concerns I raised in the previous section about whether the problem Williams addresses is well motivated. The second point was that it is not the case that a subject who merely desires a certain objective 'must' adopt purpose means to obtain that objective. I suggested, on Williams' behalf, that we might avoid this difficulty by specifying that *A* does not merely want a true belief about whether *p*, but has *decided* to (try to) arrive at a true belief about whether *p*. The third point was that there is, in any case, a fundamental flaw in the argumentative strategy Williams employs. It cannot be inferred from the fact that an agent must pursue an end by means of a certain type that his objective is to achieve the end *by* means of that type. The assumption that this kind of inference is legitimate is integral to Williams' argument. Since it is hard to see how we might try to solve the problem Williams sets for us other than by appealing to an inference of this kind, it appears that this problem is insoluble. It is therefore fortunate that the line of thought that leads to this problem is badly flawed.

3. The Problematic Reasoning

The point that there are ways of having or possessing the truth besides knowing cannot be used to show that a truth-seeker will have got what he wants even if he arrives at a state that falls short of knowledge. There is therefore no need to explain

why an enquirer should want to know given that, on the face of it, he will have got what he wants even if arrives at only (e.g.) a true belief. However, the point that there are non-epistemic ways of specifying what the enquirer wants in wanting the truth raises the possibility of another kind of challenge to the rationality of seeking knowledge. Why should an enquirer ever want to have the truth in the manner of knowing the truth? Why doesn't he restrict himself to aiming at the truth in more modest ways?

Once again, put so baldly, the question invites the retort, 'Well, why shouldn't he?' The word 'modest' hints at an answer to this question. Knowledge, one might think, is difficult to attain; it is easier to come by true beliefs. However, it is difficult to identify respects in which knowledge is superior to mere true belief. Why, then, should an enquirer bother himself with the acquisition of knowledge? Why not make life easier by sticking to the simpler task of acquiring true beliefs? This argument is crude as things stand, but it should already be clear that it does not depend on the assumption that all a truth-seeker basically wants is to arrive at true beliefs. An analogy helps to bring this out. Imagine an archer confronted with two targets, one significantly further away, and so harder to hit, than the other. Asking Williams' question is a bit like asking, 'Why should the archer want to hit the target that's further away, given that he will have got what he wants if he only hits the closer target?' This question is odd: if all he wants to do is hit the closer target, why *should* the archer want to hit the one that's further away? Asking the current question, by contrast, is more like asking, 'Why should the archer shoot at the target that's further way, instead of merely shooting at the one that's closer, given that hitting either of them is sufficient to win the competition?' This question does not depend on the assumption that all the archer basically wants is to hit the closer target. It relies only

on the assumption that, relative to the archer's broader objective of winning the competition, hitting the closer target is just as good as hitting the further target. The analogue of this assumption in the case of true belief and knowledge isn't obviously false. As the Platonic Socrates points out, it isn't easy to identify respects in which knowledge is any better than true belief relative to our broader objectives. True belief appears to be just as good as knowledge for any practical purpose.

Before I go any further, I want to set out the argument for the claim that enquirers should restrict themselves to aiming at true beliefs in more detail. The argument relies on two basic assumptions. The first is that true belief is just as good, given the enquirer's broader objectives, as knowledge. The second is that it is more difficult for the enquirer to arrive at knowledge than true beliefs. What is the justification for the premise that it is more difficult for an enquirer to arrive at knowledge than at true beliefs? It is natural here to appeal to the claim that truly believing that p is necessary but insufficient for knowing that p . It can then be argued that because knowledge is, in this sense, a more demanding condition than true belief, it is more difficult for the enquirer to bring about. There is then a further general assumption, to the effect that if objective X is just as good as objective Y , but objective Y is harder to achieve than objective X , then the agent shouldn't aim to achieve objective Y ; he should, at most, aim to achieve objective X . If these assumptions hold, it follows that it is rationally sub-optimal for an enquirer to adopt the aim of acquiring knowledge; at most, he ought to adopt the aim of arriving at true beliefs.

Since I will be discussing this argument extensively in what follows, it will be helpful to have it set out in a formal way. I will refer to it as 'the problematic reasoning', to indicate that its conclusion is one that we have at least *prima facie*

reason to reject. (Following Williams, I will assume, for simplicity's sake, that the enquirer wants the truth on a question of the form 'Is it the case that p ?'.)

The Problematic Reasoning

1. Having a true belief about whether p is just as good for the enquirer as knowing whether p , because true belief is just as true, and just as useful, as knowledge.
2. Having a true belief about whether p is necessary but insufficient for knowing whether p .
3. Therefore (from (2)), it is more difficult for an enquirer to arrive at knowledge of whether p than it is for an enquirer to arrive at a true belief about whether p .
4. If objective X is just as good as objective Y , but it is more difficult for the agent to achieve objective Y than objective X , then it is rationally sub-optimal for the agent to adopt Y as his aim; he should, at most, adopt X as his objective.
5. Given (1), (4) and (5), it is rationally sub-optimal for an enquirer to adopt the aim of arriving at knowledge of whether p ; he should, at most, adopt the aim of arriving at a true belief about whether p .

Notice that the conclusion does not merely say that it is *sometimes* rationally sub-optimal to adopt the aim of acquiring knowledge. It is not obvious that this claim contradicts anything we ought to accept. When knowledge is not in the offing it is sometimes sensible to hazard a guess. Rather, the conclusion is that it is *always* rationally sub-optimal for an enquirer to seek knowledge of the answer to a question.

This claim does seem at odds with what we naturally believe, for it implies that ordinary human beings violate standards of rationality whenever they search for knowledge. If we wish to avoid this conclusion, some way to resist the problematic reasoning must be found.

We have various options here. The reasoning assumes that having a true belief about whether p is just as good as knowing whether p , and thus that a subject who knows that p is not in a superior state to a subject who merely truly believes that p . Thus, those who take themselves to have a solution to the comparative state question will reject premise (1). But it is important to appreciate that this is only one way of responding to the reasoning. An alternative response would be to reject the premise that having a true belief about whether p is necessary but insufficient for knowing whether p . One might also object to the inference from (2) to (3), or to the principle of rationality expressed by premise (4).

In the next section I will argue that we cannot adequately respond to the problematic reasoning by rejecting premises (1) or (2), or by arguing that the inference from (2) to (3) is invalid. This is not because I think that these premises are entirely unobjectionable, or because I think that (3) is a straightforward logical consequence of (2); on the contrary, I shall be arguing, regarding (1), that there is a respect in which knowledge is superior to mere true belief, and that the inference from (2) to (3) is clearly invalid. However, though correct, these objections fail to go to the heart of the matter; it is possible to modify the reasoning so that it avoids these objections whilst still entailing that, in a worryingly wide range of cases, it is rationally sub-optimal for the enquirer to engage in the pursuit of knowledge.

4. Some Inadequate Responses to the Problematic Reasoning

4.1 Is True Belief Just as Good as Knowledge?

The first premise of the problematic reasoning asserts that having a true belief about whether p is just as good as knowing whether p . Some will object that this claim is false. According to them, it is possible identify respects in which knowledge is superior to mere true belief, and once this is appreciated the problematic reasoning ceases to pose a challenge to the rationality of our interest in knowledge. Obviously, I cannot pretend to have anticipated every possible answer to the comparative state question, or even to have studied all of the extant proposals about its solution. Thus, I certainly cannot *prove* that there is no solution to the comparative state problem that defuses the problematic reasoning. All I shall try to show is that one prominent response to the comparative state question—Williamson’s (2000) elaboration of Plato’s original solution—is inadequate when it is considered as a response to the problematic reasoning. Furthermore, it fails for a reason that casts doubt more generally on the idea that we can adequately respond to the problematic reasoning by denying its first premise.

Williamson’s proposal refines the old Platonic suggestion that knowledge is superior to mere true belief because it is more stable. Plato introduces the proposal using the metaphor of the statues of Daedalus, which run away unless they are tied down:

SOCRATES: There’s as little point in paying a lot of money for an unrestrained statue of his as there is for a runaway slave: it doesn’t stay put. But Daedalus’ pieces are so beautiful that they’re worth a great deal if they’re anchored. What am I getting at? I mean this to be an analogy for true beliefs. As long as they stay put, true beliefs too constitute a thing of beauty and do

nothing but good. The problem is that they tend not to stay for long; they escape from the human soul and this reduces their value, unless they're anchored by working out the reason. And this anchoring is recollection, Meno, my friend, as we agreed earlier. When true beliefs are anchored, they become pieces of knowledge and they become stable. That's why knowledge is more valuable than true belief, and the difference between the two is that knowledge has been anchored. (97e–98a)

The account of the difference between knowledge and mere true belief that Plato offers in this passage—that the subject who knows has ‘worked out the reason’—is obscure, but the basic idea that knowledge is more valuable than mere true belief because it is more durable or stable seems clear enough independently of further elaboration of that proposal. At this point, then, there are two questions we need to consider. Firstly, is it true, as Plato says, that knowledge is more stable—i.e. less likely to be lost or extinguished—than mere true belief? Secondly, if knowledge is more stable than mere true belief, does this identify a respect in which knowledge is superior to mere true belief that undermines the problematic reasoning? My answers to these questions will be ‘yes’ and ‘no’, respectively.

In thinking about ways in which knowledge might be more stable than mere true belief, it helps to begin by considering a passage from the discussion of Williams' that I criticized earlier, in which he reflects on the relation between his problem and the Meno problem:

This question [i.e. why should *A* want to know whether *p*, given that he is initially specified only as wanting a true belief about whether *p*?] is superficially like a very old one, raised in Plato's dialogue, the *Meno*:

wherein lies the superiority of knowledge over true belief? It does not lie, as Socrates quickly points out to a confused Meno (97c), in knowledge's always being true—*true* belief is just as true as knowledge. Rather, Socrates suggests, knowledge—which he connects with systematic understanding—will not run away: a point which we may take in the sense more interesting for the theory of knowledge, that knowledge cannot rationally be rendered doubtful, rather than as the blankly psychological proposition (in any case, surely, very dubious) that one is more disposed to forget what one merely believes than what one knows. (2005: 24)

The crucial point in this passage is that there is more than one way in which true belief can be lost. One can simply forget what one truly believes; alternatively, one can be led to abandon one's true belief in the face of evidence that suggests that it is false, or at least not adequately supported. As Williams says, the suggestion that one is more likely to forget what one merely truly believes than what one knows is doubtful at best. Suppose, for example, that we think of the difference between knowledge and true belief in something like the way that Plato appears to in the *Meno*, so that knowing that p requires (very roughly) understanding why it is true that p . Understanding why a proposition is true can be an aid to remembering it. One might think, for example, that someone who has followed Socrates' diagrammatic 'proof', in the *Meno*, that the square constructed on the diagonal of a square is double in area is more likely to remember that proposition than someone who has simply been told it is true, without being given any further understanding of why. However, understanding why a proposition is true is only one way of cementing that proposition in memory, and perhaps not even the most effective way. As Craig points out, 'effect of early upbringing, emotive ties or Humean psychological

mechanisms may be just as good, and better' (1990: 7). It is far from clear that, in general, knowledge is less likely to be forgotten than mere true belief. This suggests that if there is to be anything in Plato's claim at all, it must rather concern the differential tendencies of knowledge and true belief to be undermined in some other way—e.g. by being rationally rendered doubtful.

The proposal that Williams makes on Plato's behalf at this point is the strong one that 'knowledge cannot rationally be rendered doubtful'. On the face of it this claim is too strong: it does not seem difficult to imagine cases in which a subject who knows that p can be led, rationally, to doubt that p .¹² But even if this claim cannot be sustained, one might still maintain the weaker one that knowledge is *less likely* than mere true belief to be undermined by future evidence. This is indeed what Williamson suggests:

Present knowledge is less vulnerable than mere present true belief to *rational* undermining by future evidence, which is not to say that it is completely invulnerable to such undermining. If your cognitive faculties are in good order, the probability of your believing p tomorrow is greater conditional on your knowing p today than on your merely believing p truly today (that is, believing p truly without knowing p). Consequently, the probability of your believing p tomorrow is greater conditional on your knowing p today than on your believing p truly today. (2000: 79)

Williamson illustrates his proposal by adapting Plato's own example:

One can lose a mere true belief by discovering the falsity of further beliefs on which it had been essentially based; quite often, the truth will out. One

¹² For a plausible counterexample, see Kripke (2011: 35–36).

cannot lose knowledge that way, because a true belief essentially based on false beliefs does not constitute knowledge. For example, I might derive the true belief that this road goes to Larissa from the two false (but perhaps justified) beliefs that Larissa is due north and that this road goes due north; when dawn breaks in an unexpected quarter and I realize that this road goes south, without having been given any reason to doubt that Larissa is due north, I abandon the belief that this road goes to Larissa. Since that true belief was essentially based on false beliefs, it did not constitute knowledge. (2000: 78)

The point that mere true belief is more vulnerable to being rationally undermined by future evidence does not show that the probability of your believing p tomorrow is greater conditional on your knowing p today than on your merely believing p truly today if your belief that p is ‘profoundly dogmatic’—i.e. if you are not disposed to revise your belief that p in response to new evidence. This is why Williamson imposes the condition that the subject’s ‘cognitive faculties’ must be in ‘good order’. Furthermore, to establish the thesis that the probability of your believing p tomorrow is greater conditional on your knowing p today than on your merely believing p truly today we also need the assumption that there is at least some chance of the kind of evidence that could rationally undermine a mere true belief that p , but not knowledge that p , coming into your possession. This assumption does not seem problematic. Provided, then, that we are dealing with subjects who are by and large rational, it seems reasonable to accept the hypothesis that the probability of a subject believing that p tomorrow is greater conditional on him knowing that p today than it is on him

merely believing that p truly today. This is to endorse, in a somewhat qualified way, Plato's claim that knowledge is more stable than mere true belief.¹³

Haven't I just conceded that there is a respect in which knowledge differs from mere true belief? And isn't this difference of obvious value to us? After all, it is only for as long as true beliefs stick around that they can continue to guarantee correctness of action, and we have just seen that true beliefs that constitute knowledge are more stable than those that don't. So hasn't the paradoxical conclusion that we ought to be interested only in acquiring true beliefs been avoided? I don't think that it has. The problem is that, even conceding Williamson's point, our interest in knowledge appears to be far more widespread than it ought to be given only the rather minor respect in which knowledge is superior to mere true belief that has been identified. This point is not a novel one; it is eloquently expressed, for example, by Craig:

Whether the stabilisation of true beliefs is important or not depends on which beliefs we are considering, and the circumstances of the agent—many beliefs are required for the guidance of single, 'one-off' actions under circumstances which will not recur, and once the particular occasion is past there is no obvious value at all in their persistence. (I might now need a true belief about the time; but that this belief should persist, so that tomorrow I will still know

¹³ Jonathan Kvanvig has objected to Williamson's proposal on the grounds that 'there are other ways in which knowledge can be undermined but true belief remains unaffected. In particular, knowledge can be undermined at a later time by future changes of which one is unaware, where true belief is retained. For example, my mathematical knowledge might be undermined tomorrow by the sincere testimony of a renowned mathematician to the effect that what I believe is false. Until such testimony is rendered, I have such knowledge, but I lose it when the defeating testimony is given, even though I am unaware that such testimony has occurred. I thereby lose my knowledge but not my true belief' (2003: 15). But the point that there are ways of losing knowledge that are not ways of losing (true) belief is irrelevant, because Williamson's claim is only that the probability of my *believing* that p tomorrow is greater conditional on my knowing that p today than it is on my truly believing that p today.

what the time was today, at the moment when I wanted to know it, may be of no interest to me whatever.) (1990: 7)

In terms of the problematic reasoning, the issue is that we often seek knowledge of the answers to questions even in situations in which arriving at *stable* true beliefs about the answers is not important. But in such scenarios, for all that has been said, true belief is just as good as knowledge. Consequently, we can run a modified form of the problematic reasoning to show that, *at least in these scenarios*, it is rationally sub-optimal for enquirers to seek knowledge and not merely true beliefs.

Perhaps it will be objected here that enquirers can rarely be certain that the persistence of true belief is unimportant. Tomorrow someone might offer me £100 if I can remember what the time was yesterday, at the moment I wanted to know it. This being so, caution recommends seeking knowledge. The problem with this line of argument is that it shows at most that rational enquirers ought to balance the cost of seeking a more demanding objective, knowledge, against the possible benefit of arriving at more stable true beliefs. It is implausible that the balance of costs and benefits will always recommend the cautious strategy of seeking knowledge. Furthermore, it is striking that the need to balance the cost of seeking a more demanding objective, knowledge, against the benefit of arriving at more stable true beliefs is one that finds no trace in the thinking of ordinary enquirers: one never thinks, ‘Given that circumstances in which a true belief about whether p would be useful to me are highly unlikely to recur, perhaps I ought only to seek a true belief about whether p , and not knowledge of whether p ’.

Although this objection has arisen here specifically as an objection to Williamson’s proposal, it should be clear that it is merely a particular application of a much more

general argument. A solution to the comparative state question will have to identify some feature, F (e.g. greater stability), that knowledge possesses but mere true belief lacks, in virtue of which knowledge is more valuable than mere true belief. If there are cases in which enquirers seek knowledge without violating standards of rationality, but have no interest in F , then that solution to the comparative state question will fail to defuse the problematic reasoning when it is run for those cases. Hence, we will be forced to accept the conclusion that, in those cases at least, it is rationally sub-optimal for enquirers to seek knowledge. Given the generality of this kind of challenge, I think we should be sceptical of the idea that we can adequately respond to the problematic reasoning by denying its first premise.

4.2 Is True Belief Necessary but Insufficient for Knowledge?

According to premise (2) of the problematic reasoning, having a true belief about whether p is necessary but insufficient for knowing whether p . This follows from the claim that truly believing that p is necessary but insufficient for knowing that p , provided we accept the further assumptions that (a) knowing whether p is a matter of knowing that p , if p , and knowing that $not-p$, if $not-p$, and (b) that having a true belief about whether p is a matter of believing that p , if p , and believing that $not-p$, if $not-p$. The argument for the insufficiency claim is now as follows. Suppose that S has a true belief about whether p . It follows from (b) that S either truly believes that p or truly believes that $not-p$. Given that truly believing that p is insufficient for knowing that p , however, S may fail to know that p or to know that $not-p$. Since knowing that p or knowing that $not-p$ is necessary for knowing whether p , S may therefore fail to know whether p , despite having a true belief about whether p . The argument for the necessity claim proceeds in a similar fashion. Suppose that S knows whether p . It follows from (a) that S either knows that p or else knows that $not-p$. If

true belief is necessary for knowledge, then S either truly believes that p or else truly believes that *not-p*. Given (b), S must therefore have a true belief about whether p . So having a true belief about whether p is necessary for knowing whether p .

I will not here go into the question of whether we should accept assumptions (a) and (b). However, I do wish to say a little more about the assumption that truly believing that p is necessary but insufficient for knowing that p . Should we accept this claim?

In considering this question, a simple but important point is that, in saying that truly believing that p is necessary but insufficient for knowing that p , one doesn't commit oneself to the further claim that it is possible to specify non-circular necessary and sufficient conditions for knowing that p . As Williamson points out, 'a necessary but insufficient condition need not be a conjunct of a non-circular necessary and sufficient condition' (2000: 3). The example Williamson uses to demonstrate this point is that of being coloured and being red. Being coloured is a necessary but insufficient condition for being red, but we cannot specify a non-circular necessary and sufficient condition for something to be red by conjoining being coloured with other conditions that are stated without reference to red. (The most obvious suggestion—that x is red if and only if (a) x is coloured and (b) x is not green, or blue, or yellow, or pink, etc.—is clearly objectionable. Even supposing the account to be adequate in other respects (for example, that the open-endedness of (b) is unobjectionable, or can be repaired), it is circular. In specifying condition (b) we use other colour concepts—e.g. being green—so the question arises of whether we can specify necessary and sufficient conditions for the application of these other colour concepts without using the concept of being red. Using the same strategy we used in the case of red— x is green if and only if (a) x is coloured and (b) x is not *red*, or blue, or yellow, or pink, etc.—immediately results in circularity.) Thus, the much

advertised failure of philosophers to succeed in specifying a non-circular necessary and sufficient condition for knowing that p , by conjoining truly believing that p with further conditions, is not a reason to doubt the claim that truly believing that p is necessary but insufficient for knowing that p .

The premise that truly believing that p is necessary but insufficient for knowing that p is equivalent to the following conjunction: necessarily, if S knows that p then S believes that p and p , but it is not the case that, necessarily, if S believes that p and p , then S knows that p . Thus, one can deny the premise either by denying the former necessity claim, or by denying the latter insufficiency claim, or by denying both. The claim that truly believing that p is necessary for knowing that p has been denied by some philosophers on the basis that believing that p is unnecessary for knowing that p .¹⁴ I do not want to go into these arguments here, however, because even if there are cases of knowledge without belief, it is not clear that merely denying the necessity of true belief for knowledge permits an adequate response to the problematic reasoning. In that reasoning, the assumption that true belief is necessary but insufficient for knowledge is used to establish the further claim that it is more difficult to acquire knowledge of whether p than it is to acquire a true belief about whether p . Suppose one thinks that the following form of argument is valid: if the obtaining of condition C is necessary but insufficient for the obtaining of condition C^* , then it is more difficult for an agent to bring about condition C^* than condition C . (I consider the validity of this form of inference in § 4.3.) If one thinks that truly believing that p isn't necessary for knowing that p , then one cannot establish the claim that it is more difficult to acquire knowledge of whether p than it is to acquire a true belief about whether p simply by appealing to the validity of this form of argument. However,

¹⁴ See, for example, Radford (1966). Armstrong (1969) objects to Radford's argument.

one can still establish the claim that it is more difficult to acquire knowledge of whether p than it is to acquire *either* knowledge of whether p *or* a true belief about whether p ; trivially, having a true belief about whether p *or* knowing whether p is necessary but insufficient for knowing whether p , provided we assume that having a true belief about whether p is insufficient for knowing whether p . We can then run the problematic reasoning much as before, to establish the modified conclusion that it is rationally sub-optimal to seek knowledge of whether p rather than merely seeking *either* knowledge of whether p *or* a true belief about whether p . This conclusion also seems unacceptable. Thus, denying premise (2) will only provide a satisfying response to the reasoning if it is denied, more specifically, on the grounds that true belief is actually sufficient for knowledge.

The usual way to argue for the claim that true belief is insufficient for knowledge is to describe examples in which a subject putatively truly believes that p without knowing that p . The following passage from the *Stanford Encyclopedia of Philosophy* is illustrative:

Why not say that knowledge is true belief? The standard answer is that to identify knowledge with true belief would be implausible because a belief might be true even though it is formed improperly. Suppose that William flips a coin, and confidently believes on no particular basis that it will land tails. If by chance the coin does land tails, then William's belief was true; but a lucky guess such as this one is no knowledge. (Ichikawa and Steup 2014: §1.3)

Crispin Sartwell has argued, contrary to orthodoxy, that such examples are inconclusive. He makes two main points. The first is that in at least some examples

of this kind it is not clear that the subject really does believe the proposition at issue.

Thus, he says,

[A]rguments to the effect that some third condition is required for knowledge often play on an insufficiently rich notion of belief. Such arguments, again, often take the form simply of pointing out that a lucky guess does not count as knowledge. But of course, in the usual case, a lucky guess is not even a belief. (1991: 159)

Relatedly, he points out that it is not sufficient for a subject to believe that p that she is merely disposed to produce some conventional sign of assent in response to utterances that express the proposition that p , or disposed to act as though p is the case. Secondly, in cases in which the subject is conceded to have a true belief that p despite the fact that she apparently fails to know that p , Sartwell tries to explain away the intuition that the subject does not know that p by identifying factors that make it natural to judge that the subject doesn't know that p , despite the fact that on his view she does. For example, the intuition that a mental patient who believes that $2 + 2 = 4$ because the voices in her head have told her so doesn't know that $2 + 2 = 4$ is explained away on the grounds that, in saying that she doesn't know, we are 'reaching for a truth by means of a literal falsehood. The truth we are reaching for is that knowledge-claims made by the mental patient on controversial matters ought to be regarded as highly suspect' (1991: 162). They are suspect because the procedures the patient uses to arrive at beliefs include a procedure, listening to the voices in her head, which is liable to produce false beliefs, and false beliefs cannot constitute knowledge.

Sartwell is quite right to point out that lucky guesses are not beliefs, and that merely acting as though p is the case, or producing conventional signs of assent in response to expressions of the proposition that p , is not sufficient for believing that p . However, these claims show at most that philosophers ought to be a bit more careful in their choice and description of counterexamples to the sufficiency of true belief for knowledge, to make sure that it is clear that the relevant subjects really do satisfy the conditions for belief. Furthermore, Sartwell's explanation of why it is intuitive to judge that the 'mental patient' does not know that $2 + 2 = 4$ clearly doesn't apply across the board. There are scenarios in which we are inclined to judge that the subject doesn't know that p , despite the fact that she truly believes that p , in which there is no insinuation that the subject has a general tendency to arrive at beliefs by unreliable methods. Gettier cases are an obvious example. It seems to me, therefore, that Sartwell fails to identify compelling reasons to question our intuitions about purported counterexamples to the thesis that truly believing that p is insufficient for knowing that p .

As an aside, it is interesting to note that in the course of evaluating an argument for the claim that knowledge is at least justified true belief, Sartwell inadvertently brings to light a consideration that suggests that true belief is insufficient for knowledge. The argument for the claim that knowledge is at least justified true belief is that 'it is always legitimate, when someone claims to know something, to ask *how* she knows it' (1991: 159). Although Sartwell doesn't spell out the argument explicitly, presumably the idea is that when you ask someone how she knows that p you are asking her to justify her belief that p , and it wouldn't always be legitimate to ask her to do this unless knowing that p entailed having a justified belief that p . In response to this argument, Sartwell questions the assumption that it is always legitimate to ask

someone who claims to know something *how* she knows it, but he also argues that, even if this assumption were true, it ‘would not in itself establish a disanalogy between knowledge and mere true belief, or even between knowledge and mere belief, or even between knowledge and mere assertion’ (ibid.). In all these cases, he says, the request for justification is equally legitimate—it is quite familiar, after all, to ask ‘Why do you believe that?’, or ‘What’s your reason for saying that?’ Sartwell notes a further disanalogy, however, between the case of knowledge and the case of belief or assertion. If someone cannot adequately answer the question ‘How do you know that *p*?’, it might be appropriate to respond by saying, ‘You don’t know after all!’ However, if someone cannot adequately justify his belief or assertion it is not normally appropriate to respond by saying ‘You didn’t really assert that *p* after all!’, or ‘You don’t really believe that *p*!’ If someone cannot adequately respond to the request for justification this may show that he does not know that *p*, but it does not show that he does not believe that *p* or has not asserted that *p*. Sartwell insists, though, that ‘no disanalogy is established between knowledge and *true* belief [my emphasis], if the demand for a justification is thought of as a demand that the proposition which one claims to know be supported, that is, that reasons should be given to regard it as true’ (ibid.). This is where Sartwell inadvertently brings to light a consideration that speaks in favour of the insufficiency of true belief for knowledge. Suppose, for example, that I wish to know whether *p*, and upon asking you about the matter you confidently assert to me that *p*. When I ask you how you know that *p*, however, the reasons you give are quite inadequate. Your failure to provide an adequate response to the request for justification raises doubts both about whether you know that *p* and about whether you truly believe that *p*, but there is an important difference: whereas your failure to provide an adequate justification may

show, quite decisively, that you do not really know that p , it shows at most that I ought to suspend judgement about whether you truly believe that p . If truth belief really were sufficient for knowledge, this difference would not exist. Things are even clearer in the situation in which I know that p and I wish to check whether you know that p —that is to say, when I am in the situation of the examiner. Upon establishing that you take yourself to know that p , I may ask, ‘How do you know that p ?’, or ‘Why do you believe that p ?’ Your failure to answer either question in a satisfactory way may show that you do not really know that p , but it does not speak against the fact that you truly believe that p . So, contrary to what Sartwell says, there is an important disanalogy between knowledge and true belief with respect to the request for justification, which suggests that true belief is not sufficient for knowledge.

In this subsection I have argued that we cannot respond to the problematic reasoning by rejecting its second premise, that having a true belief about whether p is necessary but insufficient for knowing whether p . Merely rejecting the claim that having a true belief about whether p is necessary for knowing whether p is inadequate, because even if this is right it remains the case that having a true belief about whether p *or* knowing whether p is necessary but insufficient for knowing whether p (assuming that true belief is insufficient for knowledge). It can still be argued, therefore, that an enquirer ought at most to adopt the disjunctive objective of arriving at a true belief about, or knowledge of, whether p . The more important claim for the purposes of the problematic reasoning is that having a true belief about whether p is insufficient for knowing whether p . Sartwell contends that the standard argument for this claim—that a subject’s belief about whether p might be true by luck, and so fail to constitute knowledge of whether p —is inconclusive. However, Sartwell’s objections, as we have just seen, are unconvincing. It is true that philosophers sometimes lazily choose

examples (e.g. a coin toss) in which it is implausible that the subject genuinely believes the target proposition (e.g. that the coin will land tails), but it is not difficult to construct cases in which the subject clearly believes the target proposition. Sartwell's diagnosis of why it sometimes sounds right to say that a subject who truly believes that p does not know that p —that the knowledge claims made by the subject are suspect because she is arriving at beliefs by unreliable methods—does not apply in all cases, because sometimes it sounds right to deny that a subject knows that p (e.g. in a Gettier case) even when there is no insinuation that the subject is arriving at beliefs by unreliable methods. Thus, for all Sartwell says, the standard argument for the claim that true belief is insufficient for knowledge remains compelling. Furthermore, in his attempt to explain why, even supposing that it is correct that 'it is always legitimate, when someone claims to know something, to ask *how* she knows it', this still would not show that true belief is insufficient for knowledge, Sartwell inadvertently brings to light a further consideration that supports the insufficiency claim.

4.3 Does (3) Follow From (2)?

In the problematic reasoning, the role of the assumption that having a true belief about whether p is necessary but insufficient for knowing whether p is to establish the further thesis that it is more difficult for an enquirer to arrive at knowledge of whether p than it is for the enquirer to arrive at a true belief about whether p . It is only if we have this further thesis, alongside the claim that having a true belief about whether p is just as good as knowing whether p , that we can exploit the principle about rationality expressed by (4) to obtain the conclusion that it is rationally sub-optimal for an enquirer to aim at the acquisition of knowledge. Let us say that condition C is more demanding than condition C^* just in case, necessarily, if

condition C obtains then condition C^* obtains, but it is not the case that, necessarily, if condition C^* obtains then condition C obtains. If having a true belief about whether p is necessary but insufficient for knowing whether p , the condition that one knows whether p is more demanding, in this sense, than the condition that one has a true belief about whether p . But does it follow from this that it is harder for an enquirer to arrive at knowledge of whether p than it is for an enquirer to arrive at a true belief about whether p ? Are more demanding conditions necessarily more demanding on the agent?

The correct response to this question is ‘No, of course not’. My bedroom has white walls. Suppose that R is the condition that BW’s bedroom has red walls and PC is the condition that BW’s bedroom has walls of a primary colour. The obtaining of condition PC is necessary but insufficient for the obtaining of condition R : necessarily, if R obtains then PC obtains,¹⁵ but it is not the case that, necessarily, if PC obtains then R obtains. Thus, R is a more demanding condition than PC in the sense identified above. However, it obviously does not follow from this that it is more difficult for me to paint my walls red than it is for me to paint my walls a primary colour, in the ordinary sense of ‘difficult’. A more demanding condition is not necessarily more demanding on the agent.

It does not follow from the fact that true belief is necessary but insufficient for knowledge that it is more difficult for an enquirer to arrive at knowledge of whether p than it is for an enquirer to arrive at a true belief about whether p . The routes to knowledge may be no more demanding than the routes to true belief. However, it is

¹⁵ This claim depends on the assumption that red is necessarily a primary colour, but the point I am making does not depend on this assumption; if someone wishes to dispute the claim that red is necessarily primary, they should simply understand PC to be the disjunctive condition that BW’s bedroom has walls that are red, blue or yellow.

important to distinguish the claim that finding out (i.e. acquiring knowledge of) whether p is *always* more difficult than acquiring a true belief about whether p from the claim finding out whether p is *sometimes* more difficult than acquiring a true belief about whether p . The first claim, as we have just seen, is not entailed by the premise that the condition that one knows whether p is more demanding than the condition that one has a true belief about whether p . Furthermore, it is not obvious how else one might argue for this general claim. One might still think, however, that this premise can be used to argue for the weaker claim that finding out whether p is sometimes more difficult than acquiring a true belief about whether p . After all, if one can have a true belief about whether p without knowing whether p , then presumably there will at least sometimes be procedures one can follow that would result in one having a true belief about whether p but not in one knowing whether p . Furthermore, presumably at least sometimes these procedures will also be easier for an enquirer to follow than any procedure the enquirer could follow to arrive at knowledge of whether p .

It is fairly straightforward to think of examples in which these conditions are met. Imagine that May's calculator malfunctions in the following way: although the other buttons work normally, pressing the '=' sign causes it to display the figure '56,088'. May does not know this—as far as she is concerned, her calculator works just like any other. Wanting the truth about the product of 123 and 456, May taps the problem into the calculator and, upon reading the display, forms the true belief that $123 \times 456 = 56,088$. I take it that, although May now truly believes that $123 \times 456 = 56,088$, she does not know that $123 \times 456 = 56,088$; it is pure luck that the belief she has acquired from the calculator is true. Furthermore, the procedure May used to arrive at a true belief was easier for May to follow than any procedure she could have

followed that would have resulted in her knowing that $123 \times 456 = 56,088$. Working out the problem on paper would have required significantly more time, skill and mental effort, for example. Although the example is an artificial one, it is plausible that it will quite often be the case that there are procedures available to an enquirer that would result in him holding a mere true belief about the answer to some question, and which are in some sense easier for him to follow than any procedure he could follow that would result in him knowing the answer. Often, for example, the easiest way for an enquirer to arrive at a true belief about the answer to a question will be by consulting another person, but if this other person has a mere true belief about the answer to the question, then it is plausible that the enquirer will not acquire knowledge that p in forming the true belief that p in response to the speaker's assertion that p .

What are the implications of these points for the problematic reasoning? (3), as it stands, is unjustified: we cannot infer from the premise that true belief is necessary but insufficient for knowledge that, quite generally, it is more difficult for an enquirer to arrive at knowledge of whether p than it is for an enquirer to arrive at a true belief about whether p . However, (2) can be used to support a weaker version of (3): at least sometimes, there is a procedure available to an enquirer that would result in him having a mere true belief about whether p that is easier for the enquirer to follow than any procedure that would result in him knowing whether p . We can then run the reasoning as before to obtain the weaker conclusion that, at least in cases of this kind (in which, we might add, the stability of true belief is unimportant), it is rationally sub-optimal for the enquirer to adopt the aim of acquiring knowledge of whether p .

Is this weaker conclusion still problematic? The idea that it isn't worth seeking knowledge of the answer to some question in certain circumstances is not an unfamiliar one. Sometimes a mere estimate about the answer to the question will do. However, I think there is something odd about the idea that an enquirer shouldn't seek knowledge of the answer to some question *because* it is easier for him to arrive at a true belief about the answer. Like considerations to do with the stability of true belief, there does not appear to be any trace of this consideration in our ordinary thinking about enquiry. I shall be arguing that we should not accept even this weaker conclusion.

5. Does (4) Apply in the Case of True Belief and Knowledge?

My aim in this section is to highlight what seems to me the most significant objection to the problematic reasoning. In general terms, the objection is that certain conditions whose satisfaction we presuppose in assenting to (4) cannot be satisfied in the case of knowledge and true belief. It follows that (4) cannot be used to force the conclusion that it is rationally sub-optimal for an enquirer to seek knowledge, given that he could merely aim to acquire true beliefs. *Why* these conditions cannot be satisfied in the case of knowledge and true belief is not something I will attempt to explain just yet. This is because the correct explanation of why this is so depends on claims about the nature of belief and judgement that I haven't established so far. I will return to this issue in Chapter 4, after I have discussed the relevant issues in Chapters 2 and 3. For now, my aim is just to show that there are further conditions whose satisfaction we take for granted in accepting (4), and that it is at least not obvious that these conditions are satisfied in the case of true belief and knowledge.

For ease of reference, here is principle (4) again:

If objective *X* is just as good as objective *Y*, but it is more difficult for the agent to achieve objective *Y* than objective *X*, then it is rationally sub-optimal for the agent to adopt *Y* as his aim; he should, at most, adopt *X* as his objective.

(4) is made attractive by cases like the following. I am taking part in an archery competition. It is almost the end of the match. My opponent, who has used all his arrows, is slightly ahead of me, but I still have one more shot. I have a choice of two targets, one closer than the other. Hitting the farther target earns me more points, but I will get enough points to beat my opponent if I hit only the closer one. Which target should I aim for? It seems obvious that I should shoot at the closer one; that way, I maximize my chances of winning. This verdict is in line with principle (4). Hitting the closer target is just as good as hitting the farther target, from the point of view of winning the match, but it is more difficult for me to hit the farther target; hence, I should aim for only the closer one. Before we conclude that this principle can be used to obtain the conclusion that in some circumstances enquirers ought to aim only for true beliefs, however, we should consider whether the application of the principle in this case depends on features that are absent in the case of true belief and knowledge.

One very basic feature of the archery example is that I recognize that there are ways in which I can attempt to achieve the more modest objective that are not ways in which I can achieve the more demanding objective. In particular, I recognize that shooting at the closer target is a way of hitting the closer target but not a way of hitting the farther target. It is because I am capable of recognizing that there are ways in which I can attempt to achieve the more modest objective that are not ways in which I can achieve the more demanding one that I can potentially avail myself of

less exigent means by adopting only the more modest end. Is this condition satisfied in the case of true belief and knowledge? That is to say, can I recognize that a certain procedure is potentially a way of arriving at a true belief about whether p but not a way of finding out whether p ? In a theoretical sense, I can. I recognize that inferring that p from premises that include one or more propositions that I falsely believe to be the case is a way in which I might arrive at a true belief about whether p , but not a way in which I might find out whether p . However, it does not follow from this that I can follow this procedure *thought of as such* in an attempt to arrive at a true belief about whether p . There appears to be something in the thought that I can use a certain procedure to attempt to arrive at a true belief about whether p only if I regard that procedure as a potential way of finding out whether p . If this is right, then in the case of true belief and knowledge I cannot avail myself of less exigent means by adopting the less demanding end; even if I merely adopt the aim of true belief, I can pursue my aim only by using methods that I regard as potential ways of discovering, or finding out, the truth.

The example I described in the previous section provides an illustration of this point. If May knew that her calculator was malfunctioning in the way I described, she could not use the calculator to arrive at beliefs about the answers to arithmetical questions. One might think that May can arrive at beliefs about the answers to arithmetical questions using her calculator only because she regards the calculator as a way of finding out the answers to those questions. This is why she cannot make life easier for herself by adopting only the aim of arriving at a true belief about the answer. Even if she adopts this aim, the only way she can attempt to achieve her aim is by using methods that she regards as ways of finding out the answer. It is easier for her to arrive at a true belief about the answer than knowledge only because she

mistakenly regards a certain method as a way of finding out the answer, which in fact happens to yield a true belief about the answer. Although it can be argued, on the basis of the claim that true belief is necessary but insufficient for knowledge, that it is sometimes easier for an enquirer to arrive at true beliefs than at knowledge, it does not follow that it is possible for an enquirer to recognize *as such* the more efficient procedures that deliver only true belief.

To forestall misunderstanding, I should emphasise that I am not here endorsing the claim that I can use a certain procedure to arrive at a true belief about whether p only if I regard that procedure as a way of finding out whether p . As we will see in Chapter 4, there appear to be counterexamples to this claim. All I mean to assert is the following. Firstly, one thing we take for granted in appealing to (4) to explain why an agent should adopt only a more modest objective is that the agent is able to recognize that there are ways in which he can attempt to achieve the more modest objective that are not ways in which he can achieve the more demanding objective. Secondly, it is not obvious that this condition can be satisfied, at least in a straightforward way, in the case of true belief and knowledge. In fact, the proposal that an enquirer can use a certain procedure to arrive at a true belief about whether p only if he regards that procedure as a way of finding out whether p is too simple. Nevertheless, it gets at something important, and helps us to see why something that we take for granted in applying principle (4) may be problematic in the case of true belief and knowledge. To make further progress here, however, we need to have a deeper understanding of the nature of belief and judgement. This is what I shall attempt to provide in the next two chapters.

6. Summary

In this chapter I have distinguished two different challenges someone might have in mind in pressing the activity question. The first, which is raised by Williams in *Descartes*, was based on the argument that since true belief is just as true as knowledge, an enquirer will have got what he wants if he merely arrives at a true belief about the answer to his question. We saw that this argument is fallacious. There is therefore no need to explain why an enquirer who is initially specified only as wanting a true belief about whether p in fact wants to know whether p . Furthermore, Williams' attempt to explain why this is so appears to be decisively flawed. We should reject Williams' challenge. The second challenge that someone might have in mind in pressing the activity question is more interesting. I expressed this challenge in the form of a paradoxical argument, the problematic reasoning. I went on to argue that various responses to this argument are inadequate. Although there is a respect in which knowledge is superior to mere true belief—it is less likely to be undermined by future evidence, and therefore more stable—this feature of knowledge is not always of value to enquirers. Disputing the claim that having a true belief about whether p is necessary but insufficient for knowing whether p does not appear to be promising either. The inference from this claim to the further idea that it is harder for an enquirer to arrive at knowledge than at true beliefs is problematic, but it is still possible to defend on this basis the weaker claim that there is at least sometimes a procedure an enquirer can use to arrive at a true belief about whether p that is easier for him to follow than any procedure he could use to arrive at knowledge of whether p . The real problem with the reasoning, I proposed, lies in the idea that the principle expressed by (4) can be applied in the case of true belief and knowledge. I argued that, in applying that principle, one thing we presuppose is that

it is possible for an agent to recognize that there are ways in which he can attempt to achieve a more modest objective that are not ways in which he can achieve a more demanding objective. I went on to suggest that it is not obvious that this presupposition holds in the case of true belief and knowledge. The correct response to the reasoning lies in questioning this presupposition in relation to enquiry. To make further progress here, however, we need a deeper understanding of the nature of belief and judgement.

Chapter 2

Believing, Assuming and Judging

In the Introduction I argued that philosophical investigation of the value of knowledge should begin with the activity question: why do we desire and seek knowledge, and not merely true beliefs, given that true belief is just as true, and just as useful, as knowledge? It is not entirely obvious why the (alleged) facts that true belief is just as true, and just as useful, as knowledge, should be taken to show that enquirers should at most be interested in acquiring true beliefs. In the previous chapter I suggested that someone might be motivated to press the activity question by the problematic reasoning. After considering several different replies to this reasoning, I concluded that the most promising response is to question a principle of rationality on which the reasoning depends. According to that principle, embodied in the fourth premise of the reasoning, if objective *X* is just as good as objective *Y*, but it is more difficult for an agent to achieve objective *Y* than objective *X*, then it is rationally sub-optimal for the agent to adopt objective *Y* as his aim; he should, at most, adopt *X* as his objective. We seem to be disposed to appeal to this principle (or something like it) in explaining why, in certain circumstances, an agent ought to adopt one objective rather than another. It is consistent with this, however, that the application of the principle in these circumstances depends on certain conditions being met that aren't satisfied in the case of true belief and knowledge. If these conditions aren't satisfied in the case of true belief and knowledge, then the principle cannot be used to show that enquirers ought only to be interested in acquiring true beliefs.

Whether or not it is legitimate to appeal to principle (4) in relation to true belief and knowledge is not a straightforward matter to determine. Rather than tackling this question head on, I want to begin by considering a closely related issue that is also of considerable independent interest, namely, the nature of belief. Issues about the nature of belief and judgement will occupy us for the next two chapters. My ultimate aim in developing an account of belief and judgement is to show how it can be used to explain why principle (4) does not apply in the case of true belief and knowledge. But this will have to wait until Chapter 4. Until then, our attention will be directed away from the problematic reasoning onto more basic issues about the metaphysics of belief, judgement, and related cognitive mental states and events.

Although I will be allowing the problematic reasoning to recede into the background for the time being, the more general theme of the importance of considering the position of an enquirer will continue to play a major role in my discussion. Enquirers are cognitively sophisticated subjects, and, as I explain in § 1, focussing on cognitively sophisticated subjects in a philosophical investigation of the nature of belief brings certain issues to the fore. In particular, issues arise about how the beliefs of such subjects are related to the sophisticated forms of mental activity they are able to engage in, including forms of practical reasoning and theoretical deliberation, and about how we are to mark distinctions between belief proper and other forms of acceptance that such subjects are capable of adopting. § 2 considers the relationship between believing that *p* and consciously judging that *p*. I argue that a necessary condition of an enquirer having the outright (as opposed to the repressed) belief that *p* is that he is disposed to judge that *p*. §§ 3–5 examine the nature of conscious judgement by investigating the difference between judging that *p* in one's reasoning and supposing that *p* for the sake of argument. In § 3 I argue that

we cannot account for the difference simply in terms of the idea that judgement, unlike supposition, involves the acquisition/manifestation of the corresponding belief. Such an account cannot explain why it is hard to make sense of the idea that a subject might be in error about the nature of his own conscious thinking by, for example, taking himself to be supposing that p for the sake of argument when he is really judging that p in his reasoning. § 4 considers the proposal, made by Shah and Velleman (2005), that the difference between judging that p and supposing that p for the sake of argument is a matter of a difference in the aim with which the subject affirms the proposition that p in his reasoning. Although this proposal is, on the face of it, well-placed to explain the ‘privileged’ first-person epistemology of conscious thinking, it is implausible that when a subject judges that p he affirms that p with a distinctive type of intention—e.g. the intention to affirm that p only if p , or to affirm that p only if he knows that p . § 5 introduces Matthew Soteriou’s (2013) account of the difference between judgement and supposition for the sake of argument, and shows how this account can be used to explain why it is hard to make sense of the idea that a subject might be in error about the nature of his own conscious thinking. The key thought here is that whether one is judging that p in one’s reasoning or supposing that p for the sake of argument is partly a matter of what one regards as the source of the constraint of treating p as true in one’s reasoning (in particular, on whether or not one regards the constraint of treating p as true as a self-imposed constraint), and how one regards that constraint will be constitutively connected with one’s higher-order beliefs about what one is up to in reasoning on the assumption that p . This is why it is hard to make sense of the idea that a subject might be in error about whether he is judging that p or supposing that p for the sake of argument in reasoning on the assumption that p . § 6 returns to the issue of how we are to

distinguish belief from other forms of acceptance that a sophisticated enquirer is capable of engaging in. Although it may be possible to account for the ‘beliefs’ of more primitive, non-judging subjects in terms of what Velleman (2000) calls the ‘purely motivational conception’ of belief, this conception is inadequate as an account of the beliefs of enquirers. When we speak of belief in relation to enquirers, we normally intend to pick out a state that is distinct from the state one is in when one merely assumes for practical purposes that a certain proposition is true, but the state of assuming that p for practical purposes also plays the motivational role identified by the purely motivational conception of belief. I suggest that one important difference between believing that p and assuming that p for practical purposes is that when one reasons on the assumption that p because one is assuming that p for practical purposes, one once again regards the constraint of treating p as true as a self-imposed constraint. By contrast, when one reasons on the assumption that p because one believes that p one judges that p in one’s reasoning, and when one judges that p in one’s reasoning one does not regard the constraint of treating p as true as a self-imposed constraint.

1. Enquiry and Belief

Our present interest in the nature of belief stemmed from an interest in resolving philosophical issues about the value of knowledge. If the claims I have made about the relationship between the activity question and the comparative state question are correct, then the kind of believer we should be focussing on, in considering the nature of belief in the context of concerns about the value of knowledge, is a cognitively sophisticated enquirer. This is an important point, so let me say a little more about what I mean in saying this, and why I take it to be true.

In the Introduction I argued that the activity question ('why do we desire and seek knowledge, and not merely true beliefs, given that true belief is just as true, and just as useful, as knowledge?') is more basic than the comparative state question ('why is it better to know that p than to have a mere true belief that p , given that true belief is just as true and just as useful as knowledge?'). The only obvious reason to accept the presupposition of the comparative state question, that it is better for a subject to know that p than it is for the subject to have a mere true belief that p , is that we have to make this claim in order to provide an adequate response to the activity question. Since, however, it is not obvious that we do have to endorse this claim in order to respond to the activity question (and, indeed, hard to see how this claim can be vindicated, in a sufficiently strong form, to provide an adequate response to that question), philosophical reflection about the value of knowledge should begin with the activity question. Now, unlike the comparative state question, which can be posed in relation to any subject for whom we can mark a distinction between mere true belief and knowledge, the activity question explicitly concerns an enquirer. Enquirers are cognitively sophisticated subjects. A subject who is capable of engaging in enquiry is a subject who is capable of recognizing that he lacks the truth about the answer to a certain question, and of adopting purposive means in order to obtain it. These general capacities plausibly entail the enquirer's possession of more specific abilities. An enquirer is capable of adopting higher-order attitudes about his own mental states—e.g. of recognizing that he does not know whether p , or have a true belief about whether p . He is also capable of adopting purposive means in order to arrive at a true belief about, or knowledge of, whether p . Arguably, this implies the capacity to engage in forms of practical reasoning and theoretical deliberation, which in turn imply that the enquirer has the capacity to make conscious decisions

and judgements. It might be thought, furthermore, that given that the enquirer is capable of recognizing that he does not know, or have a belief about, the answer to a certain question, he will have the capacity to engage in forms of acceptance besides believing and knowing. Recognizing that he does not know the answer to a question, an enquirer may, for example, decide to make a certain assumption about the answer to that question for practical purposes, or decide to accept a certain assumption about the answer for the sake of argument, or decide to hazard a guess about the answer.

There are large issues here. To give a more detailed philosophical account of how these various capacities are related to one another—e.g. of how the capacity for higher-order epistemic attitudes is related to the capacity to engage in forms of acceptance besides believing and knowing, or of how the capacity for conscious judgement and decision is related to the capacity to engage in theoretical reasoning and practical deliberation—would require much further work. I shall have more to say about some of these issues later. For now, however, the important point is that a subject capable of engaging in enquiry is a subject who has these capacities, and it seems plausible, on the face of it, to think of these capacities as being related to one another in significant and interesting ways.

Focussing on cognitively sophisticated enquirers brings certain issues about the nature of belief to the fore. Here, I will concentrate on two such issues in particular. Firstly, as I said, an enquirer is a subject who has the capacity to engage in forms of conscious thinking, such as practical reasoning and theoretical deliberation. One question we can ask, then, is how an enquirer's beliefs are related to his conscious thinking. A number of authors have suggested, in connection with this issue, that there is an intimate connection between believing that *p* and consciously judging that *p*. I shall be considering this matter in the next section. Secondly, an enquirer is a

subject who has the capacity to engage in forms of acceptance besides knowing, believing and judging. He can, for example, make assumptions for practical purposes, or assumptions for the sake of argument. Once this has been noted, it is natural to ask how believing that p differs from merely assuming that p for practical purposes, or how judging that p differs from merely supposing that p for the sake of argument.¹⁶ These issues will be the focus of §§ 3–6. In effect, I will be arguing that the key to resolving the second issue lies in resolving the first. At least part of what distinguishes having the outright belief that p from merely assuming that p for practical purposes, I will suggest, is that a subject who has the outright belief that p is disposed to treat the proposition that p in a distinctive way in his conscious thinking—he is disposed to judge that p . The significance of this claim is not apparent at first sight, however, since it may be said that a subject who is merely assuming that p for practical purposes is also disposed to ‘judge’ that p . I will clarify what is distinctive of the kind of judging associated with believing by considering the difference between judging that p in one’s reasoning and merely supposing that p for the sake of argument. (The distinctive nature of judging will continue to concern us in the next chapter, where I will be concerned with the question of how judging that p differs from merely guessing that p .)

To begin with, how should we conceive of the relationship between believing that p and consciously judging that p , when the kind of believer we are concerned with is an enquirer?

¹⁶ Earlier discussions of this kind of issue include Cohen (1992), Velleman (2000), and Shah and Velleman (2005).

2. Belief and Judgement

Believing that p is a mental state. As such, belief is distinct from the thoughts one has when one is consciously thinking about something—e.g. the thoughts one has when one attempts to calculate the product of 13 and 17 in one's head. Such thoughts are events that occur at certain times, rather than states that obtain over intervals of time. Although believing that p is a state and not an event, it is commonly assumed that there is a type of occurrent thought that is closely associated with belief. Philosophers normally reserve the term 'judgement' or 'conscious judgement' to refer to this type of mental event. When you work out that the product of 13 and 17 is 221, you consciously judge, in this sense, that $13 \times 17 = 221$. Likewise, when I recall that I had Weetabix for breakfast this morning, I consciously judge that I had Weetabix for breakfast this morning. (I do not mean to imply by these examples that this kind of mental event occurs only when one succeeds in (e.g.) working out or remembering something. If I make a mistake in my calculations, I might falsely judge that $13 \times 17 = 227$; likewise, if I misremember what I had for breakfast, I might falsely judge that I had Shreddies.)

What is the relationship between believing that p and consciously judging that p , in this sense? Clearly, one can believe that p at a time when one is not judging that p . My belief that Madrid is the capital of Spain does not go out of existence when I stop consciously thinking that Madrid is the capital of Spain, and it continues to obtain even when I am in a state of dreamless sleep, and am thus not consciously thinking about anything. However, it is consistent with this that believing that p at least requires being disposed to judge that p , when, for example, the question of whether p arises. I am disposed to judge that Madrid is the capital of Spain in response to the

question, 'What is the capital of Spain?' Should we then accept this more modest proposal?

Quassim Cassam identifies two common objections to this suggestion in the following passage:

[I]t might be suggested that someone who believes that p must at least be disposed to judge that p . This has some plausibility, but is still not obviously correct. One can imagine someone who finds it psychologically impossible mentally to affirm to herself that p but who nevertheless believes that p . She has no disposition to judge that p , even when explicitly asked whether p , but she does in fact believe that p . If some non-human animals are capable of belief but not judgement then that would be another reason not to regard the belief that p as a disposition to judge that p , or as requiring the disposition to judge that p . (2010: 83)

Do these objections show that we should reject the claim that believing that p involves being disposed to judge that p ? In particular, do they show that we should reject this claim even when the kind of believer we are concerned with is an enquirer? It should be clear that concerns about the possibility that animals have beliefs despite being incapable of making judgements are irrelevant in the present context. The kind of believer we are concerned with is an enquirer, and an enquirer is a subject with the capacity for judgement. Even if it is possible for an animal without the capacity for judgement to believe that p without being disposed to judge that p , it does not follow that it is possible for an enquirer to believe that p without being disposed to judge that p . However, this still leaves the other kind of objection identified by Cassam, that of a person who believes that p but nevertheless find it

psychologically impossible to judge that p . This subject does have the capacity for judgement, and more generally (we may presume) the capacity to engage in enquiry, so we cannot duck this objection by pointing out that we are presently only interested in the conditions under which an enquirer believes that p .

Cassam does not go on to explain, in more detail, the type of case he has in mind here, but presumably he is thinking of examples of so-called 'repressed belief'. Here is an example of this possibility. The manager of an engineering firm might profess to think that female engineers are just as capable as their male counterparts. Nevertheless, it might be glaringly obvious, from the way he assigns tasks to his employees, that he does not really believe this. Whenever a particularly challenging or important job comes up, he assigns it to one of the men, and the more basic, menial tasks are generally delegated to the women. His actions make it clear that his actual belief about the competence of female engineers is the very opposite of his professed belief, and that what he really thinks is that female engineers are less capable than male ones. The manager needn't be being intentionally dishonest about what he believes; he is not merely pay lip service to gender equality in the workplace whilst secretly thinking that it is obvious that male engineers are better, and that everyone would be able to see this were it not for the undue influence of feminism. It is rather that his decisions are influenced by an unconscious prejudice that he has. Thus, although the manager is not disposed to judge that male engineers are better than their female counterparts, and would indeed find it psychologically impossible to do so, there is nevertheless a sense in which he believes this to be the case.

As Soteriou observes (2013: 349), examples like this one draw our attention to the fact that the psychological state of belief can be manifested in different ways, of which consciously judging that p is only one. Someone can also manifest the belief

that p in his actions and decisions. This being the case, it appears to be conceivable that the best explanation of someone's actions and decisions might be one that attributes him the belief that p , even though he is not disposed to judge that p , and is perhaps even disposed to judge that *not-p*.

I do not wish to dispute the claim that cases like the one I have just described are cases of belief. What such cases should prompt us to do, however, is mark distinctions between different *kinds* of belief. Cases of repressed belief are clearly importantly different from cases of ordinary factual belief—e.g. the belief that Madrid is in Spain, or the belief that it rained on Monday. Plausibly, the difference between ordinary factual belief (or 'outright belief') and repressed belief is that the former but not the latter involves a disposition to make the corresponding judgement. Having marked the distinction between outright and repressed belief, a question we might go on to consider is what kind of belief is at issue in the problematic reasoning. In Chapter 1 I introduced the reasoning using the general notion of belief, but it might be thought that the reasoning is in fact only plausible in relation to a particular kind of belief. This is what I now want to suggest.

When a subject attempts to find out or work out the answer to a question, his aim is clearly to arrive at *explicit* knowledge of the answer to that question. Thus, in thinking about whether he should not really go after some 'lesser' state, which is allegedly just as good but easier to attain, what we should be considering is whether this lesser state is just as good but easier to attain than the state of having explicit knowledge of the answer to some question. Having a repressed true belief about the answer to some question is obviously not as good as having explicit knowledge of the answer to that question. Although repressed beliefs influence the behaviour of the agent, they do not dispose the agent to treat their contents as reasons for belief

and action. Repressed beliefs do not dispose the agent to treat their contents as premises in his practical reasoning and theoretical deliberation. Consequently, they are significantly less useful to the agent. They are not available to guide his planning and decision making, or to guide further enquiry. It is really only outright belief that may be claimed, with some credibility, to be as good as explicit knowledge. Thus, the first premise of the problematic reasoning is plausible only if 'belief' is understood to mean *outright* belief.

Once the problematic reasoning is understood to concern outright belief, an issue it is relevant to consider in relation to that reasoning is the conditions under which someone will acquire an outright belief about whether p as a consequence of engaging in enquiry into whether p . Since having the outright belief that p involves being disposed to consciously judge that p , acquiring an outright belief about whether p involves acquiring a disposition to consciously judge whether p . Thus, in considering the conditions under which someone will acquire an outright belief by engaging in enquiry, one thing that is relevant is the nature of conscious judgement itself. This issue will be central to the rest of this chapter.

The nature of judgement is not a straightforward matter to determine. However, an apparently secure starting point is the thought that judging that p contrasts with other ways of consciously affirming the proposition that p , such as supposing that p for the sake of argument. What is the difference between the situation in which one judges that p in one's reasoning and the situation in which one merely supposes that p for the sake of argument? Matthew Soteriou discusses this question in Chapter 11 of *The Mind's Construction*. There, he argues that a significant difference between the situation in which one supposes that p for the sake of argument and the situation in which one judges that p in one's reasoning is that, in the case of the former but not

the latter, the constraint of treating p as true is a self-imposed constraint on one's reasoning. I shall consider Soteriou's proposal in more detail in a moment. Before I do so, however, I want to consider some more simple proposals that might be made about the difference between judgement and supposition. Doing so helps to bring out the motivation for Soteriou's view.

3. Judging, Supposing and the Acquisition/Manifestation of Belief

An obvious way of attempting to capture the difference between judging that p and supposing that p for the sake of argument is to appeal to a difference in the way in which these cognitive mental events are related to the subject's cognitive mental states. It is sometimes suggested, for example, that judging that p involves the acquisition of the belief that p . One certainly doesn't acquire the belief that p in merely supposing that p for the sake of argument. This proposal is incompatible with the idea that a subject can judge that p when she already believes that p , however. When I recall that Madrid is in Spain I consciously judge that Madrid is in Spain, but I obviously do not acquire the belief that Madrid is in Spain; I can recall that this is so only if I already believe it. We can deal with this problem by refining the original proposal about how we are to distinguish between judgement and supposition, in the following way: when a subject judges that p he acquires *or manifests* the belief that p , but he does not acquire or manifest the belief that p merely in supposing that p for the sake of argument. It can now be said that although I do not acquire the belief that Madrid is in Spain when I recall that Madrid is in Spain, I do manifest this belief.

Arguably, this proposal succeeds in identifying a genuine difference between the situation in which one judges that p in one's reasoning and the situation in which one merely supposes that p for the sake of argument. What I wish to deny, however, is

that we can provide a *complete* account of the difference between judgement and supposition simply by appealing to the idea that the former, but not the latter, involves the acquisition/manifestation of the corresponding state of belief. The problem with this view is that it cannot explain a certain feature of the subject's awareness of the nature of his own conscious thinking. Before I go on to develop this objection, however, I want to briefly consider a more basic concern one might have about the acquisition/manifestation proposal.

The claim that one cannot form or manifest the belief that p merely in supposing that p for the sake of argument seems unobjectionable. The more controversial aspect of the current proposal is that judging that p invariably involves the formation or manifestation of the belief that p . A number of prominent authors have argued that it is possible for a subject to judge that p even if she does not believe that p . The following example, due to Christopher Peacocke, is frequently cited as a justification for this claim:

Someone may judge that undergraduate degrees from countries other than her own are of an equal standard to her own, and excellent reasons may be operative in her assertions to that effect. All the same, it may be quite clear, in decisions she makes on hiring, or in making recommendations, that she does not really have this belief at all. (1999: 242–243)

If it is possible to judge that p when one does not even believe that p , then it is certainly possible to judge that p without forming or manifesting the belief that p . Peacocke's example is not decisive, however. One might think that the subject he describes holds *both* the belief that undergraduate degrees from countries other than her own are of an equal standard to her own *and* the repressed belief that

undergraduate degrees from countries other than her own are *not* of an equal standard to her own. Admittedly, this stance implies that it is possible to simultaneously believe that p and to believe that $not-p$, but then again it is not obvious that we should deny this possibility. It seems to me, therefore, that one might defensibly maintain the view that an important difference between judging that p in one's reasoning and supposing that p for the sake of argument is that, in the case of the former but not the latter, the subject acquires or manifests the belief that p .

A more serious reservation about the proposal, I think, is that even if it succeeds in identifying a genuine difference between judgement and supposition, in an unsupplemented form it cannot explain certain features of judgement and supposition that a philosophical account of these cognitive mental events should be able to explain. Obviously, this general complaint might be developed in a number of different ways. Here I will focus on issues to do with the thinker's awareness of the nature of her own conscious thinking. It is hard to make sense of the idea, I suggest, that a thinker might mistakenly take himself to be supposing that p for the sake of argument when he is really judging that p in his reasoning. Likewise, it is difficult to make sense of the idea that a thinker might mistakenly take himself to be judging that p in his reasoning when he is really only supposing that p for the sake of argument. How are we to account for these facts? Can we explain them if all we have at our disposal is the idea that one act, but not the other, involves the acquisition/manifestation of the corresponding belief?

Before we go any further, a qualification is in order. I have just suggested that it is difficult to make sense of the idea that a thinker might mistakenly take himself to be supposing that p for the sake of argument when he is really judging that p in his

reasoning. It is important that this claim is limited to the mental act of supposing that p for the sake of argument. It might be held that, in addition to the mental act of supposing that p for the sake of argument, we should also acknowledge the existence of a linguistic act of supposing that p for the sake of argument, in something like the way we acknowledge both the mental act of judging that p and the speech act of asserting that p . One performs this linguistic act when one begins one's utterance by saying, 'Suppose...' There appear to be good reasons for allowing the possibility that a speaker might be mistaken about the character of his act where the linguistic act of supposition is concerned. Whether a speaker has performed the linguistic act of supposing that p for the sake of argument depends on the conventional significance of the words he has uttered, and the speaker might have mistaken beliefs about the conventional significance of his words, or unwittingly fail to utter the words he intends to. For example, if the speaker mistakenly believes that beginning his utterance with the words 'I believe...' is a way of conveying that the force of his utterance is to be taken as that of mere supposition, he might take himself to have supposed that p when he has really made a (qualified) assertion that p . But this possibility does not appear to arise in the case of the mental act of supposing that p for the sake of argument. It seems odd to suggest that I can be mistaken about what I am up to in introducing and treating p as a premise in my reasoning, where the reasoning is in question is a conscious mental activity.

Can the proposal that supposing that p for the sake of argument is merely a matter of introducing and treating p as a premise in one's reasoning without acquiring/manifesting the belief that p explain why it is difficult to make sense of the idea that a subject might mistakenly take himself to be supposing that p for the sake of argument when he is really judging that p in his reasoning? It is hard to see how it

can do so. In general, it seems possible for a subject to manifest the belief that p in ϕ -ing without realizing that he is doing so. Furthermore, it appears to be possible for a subject to take himself *not* to be manifesting the belief that p in ϕ -ing, even though his ϕ -ing is in fact a manifestation of his belief that p . Think again of the subject in Peacocke's example. In her hiring decisions and the recommendations she makes, she manifests the belief that undergraduate degrees from countries other than her own are not of an equal standard to her own, but she is not aware that she is doing so. Presumably, if the subject were challenged over her biased hiring decisions and recommendations, she would deny that they were due to a prejudiced belief she held. This suggests that, in general, whether one is manifesting the belief that p in ϕ -ing is not constitutively dependent on one's higher-order beliefs. This, in turn, suggests that whether one is manifesting the belief that p in introducing and treating p as a premise in one's reasoning does not constitutively depend on whether one believes that one is performing the mental act supposing that p for the sake of argument. If there is no constitutive dependence here, it is difficult to see why one cannot mistakenly take oneself to be merely supposing that p for the sake of argument when one is, in fact, manifesting the belief that p in reasoning on the assumption that p .

A philosophical account of the distinction between judgement and supposition for the sake of argument that merely appeals to the idea that the former, but not the latter, involves the acquisition/manifestation of the belief that p cannot explain why it is hard to make sense of the idea that a subject might mistakenly take himself to be supposing that p when he is really judging that p in his reasoning. What other potential differences are there, then, between judging that p and supposing that p ? And how might these differences help to explain why this kind of mistake appears to be impossible?

4. The Teleological Conception of Cognitive Mental Acts

David Velleman (2000) proposes that belief is distinguished from other cognitive attitudes by its possession of a distinctive aim. Belief aims at truth. When one believes that p , one accepts that p with the aim of accepting that p only if p . One does not have this aim when one accepts that p merely by (e.g.) imagining that p . Shah and Velleman (2005) extend this proposal to cognitive mental events:

[W]hy can't one arbitrarily affirm that p in such a way as to make a judgment? The reason is that an affirmation that p qualifies as a judgment, rather than a mental fiction or hypothesis, only when it is aimed at getting the truth value of p right – aimed, that is, at presenting p as true only if it really is true. (2005: 504)

Shah and Velleman go on to make it clear that they regard the presence of a truth-aim as sufficient for a mental affirmation of p to qualify as a judgement that p . This aim is absent when one merely supposes that p for the sake of argument. When one supposes that p , one's aim is to determine or indicate the consequences of p —perhaps, more specifically, to show that p implies a contradiction, and therefore must be false—not to affirm that p only if p is true.

The proposal that judgement aims at the truth raises many questions, but I want to focus, to begin with, on the question of how it might help us to account for the subject's awareness of the nature of her own cognitive mental acts. Relevant here, I think, is Shah and Velleman's further claim that the aims of the subject's cognitive mental acts, unlike the aims of his cognitive mental states, are always realized

intentionally.¹⁷ Thus, for example, they are committed to the more specific view that when one judges that *p*, one does so with the intention of affirming that *p* only if *p* is true. This claim might be thought to be relevant to explaining why it is hard to make sense of the idea that a subject might mistakenly take himself to be judging that *p* when he is really only supposing that *p* for the sake of argument in his reasoning. It is plausible that agents have privileged access to the intentions with which they act. At least normally, when you ask someone why they are acting in a certain way, where that question is understood as a request to be told what their purpose or intention is in so acting, they are in a position to tell you what their intention is. The awareness we have of the character of our conscious thinking might be thought to be a consequence of this. In consciously affirming a proposition, a thinker is aware of the intention with which he is affirming that proposition, and this is what gives him access to the nature of his own conscious thinking—e.g. to whether he is judging that a certain proposition is true in his reasoning, or is instead merely supposing that it is true for the sake of argument. It is difficult to make sense of the idea (e.g.) that a subject might mistakenly take himself to be supposing that *p* for the sake of argument when he is really judging that *p* in his reasoning because it is difficult to make sense of the idea that a subject might be mistaken about the intention with which he is affirming that *p* in his reasoning. On the face of it, then, a teleological account of the distinctions between the cognitive acts that constitute our conscious thinking is well-positioned to explain why, on the face of it, a subject cannot make certain kinds of mistake about what he is up to in his own conscious thinking.

¹⁷ They presuppose this view in the following footnote: ‘Because judgement is an *act*, it differs from belief in that it necessarily has a literal aim. Hence, there is no problem accounting for judgment’s standard of correctness: its standard of correctness is just the criterion of success associated with the intention with which it is made.’ (Shah and Velleman 2005: 531, note 21)

If this kind of account of our awareness of the nature of our own conscious cognitive acts is to explain why it is difficult to make sense of the idea that a subject might mistakenly take himself to be performing some other type of cognitive act when he judges that p or supposes that p for the sake of argument, it must be the case that when a subject judges that p or supposes that p he affirms that p with a distinctive type of intention. If a subject can affirm that p with the same intention both when he (e.g.) supposes that p and when he affirms that p in some other way, it cannot be his awareness that he is affirming that p with that intention that accounts for his knowledge that he is supposing that p rather than affirming that p in that other way. As I pointed out earlier, in the case of the act of supposing that p for the sake of argument, the most obvious proposal is that the subject affirms that p with the aim of determining the consequences of p . A concern one might have here is that it also appears to be possible for a subject to judge that p in his reasoning with this intention. However, it might be argued here, as Shah and Velleman in effect suggest, that when the subject judges that p in his reasoning he affirms that p with an additional intention that is absent when he supposes that p for the sake of argument. When the subject judges that p in his reasoning he affirms that p with the aim of affirming that p only if p is true, but he does not affirm that p with this intention when he merely supposes that p for the sake of argument.

A concern one might have at this point is that there appears to be another type of cognitive act, distinct from judgement, that also aims at the truth. According to Owens (2003), when one guesses that p one does so with the purpose of guessing that p only if p is true. Whether the guesser necessarily guesses with the intention of guessing truly is a question I will be considering in more detail in the next chapter. For the purposes of the present objection, however, it is sufficient that on at least

some occasions on which a subject guesses that p he does so with the intention of guessing truly. This much seems undeniable. If this is right, then Shah and Velleman are wrong that it *suffices* for an affirmation that p to constitute a judgement that p that the subject affirms that p with the truth-aim. On some occasions on which a subject affirms that p with the aim of affirming that p only if p is true, he merely guesses, and does not judge, that p . Furthermore, it cannot be just in virtue of my awareness that I am affirming that p with the truth-aim that I am aware that I am judging that p ; this aim can also be present when I am merely guessing that p , and not judging that p .

One response to this difficulty would be to appeal to the idea that judging aims, not just at truth, but at knowledge: that is to say, when I judge that p I affirm that p with the intention of affirming that p only if I know that p . Guessing certainly doesn't aim at knowledge: guessing is what I resort to when I don't know the answer to a question, or at least take myself not to know the answer.

In any case, I think that there is a more serious objection to the idea that, when a subject judges that p , he does so with a characteristic kind of intention. Suppose, for the purposes of argument, that the proposal on the table is that when a subject judges that p he affirms that p with the intention of affirming that p only if he knows that p . Is this plausible? Hardly. Typically, a subject who judges that p does not intend to affirm that p only if he knows that p . This would imply that, when one judges that p , the content that p must *already* have occurred to one, upon which one adopts the intention of affirming that p only if one knows that p . But this leaves us with a picture upon which one's conscious thinking is, as it were, pre-empted by a further underlying process of thought. This seems absurd.

It might be replied to this that rather than claiming that the subject forms an intention to affirm a particular proposition only if it is one that he knows each time he makes a conscious judgement, which is implausible, we should instead regard the subject as having a general intention to judge as true only those propositions that he knows. One concern about this proposal is that it appears to result in circularity. How should we specify the content of the subject's general intention? It might be suggested that it is an intention each subject has, for any proposition, p , to affirm that p only if he knows that p . However, this clearly won't do. Subjects often affirm propositions in their conscious thinking that they are fully aware that they don't know, and in doing so they do not violate any general intention that they have. For this proposal to work, therefore, the scope of the general intention must be restricted so that it includes all and only those conscious affirmations that constitute judgements. The obvious way to do this is to specify its content by saying that each subject has an intention, for any proposition, p , to *judge* that p only if p is true. A problem now, however, is that the account is circular: in giving an account of what it is to judge that p , we have had to appeal to an intention whose content includes the target notion of judgement.

Another concern about this proposal is that it is not clear how it can contribute towards an account of a subject's awareness that he is judging that p when he does so. This concern is especially pressing in the present context, given that our original reason for considering the teleological account was that it appears to be well placed to explain why it is hard to make sense of the idea that a subject might make certain kinds of error about what he is up to in his own conscious thinking. It appears to be well placed to explain this because it seems plausible that agents have privileged access to the intentions with which they act. However, if the proponent of the teleological account retreats to the view that there is merely a general intention to

judge as true only those propositions that one knows, then the view is no longer one upon which the subject affirms that *p* with a particular intention when he judges that *p*. Therefore, we can no longer account for the subject's awareness that he is judging that *p* by appealing to his awareness that he is affirming that *p* with a certain intention. This is not to deny that the subject has privileged knowledge of his own intentions, general or otherwise. It is merely to point out that it is unclear how awareness that I have a general intention to judge as true only those propositions that I know can account for my awareness that a particular affirmation of mine is a judgement rather than, say, a mere supposition.

Contrary to our initial impression, then, it is not clear that a teleological account of cognitive mental events can explain why it is difficult to make sense of the idea that a subject might mistakenly take himself to be supposing that *p* for the sake of argument when he is really judging that *p* in his reasoning, or mistakenly take himself to be judging that *p* when he is really only supposing that *p* for the sake of argument. It is implausible that, when I judge that *p*, I do so with the intention of affirming that *p* only if I know that *p*, and it is unclear how a general intention to judge as true only those propositions that I know can contribute towards an explanation of my awareness that a particular affirmation of mine is a judgement, rather than, say, a mere supposition. The suggestion that when I suppose that *p* for the sake of argument I do so with the intention of determining or demonstrating the consequences of *p* is more plausible, but it seems that I might also judge that *p* in my reasoning with this intention. Hence, simply being aware that I am affirming that *p* with this intention cannot be what accounts for my knowledge that I supposing that *p* for the sake of argument, rather than judging that *p*, in my reasoning.

5. Supposition, Judgement and Self-Imposed Constraints

An account of the nature of the conscious mental acts that occur in conscious thinking must be able to explain why it is difficult to make sense of the idea that a subject might be mistaken about what he is up to in his own conscious thinking. In particular, it must have the resources to explain why it is hard to make sense of the idea that a subject might mistakenly take himself to be supposing that p for the sake of argument when he is really judging that p in his reasoning, or mistakenly take himself to be judging that p when he is really only supposing that p for the sake of argument. There is reason to think that the accounts we have considered so far of the difference between judgement and supposition for the sake of argument—the acquisition/manifestation of belief account, and the teleological approach advocated by Shah and Velleman—cannot meet this requirement. Earlier on, I mentioned an account of the difference between judgement and this kind of supposition that has recently been proposed by Soteriou. In this section, I shall outline this account and show how it meets this explanatory requirement.

Soteriou points out that, both when I judge that p in my reasoning and when merely I suppose that p for the sake of argument, I reason on the assumption that p . When I reason on the assumption that p , one of the constraints that is operative in my reasoning is that of treating p as true in my reasoning. I treat p as true in my reasoning by, for example, drawing inferences from the proposition that p , and/or by introducing other propositions as a premises in my reasoning that are not inconsistent with p (unless they are entailed by p).¹⁸ According to Soteriou, when the subject supposes that p for the sake of argument, unlike when she judges that p in her reasoning, the constraint of treating p as true is a self-imposed constraint on her

¹⁸ For this point, see Soteriou (2013: 263). (My wording here follows Soteriou's own.)

reasoning. The subject imposes the constraint on her reasoning by reasoning in recognition of that self-imposed constraint. The subject manifests her recognition that the constraint of treating p as true in her reasoning is self-imposed in the way in which she treats the assumption that p in her reasoning: in particular, she treats the assumption that p as an assumption that is to be discharged by, for example, a conditional judgement or assertion that is outside the scope of the supposition.

How does this proposal help us to understand the subject's awareness of the nature of her own conscious thinking? How, in particular, does it help us to understand why it is difficult to make sense of the idea that a subject might, for example, mistakenly believe himself to be judging that p in his reasoning when he is really supposing that p for the sake of argument? Well, suppose that the subject does believe himself to be supposing that p for the sake of argument in introducing and treating p as a premise in his reasoning. Presumably, this belief will influence the way in which he treats the assumption that p in his reasoning. In particular, he will treat the assumption that p as an assumption that is to be discharged in the manner of a supposition, and in doing so he will manifest his recognition that the constraint of treating p as true is a self-imposed constraint on his reasoning. In doing so, the subject makes it the case that the constraint of treating p as true in his reasoning is a self-imposed constraint: in order to *be* a source of constraints over his own thinking, it is sufficient for a subject to *treat himself* as such. But supposing that p for the sake of argument just is a matter of so treating the assumption that p in one's reasoning. Hence, on Soteriou's account, there is plausibly a constitutive connection between the belief that one is supposing that p for the sake of argument and whether one is supposing that p for the sake of argument in introducing and treating p as a premise in one's reasoning. Soteriou's account is therefore well placed to explain why it hard to make sense of

the idea that a subject might mistakenly believe himself to be performing the mental act of supposing that p for the sake of argument.

A similar point can be made about the mental act of judging that p . Whether one is judging that p in one's reasoning depends on how one treats the assumption that p in one's reasoning. There is a sense in which, when one judges that p in one's reasoning, one regards p as something that really is the case. According to Soteriou, this is to be captured, at least in part, in terms of the idea that one does *not* regard the constraint of treating p as true as a constraint that is self-imposed. (Soteriou makes a further proposal here: that when one judges that p in one's reasoning, one regards the constraint of treating p as true as a constraint that is imposed on one by a fact that one acknowledges to obtain. The idea that we need to invoke the notion of knowledge in giving a philosophical account of the act of judgement is one that I will be considering in detail in the next two chapters.) Presumably, therefore, if the subject believes that he is judging that p in his reasoning, he will not treat the constraint of treating p as a self-imposed constraint on his reasoning, or treat the assumption that p as one that is to be discharged in the manner of a supposition. There is therefore a constitutive connection between the belief that one is judging that p and whether one is judging that p in one's reasoning. This explains why it is difficult to make sense of the idea that a subject might mistakenly take himself to be judging that p in his reasoning when he is really supposing that p for the sake of argument.

Earlier on, I pointed out that a subject who has the capacity to engage in enquiry is plausibly a subject who has the capacity to engage in forms of acceptance besides belief. Such a subject has the capacity for higher-order attitudes—e.g. to recognize that he does not know, or have a belief about, the answer to a certain question. One

might think that the capacity to recognize ‘gaps’ in one’s knowledge, or in one’s beliefs, goes along with the capacity to ‘fill those gaps’ (at least temporarily) by engaging in other forms of acceptance—e.g. by making assumptions for practical purposes, or assumptions for the sake of argument. So far, I have been concerned with marking distinctions between different types of cognitive mental event—specifically, between judgement and supposition for the sake of argument—but there are also issues about how to mark distinctions between cognitive mental states. Believing that p is a mental state, but it is plausible that we should also think of assuming that p for practical purposes as a matter of acquiring a state that plays a certain role in one’s reasoning and action. How, then, should we explain the difference between the state one is in when one believes that p and the state one is in when one is merely assuming that p for practical purposes? This is the issue I want to address in the next section. In doing so, I shall once again be drawing on Soteriou’s idea of self-imposed constraints on one’s reasoning.

6. Believing that p and Assuming that p for Practical Purposes

It is sometimes suggested that we can provide an adequate account of what it is to believe that p in simple motivational terms. To believe that p is to be disposed to behave as if p is true. How are we to understand the notion of a subject being disposed to behave as if p is true? It is well known that we cannot say, for example, that to believe that it is raining is a matter of being disposed to assert that it is raining in response to a query about the weather, being disposed to take one’s umbrella when one leaves the house, and so on; someone who wants to conceal the state of the weather from others, and who enjoys getting wet, can believe that it is raining without being disposed to do either of these things. What one is disposed to do is determined by what one desires as well as by what one believes. However, this does

not preclude a more sophisticated general motivational account of belief that takes into account the role of the subject's desires in fixing her dispositions. For example, it might be held that 'all that's necessary for an attitude to qualify as a belief is that it dispose the subject to behave in ways that would promote the satisfaction of his desires if its content were true' (Velleman 2000: 255).¹⁹ Velleman calls this view the 'purely motivational conception' of belief. The fact that a subject who believes that it is raining and who wants to get wet is not disposed to take their umbrella with them when they leave the house does not refute the purely motivational conception of belief, because in being disposed not to take an umbrella with them they are disposed to behave in a way that is conducive to the satisfaction of their desires, provided that it is actually raining.

In § 2 I considered the proposal that believing that *p* involves being disposed to judge that *p*. One objection to this proposal was that some animals appear to be capable of belief but not judgement. Although this was not a good objection to the proposal that an enquirer believes that *p* only if he is disposed to judge that *p*, it is interesting to consider such believers again in the context of the purely motivational conception of belief. When we are dealing with a believer of this kind—i.e. a believer who is incapable of judgement—it is plausible that the kind of state we report when we say that the subject 'believes' that *p* can be accounted for in purely motivational terms. For a dog, say, to 'believe' that you have thrown the ball in a certain direction is a matter of the dog being in a state that disposes it to behave in ways that would be conducive to the satisfaction of its desires if it were true that you had thrown the ball in that direction. We should not assume, however, that when we

¹⁹ For a classic statement of this view, see Braithwaite 1932–1933. I do not mean to suggest by quoting his statement of view that Velleman endorses this thesis; as we will see in more detail in a moment, he emphatically rejects it.

speak of belief in relation to enquirers we are reporting the same kind of psychological state that we are reporting when we speak of belief in relation to more primitive subjects. An enquirer is capable of engaging in forms of acceptance that are beyond the capacities of non-judging animals. An enquirer can assume for practical purposes that p , but, despite all the ‘anthropomorphic apparatus’ we project onto dogs,²⁰ the idea that dogs are capable of making assumptions of this kind seems absurd. As I shall argue in a moment, the state of assuming that p for practical purposes also plays the motivational role identified above—i.e. that of disposing the subject to behave in ways that would be conducive to the satisfaction of his desires if it were true that p —but when we speak of belief in relation to enquirers, I think that we normally intend to refer to a kind of psychological state that *contrasts* with merely assuming that p for practical purposes. It follows that we cannot give an account of this kind of psychological state, which I referred to earlier as ‘outright belief’, in terms of the purely motivational conception of belief. That conception seems fitted only to the ‘beliefs’ of more primitive subjects.

Velleman (2000) also objects to the purely motivational conception of belief on the basis that there are other attitudes that play the relevant motivational role. However, there are two respects in which my view is weaker than Velleman’s. Firstly, as I have just indicated, I am sympathetic to the view that the purely motivational conception may provide an adequate account of the ‘beliefs’ of more primitive subjects. Secondly, Velleman argues that the motivational role that the purely motivational conception asserts to be distinctive of belief is in fact common to all of the cognitive attitudes. (Since he holds that the cognitive attitudes include supposing that p and imagining that p , it follows that, on his view, a subject who merely

²⁰ The phrase comes from Williams (1973: 138).

imagines that p is disposed to behave in ways that would promote the satisfaction of her desires if p were true.)²¹ I shall not commit to this further claim here. To object to the purely motivational conception on the grounds that it fails to distinguish belief from other propositional attitudes, it is not necessary to make the strong claim that the relevant motivational role is common to *all* of the cognitive attitudes; it is necessary only to identify one propositional attitude, distinct from belief, that motivates the subject in the appropriate way. I will now argue that assuming that p for practical purposes is just such an attitude.

Sometimes one recognizes that one ought to do in a particular practical situation depends on the answer to a question one is uncertain about. One way of dealing with this problem is to try to determine the answer to the question before one acts, but sometimes this is impossible, or at least impractical. In such a situation, it is still possible to make an assumption about the answer to the relevant question. Having done so, one is disposed to behave to ways that would promote the satisfaction of one's desires if the assumption one has made were true. But making the assumption that p in these circumstances does not amount to forming the belief that p ; one is merely assuming that p for practical purposes. An example, due to Bratman, helps to clarify and reinforce the general argument:

*Building Costs*²²

I am planning for a major construction project to begin next month. I need to decide now whether to do the entire project at once or instead to break the project into two parts, to be executed separately. The rationale for the second

²¹ Velleman in fact spends the bulk of his paper arguing that imagining that p is an attitude that disposes the subject to behave in ways that would promote the satisfaction of her desires if p were true. For further critical discussion of this claim, see O'Brien (2005).

²² This name for the case is my own.

strategy is that I am unsure whether I presently have the financial resources to do the whole thing at once. I know that in the case of each sub-contractor—carpenter, plumber, and on on—it is only possible at present to get an estimate of the range of potential costs. In the face of this uncertainty I proceed in a cautious way: In the case of each sub-contractor I take it for granted that the total costs will be at the top of the estimated range. On the basis of these assumptions I determine whether I have at present enough money to do the whole project at once. In contrast, if you offered me a bet on the actual total cost of the project—the winner being the person whose guess is closest to the actual total—I would reason differently. (Bratman 1992: 6)

In this case the agent assumes, in the face of uncertainty, that the total costs for each sub-contractor will be at the top of the estimated range. Having done so, the agent is disposed to behave in ways that would promote the satisfaction of his desires if it were true that the total costs for each sub-contractor will be at the top of the estimated range. But the agent does not believe that the total costs for each sub-contractor will be at the top of the estimated range; he is merely assuming that this is so for practical purposes. Hence, a mental state with the content that p may fail to be the belief that p even if it plays the motivational role identified by the purely motivational conception of belief. It follows that the condition identified by that conception is at most necessary, but not sufficient, for an attitude to qualify as a belief.

Is there any way for the proponent of the purely motivational conception to respond to this objection? He has two options here: he can argue either that, contrary to appearances, the agent of building costs is not a mental state, with the content that

the total costs for each sub-contractor will be at the top of the estimated range, that disposes him to act in ways that would be conducive to the satisfaction of his desires if its content were true, or else he can acknowledge that the agent is in such a state, but argue that, contrary to appearances, the state in question is one of belief. The first option seems to me a non-starter. Once the agent has assumed that the building costs for each sub-contractor will be at the top of the estimated range, the assumption guides the agent's conduct in a way that would promote the satisfaction of his desires if its content were to be true. For example, if the agent calculates that if the cost for each sub-contractor is at the top of estimated range then he will not be able to afford to do the whole project in one go, he will break the project into two parts. This course of action is the one that would promote the satisfaction of his desires were it to be true that the cost for each sub-contractor will be at the top of the estimated range.

What of the second option of arguing that the agent of building costs does actually believe that the costs for each sub-contractor will be at the top of estimated range? In support of this contention, it might be pointed out that it is not wholly unnatural to say that, in the face of uncertainty about the costs of the project, the agent decides to *believe* that the costs for each sub-contractor will be at the top of estimated range. Why shouldn't we take this belief-ascription at face value? One problem with doing so is that it seems to lead to the conclusion that the agent of building costs, on a certain way developing the case, has beliefs that violate rationality constraints on belief. It is conceivable that the agent actually believes that the costs for each sub-contractor are likely to be in the middle of the estimated range. After all, he assumes that they will be at the top of the estimated range not because he thinks this is the most likely outcome, but because he wants to guard against the possibility of trying

to do the whole project in one go if it is beyond his means to do so. In relation to this version of the case, the objector is apparently committed to the claim that the agent believes both that the costs for each sub-contractor are likely to be in the middle of the estimated range, and that the costs for each sub-contractor will be at the top of the estimated range. Surely the agent shouldn't hold both of these beliefs at once. On the face of it, however, it isn't appropriate to criticize the attitudes of the agent of building costs on this basis. (Perhaps he is being overly cautious or paranoid in assuming that the costs for each sub-contractor will be at the top of the estimated range, but then the criticism is not that his *beliefs* are in tension with one another.) We should therefore reject the claim that the agent believes that the costs for each sub-contractor will be at the top of the estimated range. Rather, this is something that he is merely assuming for practical purposes.

Such examples should lead us to reject the idea that the purely motivational conception of belief succeeds in explaining how believing that p differs from other stative propositional attitudes. An agent who is merely assuming that p for practical purposes is also in a state that plays the relevant motivational role. How, then, does believing that p differ from the state one acquires when one assumes that p for practical purposes?

An enquirer who has the outright belief that p is disposed to judge that p . Merely saying this fails to identify a respect in which believing that p differs from the state one is in when one is assuming that p for practical purposes. Someone who is assuming that p for practical purposes is disposed to plan on the assumption that p , and when such a subject plans on the assumption that p because he is assuming that p for practical purposes, it may legitimately be *said* that he 'judges' that p in his

reasoning. (In ordinary language, we do not appear to reserve a special term to report the way in which a subject affirms that p when he plans on the assumption that p because he is assuming that p for practical purposes.) However, the notion of judgement in play in the present discussion is a refined philosophical one. Judging in this sense contrasts with merely supposing that p for the sake of argument. I argued earlier, following Soteriou, that part of the difference between such judgement and supposition for the sake of argument is that when the subject judges that p in his reasoning the constraint of treating p as true is not a self-imposed constraint on his reasoning. What I now want to suggest, again following Soteriou (2013: 351), is that when a subject plans on the assumption that p because she is assuming that p for practical purposes, the constraint of treating p as true is likewise a constraint on his reasoning that is self-imposed. Thus, the fact that believing that p involves a disposition to judge that p , *in this sense*, does identify a distinctive feature of belief vis-à-vis the attitude of assuming that p for practical purposes.

We should not advocate the claim that when a subject plans on the assumption that p because he is assuming that p for practical purposes, the constraint of treating p as true is self-imposed, on the basis of the premise that when a subject plans on the assumption that p because he is assuming that p for practical purposes, he merely supposes that p for the sake of argument. When you plan on the assumption that p because you are assuming that p for practical purposes, you do not merely suppose that p for the sake of argument in your planning.²³ You can suppose that p for the sake of argument in your planning, but when you do so you will be prepared to make at most conditional decisions on the basis of the assumption that p —i.e. decisions of the form: if p , then I will ϕ . By contrast, when you plan on the assumption that p

²³ This point is noted by Bratman (1992).

because you are assuming that p for practical purposes, you are prepared to make unconditional decisions on the basis of the assumption that p . So in accepting the claim that when a subject reasons (e.g. plans) on the assumption that p because she is assuming that p for practical purposes, the constraint of treating p as true is self-imposed, we commit ourselves to the idea that there are further ways of introducing and treating p as a premise in one's reasoning, besides supposition, that involve self-imposed constraints on one's reasoning.

In the case of supposition, the subject manifests her recognition that the constraint of treating p as true is self-imposed by treating the assumption that p as an assumption that is to be discharged by (e.g.) a conditional judgement or assertion. But when a subject reasons on the assumption that p because she is assuming that p for practical purposes, she does not regard the assumption that p as an assumption that is to be discharged in this way. How, then, does she manifest her recognition that the constraint of treating p as true is self-imposed? A natural proposal is that she manifests her recognition that the constraint is self-imposed by treating the assumption that p as one that she is making at least partly for practical reasons, relative to some project that she has. For example, she may recognize that she is assuming that p in a certain context because of an asymmetry in the cost of errors; if she were to plan on p when $not-p$, the consequences would be manageable, but were she to plan on $not-p$ when p , the consequences would be disastrous. By contrast, when the subject judges that p in her reasoning, she does not regard the assumption that p as one that she is making at least partly for practical reasons.

To summarize, my aim in this section has been to explain how believing that p differs from merely assuming that p for practical purposes. I began by considering

the proposal that belief is marked out from other forms of acceptance by its distinctive motivational role: a subject who believes that p is disposed to act in ways that would be conducive to the satisfaction of his desires if p were true. Although the purely motivational conception of belief may be adequate as an account of the 'beliefs' of primitive subjects who aren't capable of engaging in sophisticated forms of conscious thinking and alternative forms of acceptance, it is inadequate as an account of the beliefs of enquirers. When we speak of belief in relation to enquirers, I think that we normally intend to refer to a kind of psychological state, which I earlier referred to as 'outright belief', that contrasts with merely assuming that p for practical purposes. However, as I argued, the state of assuming that p for practical purposes also plays the motivational role that the purely motivational conception asserts to be distinctive of belief. In order to mark what is distinctive of outright belief, we can appeal to the idea that a subject who has the outright belief that p is disposed to consciously judge that p . However, more needs to be said here, because it may also be said that a subject who is merely assuming that p for practical purposes is disposed to 'judge' that p in his planning. I appealed at this point to ideas I introduced earlier on when I was discussing the difference between judging that p in one's reasoning and merely supposing that p for the sake of argument. When one supposes that p for the sake of argument the constraint of treating p as true in one's reasoning is self-imposed, and one regards it as such. This is not the case when one judges that p in one's reasoning. I suggested that, when a subject plans on the assumption that p because he is assuming that p for practical purposes, he likewise regards the constraint of treating p as true as self-imposed, although we should be careful to note that this is not because he is supposing that p for the sake of argument. Thus, although a subject who is assuming that p for practical purposes

may be said to have ‘judged’ that p in his planning, he has not judged that p in the sense that involves *not* treating the constraint of treating p as true as a self-imposed constraint on his reasoning.

7. Summary

My focus in this chapter has been on the question of the conditions under which an enquirer believes that p . I have advanced three main claims. The first is that, when an enquirer has the outright (as opposed to the repressed) belief that p , he is disposed to judge that p . We should understand the problematic reasoning to concern outright belief, since the claim that truly believing that p is just as good as knowing that p is plausible only in relation to outright belief. Secondly, when it comes to giving an account of what is distinctive of judging that p , part of what we should be appealing to is the idea that, when a subject judges that p in his reasoning, the constraint of treating p as true in his reasoning is not self-imposed, and neither does the subject regard it as such. Finally, at least part of what is distinctive of believing that p , as opposed to merely assuming that p for practical purposes, is that it involves a disposition to judge that p in this sense.

In the next chapter, I will be investigating the nature of another kind of mental act which, like supposition for the sake of argument, is distinct from judgement. The act I have in mind is that of guessing. What is the difference between guessing that p and consciously judging that p ? Is it right to think that, when a subject guesses that p , he necessarily does so with a characteristic type of purpose or intention—e.g. that of guessing that p only if p is true? These are some of the issues, amongst others, that I will be investigating. I will ultimately be arguing, in Chapter 4, that considerations about the difference between guessing and judging can be used to motivate the claim

that, when a subject judges that p , he takes himself to know that p . Given the dispositional link between outright belief and judgement, this claim is relevant to how we ought to respond to the problematic reasoning.

Chapter 3

Guessing

Although philosophers often use the term ‘guess’ and its inflections in their writings, they rarely pay explicit attention to the question of what guessing is. Perhaps this is because it is assumed that guessing is not a particularly interesting propositional attitude in comparison to attitudes such as knowing that p or believing that p . I think this assumption is mistaken. Closer inspection of guessing reveals it to be a surprisingly complex phenomenon that can only be understood through sustained philosophical reflection: in particular, it is not entirely straightforward to provide an adequate account of the conditions under which a subject is guessing that p in answering that p to some question. A consequence of this is that the claims philosophers *have* made about guessing often turn out to contain significant errors when they are subjected to careful scrutiny. My main aim in this chapter is to correct these errors by providing a more nuanced account of the conditions under which a subject is guessing that p .

In § 1, I delineate the kind of guessing that is my concern by distinguishing some different ways in which the term ‘guess’ is ordinarily used. § 2 outlines the account of guessing proposed by Roy Sorensen (1984). In § 3 I argue that it is possible for a person’s answer to a question to satisfy the conditions identified by Sorensen without being a guess about the answer to that question, because an answer to a question can satisfy Sorensen’s conditions without being subject to the standard of success constitutive of guessing. I go on to consider (and reject) the proposal that guessing necessarily aims at the truth, and that this is why truth is the constitutive standard of success for guessing. § 4 discusses the relationship between guessing,

believing and knowing. I argue, contrary to Sorensen, that guessing that p is not incompatible with knowing that p or believing that p , and that we should instead conceive of the relationship between guessing and knowing at least partly in higher-order terms.

1. Uses of ‘Guess’

The term ‘guess’ can be used in a variety of ways. Sometimes, for instance, we characterize a subject’s belief that p as a guess. In philosophy this use of the term ‘guess’ is frequently invoked in discussions of the analysis of propositional knowledge. The following passage from the *Stanford Encyclopedia of Philosophy*, which I quoted in Chapter 1, provides an illustration:

Suppose that William flips a coin, and confidently believes on no particular basis that it will land tails. If by chance the coin does land tails, then William’s belief was true; but a lucky guess such as this one is no knowledge. (Ichikawa and Steup 2013: Section 1.3)

As the passage makes clear, by describing a subject’s belief that p as a guess one conveys that the subject’s belief, if true, is true by luck, and therefore does not constitute knowledge. It should also be noted that we can use the term ‘guess’ to report doxastic states of mind aside from the outright belief that p . For example, the sentence ‘Tim’s guess is that it will rain’ is naturally understood to report the fact that Tim estimates that it will rain, and ‘Tim estimates that it will rain’ can be used to report Tim’s belief that rain is more likely than not.

My concern in this paper is not with the varieties of doxastic state that are reported or characterized by such uses of ‘guess’, for the kind of guessing that interests me is

an event and not a state. Relevant here is the use of ‘guess’ as a success verb. ‘Guess’ and its inflections are naturally understood to function as success verbs in sentences like the following:

- Olly guessed that the defendant was guilty;
- Molly guessed who was at the door;
- Maggie managed to guess where the key was hidden.

In such sentences ‘guess’ is naturally understood to report an event that occurred at a certain time, rather than some state that obtained over an interval of time. We can think of that event as the event of a subject answering a question, at least in her own mind. In the first example above, for instance, we can think of the event reported as the event of Olly answering the question of whether the defendant was guilty. Furthermore, when ‘guess’ is used as a success verb, it not only implies that the subject answered the relevant question, but also that she answered the relevant question successfully. What does success consist in here? It is not sufficient for guessing the answer to a question that one merely manages to give *an* answer to the question. Rather, one has to give *the* answer—i.e. the true answer, or at least *a* true answer, if the question has more than one true answer. Notice, however, that the understanding of ‘guess’ as a success verb in these sentences is not compulsory; I can say, without contradiction, ‘Olly guessed that the defendant was guilty, but actually he was innocent’.²⁴ More generally, we can use the term ‘guess’ to report an event of a subject answering a question that will count as a success just in case the

²⁴ Admittedly, it is more difficult to conceive of circumstances in which one could say, without contradiction, ‘Maggie managed to guess where the key was hidden, but she was mistaken’. The reason is that ‘managed’ suggests that there was at least a chance that Maggie would fail; and whilst it is readily intelligible that she might have failed to successfully guess where the key was hidden, it is not so readily intelligible that she might have failed to make any guess at all.

answer she gave is true, but also indicate that the subject was unsuccessful, or at least leave it open whether her attempt to answer the question truly was successful. This is the kind of guessing that I shall be concerned with in this chapter.

It is not sufficient for a subject to guess that p , in this sense, that she answers p to a question in conditions in which her answer will count as a success if and only if it is true. If you enquire after my father's middle name and I tell you that it is 'Evan', then the relevant criterion of success applies to my answer, but I am clearly not guessing that my father's middle name is 'Evan'. One question that needs to be addressed, then, is what further conditions need to obtain in order for a subject's answer to constitute a guess. We should also try to say something about the conditions under which a subject's answers to questions are regulated by the standard of success constitutive of guessing. In the next section I lay the groundwork for approaching these issues by outlining the account of guessing proposed by Sorensen (1984).

2. Sorensen's Account of Guessing

Strictly speaking, Sorensen does not attempt to specify necessary and sufficient conditions for any subject to guess that P , but rather necessary and sufficient conditions for a *rational* subject to make what Sorensen calls a *unique alternative guess*. A guess is a unique alternative guess 'just in case the guess is a direct answer to a question which the guesser believes has exactly one correct answer' (1984: 80). (In discussing the sense in which making a guess involves answering a question, Sorensen claims that the question needn't be 'explicitly posed' to the guesser, but 'need only express his puzzlement' (ibid.)) Since someone who guesses the answer to a question might (correctly) believe that the question has more than one correct

answer, not all guesses are unique alternative guesses. On Sorensen's view, a rational subject S makes a unique alternative guess that p_1 iff:

- (1) There is no answer to the question, p_2 , such that S believes that p_2 is more likely than p_1 .
- (2) For any answer to the question, p_3 , if S picks p_3 , then $p_3 = p_1$.
- (3) S does not believe that p_1 . (1984: 83)

Sorensen refers to (1) as the nondomination condition for guessing, (2) as the uniqueness condition for guessing, and (3) as the condition that believing is incompatible with guessing (BIG). As seems reasonable, I shall understand Sorensen's account as an attempt to specify necessary and sufficient conditions for the kind of guessing I am interested in.

I wish to make two comments about Sorensen's third condition, BIG. The first is that BIG is intuitively plausible. If I have a sincerely held false belief that the Battle of Hastings was fought in 1266—perhaps one I picked up as a result of a typographical error in a history book—then intuitively I am not guessing that the Battle of Hastings was fought in 1266 when I answer '1266' to the question 'When was the Battle of Hastings?'. The second is that Sorensen observes that the incompatibility of guessing that p with knowing that p (KIG) follows from the conjunction of BIG with the popular thesis that one knows that p only if one believes that p . This consequence speaks in favour of KIG, in Sorensen's view, because one of the most striking features of our ordinary thinking about guessing is the contrast we recognize between guessing and knowing. A natural response to a contestant getting an esoteric question right is to ask, 'Did she know that that was the answer, or was she only guessing?' Relatedly, Sorensen argues that we should accept KIG

because it ‘explains why it is so natural to infer ‘John did not know that p ’ from ‘John guessed that p ’ (p. 82). Such mundane facts make it plausible that there is an important relationship between guessing that p and knowing that p , and also that this relationship can be captured in terms of KIG, i.e., the simple claim that S guesses that p , at some time t , only if S does not know that p at t . Although Sorensen does not include KIG in his account of guessing, since it would be redundant, he both endorses KIG and takes his account to be supported by the fact that it entails KIG.

As I said earlier, Sorensen explicitly states that his claims about guessing ‘should be understood as claims limited to unique alternative guesses’ (1984: 80). Since, however, none of the objections I wish to make to his account can be avoided by taking note of this limitation, from now on I will simplify matters by ignoring it. Furthermore, Sorensen limits his analysis to *rational* subjects, by which he appears to mean subjects who always act rationally. By contrast, I wish to arrive at an account of the conditions under which *any* subject guesses that p , so I will also usually ignore Sorensen’s limitation of his account to rational subjects.

3. Guessing and its Standard of Success

It is, I think, fairly straightforward to come up circumstances in which a subject would not be guessing that p in answering that p to some question, even though he satisfies Sorensen’s conditions for doing so. I shall begin this section by stating a counterexample of this kind—i.e. a counterexample to the *sufficiency* of Sorensen’s conditions for guessing. I shall then consider how we should respond to this difficulty.

Imagine that you are taking part in a multiple-choice test with the following twist: you are rewarded for giving false answers, and punished for giving true ones. You are asked to identify the capital of Nigeria out of the following four options: Accra, Abuja, Alofi and Apia. Let us suppose that your knowledge of African capitals is poor, so that you are forced to pick one of the possible answers at random. Under such conditions, in picking the answer 'Abuja' you clearly satisfy the conditions Sorensen lays down as necessary and sufficient for guessing that Abuja is the capital of Nigeria: in particular, (1) there is no answer to the question that you believe to be more likely than Abuja, (2) there is no other answer you pick, and (3) you do not believe that Abuja is the capital of Abuja. The problem is that it seems very counterintuitive to describe you as guessing that Abuja is the capital of Nigeria in the envisaged circumstances. If this is right, then satisfying the conditions identified by Sorensen is not sufficient for you to guess that p in answering that p .

When it comes to explaining why you weren't guessing that Abuja is the capital of Nigeria, the obvious thing to say is that you weren't guessing because your answer wasn't subject to the standard of success constitutive of guessing. In the envisaged scenario you are rewarded for giving false answers, not true ones; thus, the standard of success for your answers is falsity, not truth.²⁵ It is not sufficient for a person's answer to a question to be subject to the standard of success constitutive of guessing that the person satisfies the conditions identified by Sorensen in answering as she does; in particular, a person's answer to a question can fail to be subject to this standard of correctness even if there is no answer to the question that he believes to be more likely than the one that he gives.

²⁵ Notice that, considered as an answer to the question 'Which of these cities is *not* the capital of Nigeria?', your answer *does* count as a success just in case it is true that Abuja is not the capital of Nigeria. This is why it is natural to say that, in the envisaged circumstances, you are guessing that Abuja is not the capital of Nigeria.

The proposal that truth is the standard of success for guessing is related to the proposal that guessing aims at the truth. As Shah and Velleman point out, one's goals or purposes 'establish criteria of success for the activities they regulate' (2005: 498). As I mentioned in the previous chapter, the view that it is constitutive of guessing that one guesses with the aim or purpose of guessing truly is advocated by David Owens (2003). Owens provides the following general account of what it is for a propositional attitude to aim at the truth:

Φ -ing that p aims at the truth if and only if someone who Φ s that p does so with the purpose of Φ -ing that p only if p is true. (2003: 289)

Having introduced this general account of what it is for a propositional attitude to aim at the truth, Owens goes on to claim that '[g]uessing aims at the truth in the sense just defined' (2003: 290). So, on Owens's view, S guesses that p only if S does so with the purpose of guessing that p only if p is true. The proponent of the view that guessing aims at the truth can now explain why truth is the standard of success for guessing by appealing to the idea that the guesser necessarily guesses with the purpose of guessing truly. The criterion of success for guessing is established by the guesser's own aim.

Attractive as this idea seems, it cannot be quite right. The problem is that there appear to be cases in which a subject's answers to questions are regulated by the relevant standard of success, and thus may constitute guesses about the answers to those questions, despite the fact that the subject is indifferent to whether she guesses successfully.²⁶ Imagine, for instance, that you ask me to guess the winning numbers

²⁶ Thanks to Quassim Cassam for drawing my attention to cases of this kind. The example that follows is his own.

in last night's national lottery draw. Neither of us bought a ticket, and I don't stand to gain anything by guessing correctly; it's just a game. Since I didn't see the draw, I don't know what the winning numbers were. After several unsuccessful attempts to guess the winning numbers in last night's draw, you ask me to guess the winning numbers in last week's draw. By now I'm getting tired of playing, but you insist that I carry on making guesses, and I do so. Suppose that the next answer I give is that the winning numbers in last week's draw were 2, 9, 18, 27, 33, and 46. It is intuitively plausible that in giving this answer I am guessing that the winning numbers in last week's draw were 2, 9, 18, 27, 33, and 46. But do I guess with the purpose of guessing that the winning numbers were 2, 9, 18, 27, 33, and 46 only if those were, in fact, the winning numbers? It seems not. I don't guess with the purpose of guessing truly if I don't care whether I guess truly or not, and in the envisaged circumstances I don't care; all I want to do is keep you happy by continuing to play your tedious game. If this is right then it's false that guessing aims at the truth in the sense defined by Owens: someone who guesses that p does not necessarily do so with the purpose of guessing that p only if p is true. But if we cannot appeal to the purpose of the guesser in accounting for guessing's standard of success in such cases of 'indifferent guessing', what are we to say about them?

The obvious thing to say is that what cases of indifferent guessing draw our attention to is the fact that a person's answers to questions may be subject to the standard of success constitutive of guessing in virtue of features of the situation besides the purpose with which he answers. It is plausible in example above, for instance, that my answers are subject to the standard of success constitutive of guessing because I am participating in a guessing game. One of the rules of the guessing game you and I tacitly acknowledge is that I win the game by giving the true answer to the question

asked. Winning the game is a form of success, so the criterion for winning the game establishes a standard of success for my answers. My answers are subject to this standard of success not because I answer with the purpose of answering truly, then, but because I knowingly participate in a game in which success consists in giving the true answer to the question asked. This is not to say, of course, that *whenever* a subject makes a guess about the answer to a question her answer is regulated by the standard of success constitutive of guessing in virtue of the fact that she is a witting participant in a guessing game. One can take this line about this kind of case of indifferent guessing consistently with holding that, in cases in which the subject is *not* indifferent, her answers to questions may be regulated by guessing's standard of success in virtue of the fact that she answers with the purpose of answering truly.

One concern about an account of guessing that gives up on the idea that the guesser necessarily guesses with the purpose of guessing truly is whether it can adequately account for certain doxastic constraints on guessing. It is intuitive that S cannot be guessing that *p* if S knows or believes that *p* is false.²⁷ For example, since I know that 2 is the only even prime number, I cannot guess that 4 is the only even prime. It is not difficult to explain why guessing is subject to this doxastic constraint if we accept the view that the guesser necessarily guesses with the purpose of guessing truly. One cannot do something with a certain purpose if one believes that one cannot achieve that purpose by doing that thing.²⁸ So the purpose that ensures that

²⁷ For reasons given in § 4, this claim cannot be strictly correct. However, it is false for reasons that needn't detain us here.

²⁸ Admittedly, I may open the curtains with the purpose of cleaning the windows despite the fact that I know that I cannot clean the windows simply by opening the curtains. But, in Michael Bratman's terms, this is a case in which I open the curtains with the *further* purpose or intention of cleaning the windows (see Bratman 1987: 128–130). I see my action of opening the curtains as part of a plan that includes a distinct subsequent action of cleaning the windows. Thus, I do not open the curtains with the purpose of *thereby* cleaning the windows. But when I answer a question with the purpose of answering the question truly, I do answer with the purpose of thereby answering truly; I don't merely see my act of answering as part of a plan that includes a subsequent action of answering the question

the subject's answer is regulated by guessing's standard of success also ensures that the doxastic constraint applies. But what about cases in which a subject's answer is regulated by the relevant standard of success not in virtue of the purpose with which she answers, but rather in virtue of the fact that she is participating in a guessing game? Here we cannot explain why her answers are subject to the doxastic constraint by appealing to the purpose with which she answers. It also seems clear the relevant doxastic constraint still applies: no matter how indifferent I am about whether I guess correctly, I am *not* guessing that the winning numbers were 2, 9, 18, 27, 33, and 46 if I know that these were not the winning numbers.

It might be argued that I would not be guessing in naming these numbers because I cannot continue to participate in a game consistently with deliberately trying to lose it. Since I am no longer participating in the game, and do not myself have the purpose of answering truly, my answers are no longer regulated by the standard of success constitutive of guessing, and thus do not constitute guesses. However, it's not clear that deliberately losing a game is always inconsistent with playing it: if it were, there would be no such thing as letting your opponent win, as opposed to terminating the game but pretending to play on in order to trick your opponent into thinking that he had won.

One way of explaining why one's guesses are subject to the doxastic constraint even in cases of indifferent guessing is to acknowledge that cases in which the subject answers a question with the purpose of answering truly enjoy a primary status in relation to our concept of guessing.²⁹ The thought here is this. In central or paradigm cases, a subject is guessing that p only if the subject answers that p with the purpose

truly. And, in general, it seems that I cannot do something with the purpose of *thereby* ϕ -ing if I believe that I cannot ϕ by doing that thing.

²⁹ My thinking here is indebted to Steglich-Petersen (2006).

of answering truly. In more general terms, our concept of guessing is most basically the concept of a 'goal-oriented action-type' (Steglich-Petersen 2006: 511). This purpose or goal generates a standard of success for guessing, and it is also the source of the doxastic constraint that a subject cannot be guessing that p if he know or believes p to be false. But a subject may satisfy the doxastic constraint, and his answer be subject to the relevant standard of success, despite the fact that he does not answer with the purpose of answering truly. If he satisfies these conditions, and any others required for guessing aside from having the aforementioned purpose, then there's a derivative manner in which he counts as guessing that p despite the fact that he lacks the purpose present in paradigm cases of guessing. To advocate this line of thought is to acknowledge that there's something importantly right about Owens's proposal that guessing aims at the truth: one cannot understand our concept of guessing without understanding that guessing paradigmatically involves answering a question with the purpose of answering truly. However, we are prepared to apply the concept in cases where the relevant goal or purpose is absent provided that the subject meets other conditions on performing an action of the relevant type. Thus, it would be a mistake to regard satisfying the truth-aim condition, as construed by Owens, as strictly necessary for guessing that p .

Does Sorensen's proposal that a person cannot be guessing that p in answering that p if he believes that some alternative answer to the question is more likely than p express another genuine doxastic constraint on guessing that p ? I think we have good reasons for rejecting the claim that guesses must conform to this nondomination condition. Suppose, for example, that you ask me to guess whether a tossed coin will land heads or tails. I know, and thus believe, that the coin is biased towards landing heads, but I also have a hunch that this is one of those occasions on which events

will defy the balance of probabilities, so I answer that the coin will land tails. Even supposing that I'm not so naïve as to think that my hunch is, or implies the existence of, evidence that increases the likelihood that the coin will land tails, it seems difficult to deny that in answering 'tails' I am guessing that the coin will land tails. After all, I do not know or believe that the coin will land tails, and my answer will be considered successful just in case the coin does in fact land tails. But if this right then satisfying the nondomination condition is not even a necessary condition for guessing that p , because in the envisaged circumstances I guess that the coin will land tails despite the fact that I believe it is more likely to land heads.³⁰

4. Guessing, Believing and Knowing

In the previous chapter I noted that the term 'judgement' is often used to refer to a type of occurrent thought that is closely associated with belief. You make a judgement in this sense when, for example, you work out that the product of 13 and 17 is 221, or at least take yourself to have done so. Although the conditions under which a judging occurs are far from obvious, an apparently secure starting point in investigating what these conditions are is that judging that p contrasts with other ways of consciously affirming that p , such as supposing that p for the sake of argument or guessing that p . When one supposes that p for the sake of argument, one does not judge that p in one's reasoning, and when one judges that p in one's reasoning, one does not suppose that p for the sake of argument; likewise, when one judges that p , one does not guess that p , and when one guesses that p , one does not judge that p . It is tempting to try to provide a deeper philosophical explanation of why these different ways of consciously affirming a proposition are incompatible

³⁰ Note that this case may not be counterexample to Sorensen's analysis taken on its own terms, since in this case it might be maintained that my guess is not rational, so that I don't qualify as a rational subject in Sorensen's sense.

with one another. Such an explanation will take the following form: it will identify some condition, C , that is satisfied when one judges that p , but which is not satisfied when one merely (e.g.) guesses that p .

Sorensen's proposal that guessing that p is straightforwardly incompatible with both knowing that p and believing that p is relevant here. We saw in the previous chapter that the view that judging that p essentially involves the acquisition/manifestation of the belief that p , though not obviously correct, is at least a defensible proposal. This view implies that if S judges that p , then S believes that p . If, as Sorensen claims, guessing that p is incompatible with believing that p , then we have an explanation of why it is impossible for a conscious affirmation that p to simultaneously constitute both a guessing and a judging that p : on pain of contradiction, a subject cannot both believe that p and not believe that p when she affirms that p .

It is natural to think that guessing that p is straightforwardly incompatible with both knowing that p and believing that p . This led Sorensen to include BIG in his account of guessing. However, I believe that the existence of counterexamples to both BIG and KIG is a consequence of a mundane fact about knowledge, belief and propositional memory: namely, that it is possible for a subject to know that p at t , and/or to believe that p at t , despite the fact that she is unable to recall that p at t . I will focus below on showing how this possibility can be used to refute KIG.

One illustration of the general possibility of knowing p despite being unable to recall that p is provided by cases in which the subject is in a state that precludes her recalling that *anything* is the case, consistently with states of knowing continuing to obtain in her—for example, when she is in a state of dreamless sleep, or comatose. But we also ordinarily acknowledge the possibility of a subject being unable to recall

that p despite the fact that she knows that p and is awake and attending to a question to which p is the answer. Indeed, it is a common experience to be aware of such a situation from the first-person perspective: to know that one has learnt the answer to some question, and that the information is likely still ‘in there somewhere’, despite the fact that (frustratingly, embarrassingly) one is unable to bring the answer to mind. Suppose that you are in the situation described: you know that p , despite the fact that you are awake and unable to recall that p . Someone puts to you a question to which p is the answer. You cannot remember the answer. In the end, you pick one of the possible answers at random, and say ‘ p ’. My claim is that the envisaged scenario is one in which you have guessed that p , despite the fact that you knew that p . If this is right, then KIG is false. I think that the counterexample stands even if we allow that at the time at which you answered that p you knew yourself to know the answer to the question, despite your present inability to recall it. Furthermore, since the scenario does not seem to be one in which you are necessarily guilty of any form of irrationality—intuitively, failures of memory are not failures of rationality—it also apparently refutes Sorensen’s version of KIG, upon which it is restricted to rational subjects.

The counterexample rests on the claim that the envisaged circumstances are ones in which (a) the subject guesses that p and (b) the subject knows that p . Someone who wishes to defend KIG therefore has to deny at least one of (a) and (b). Is it at all plausible to deny either of these claims?

It seems to me difficult to deny (a). There is, however, a point that might cause some confusion about the matter. When a subject knows that she knows the answer to a question, despite her inability to recall it, she will often not be able to make a guess

about the answer to the question. This is because all the answers that do come to mind will be ones she recognizes to be incorrect, and, as we saw in § 3, she cannot give an answer that she recognizes to be incorrect if her answer is to constitute a guess. However, it certainly doesn't follow from this that a subject can *never* guess that p under these circumstances. Sometimes, even when one knows oneself to know the answer, one may be unsure whether the answer that comes to mind is correct, and it seems to me that in such circumstances one may still be guessing in giving that answer. Moreover, it is inessential to the counterexample that the subject should realize that she knows the answer to the question. When she does not realize this, it is less likely that she will recognize answers that come to mind to be incorrect.

What about the option of denying (b), that the envisaged circumstances are ones in which the subject genuinely knows that p ? Presumably, someone who wishes to deny (b) will be motivated by the idea that, if a subject is unable to recall that p despite the fact that she is attending to a question to which p is the answer, then she doesn't know that p . This idea clearly contradicts our commonsensical understanding of the psychology of recall. Part of that common sense understanding is the idea that a subject might be unable to recall that p at a certain time because she is in a state that impairs her ability to access her stored knowledge that p . For example, she might be unable to recall that p , even though she is awake and attending to a question to which p is the answer, because she is tired, or drunk. She knows that p , but she is temporarily unable to retrieve her knowledge. Denying this aspect of common sense psychology has unattractive consequences. Suppose that our subject, who was previously drunk and unable to recall that p , has now sobered up. Someone puts to her a question to which p is the answer, and, recalling that p is the case ('Now I remember!'), she confidently asserts that p . How are we to account for the fact that

the subject is now able to recall that p ? If we think that when she was drunk and unable to recall that p , she didn't know that p , we will have to appeal to the idea that at some time since she was drunk, she has acquired the knowledge that p . But we do not normally think of sobering up as a way of acquiring knowledge. Furthermore, it may well be the case that she was not in a position to learn that p in any other way between being drunk and sobering up. The option of denying that the envisaged circumstances are ones in which the subject knows that p therefore appears highly implausible.

Are the circumstances described above also ones in which the subject may be guessing that p , despite the fact that she believes that p ? This follows straightforwardly if we assume that knowing p entails believing p . But notice that we do not need to rely on anything as strong as the entailment thesis here: all we need is the more minimal idea that it is at least consistent with a subject knowing that p , despite the fact that she is unable to recall that p , that she also believes that p .

KIG and BIG fail to provide a correct account of the relationship between guessing, believing, and knowing. Consequently, we cannot appeal to BIG to explain why guessing that p is incompatible with judging that p . The idea that judging involves the acquisition/manifestation of belief suggests a simple fix for this problem, however: guessing is not incompatible with believing/knowing *per se*, but only with the acquisition/manifestation of belief/knowledge. This condition does not deliver the incorrect verdict that the subject isn't guessing in cases of impaired recall. When a subject knows that p but is unable to recall that p he may still guess that p , because he will not be manifesting his knowledge that p in affirming that p ; likewise, when a subject is not able to access in memory a stored false belief that p , he may still guess

that p because he will not be manifesting his belief that p in affirming that p . The modified version of BIG can be used to explain why guessing that p is incompatible with judging that p : on pain of contradiction, a subject cannot simultaneously form/manifest the belief that p in affirming that p and not form/manifest the belief that p in affirming that p .

As I noted in the previous chapter, the claim that judging that p essentially involves the acquisition/manifestation of the belief that p is not obviously correct. Some authors have suggested that it is possible for a subject who does not believe that p to judge that p , and thus to judge that p without forming/manifesting the belief that p . The examples cited to support this claim are hardly conclusive, however. A better objection, I think, is that it is not clear that the acquisition/manifestation account can accommodate the typical incompatibility of guessing that p with believing that p without collapsing into a higher-order account. The following case can be used to raise the concern I have here. Suppose we're listening to a politics show on the radio, and that one of the pundits on the show – a well-known Labour supporter – predicts that Labour will win next year's general election. The confident manner in which he makes his prediction suggests that he genuinely believes it, and there is no reason to think he is being insincere, but the evidence he is able to cite in support of his prediction (opinion polls that marginally favour Labour, the alleged popularity of the leader amongst the general public) is far from convincing. Is the pundit, or is he not, guessing that Labour will win next year's election? It seems to me that he is obviously only guessing what the outcome of the election will be, for he certainly doesn't know that Labour are going to win. However, this verdict sits ill with the thought that believing is typically incompatible with guessing. Isn't this a case in which a subject is guessing that p despite the fact that he believes that p , and is

perfectly well aware that he believes that p ? How are we to square this with the intuition that, in the Battle of Hastings case, I am not guessing that the Battle of Hastings occurred in 1266 *because*, as a result of a typographical error, this is what I believe? One thought is that this confusion is a consequence of the fact that in ordinary language we do not use the term 'believes' in an entirely uniform way. There is a sense in which the pundit 'believes' that Labour will win—after all, this is the line he takes in his newspaper columns and on radio shows, and he feels pretty confident that it is right. But we might also want to say that the pundit does not believe, *outright*, that Labour will win, in the way in which he believes that $7 + 5 = 12$, or believes that Labour won by a landslide majority in 1997. The manifestation theorist can take this point on board by holding that it is only the outright belief that P that is typically incompatible with guessing that p . Thus, despite the fact that the pundit's prediction in some sense manifests his belief that Labour will win, it can still constitute a guess, because the 'belief' it manifests is not an outright belief but only a fairly high degree of confidence that Labour will win. However, if the manifestation theorist takes this line, then he owes us some account of how having the outright belief that p differs from having only a high degree of confidence that p . A plausible approach to this distinction is to connect it with the subject's higher-order attitudes about what he knows. For example, we might say that the pundit does not hold the outright belief that Labour will win because he is well aware that the evidence he has is not sufficient for him to know that this is the case. But if the manifestation theorist pursues this line, then it's not clear why we should prefer his account of guessing over the apparently simpler account that instead appeals directly to the subject's higher-order attitudes about what he knows, especially given that this

account avoids commitment to the controversial idea that judging that p essentially involves the manifestation of the belief that p .

In my view, we can adequately explain why guessing that p is typically incompatible with believing that p , and strictly incompatible with judging that p , if we hold that the relationship between guessing that p and knowing that p is to be accounted for in higher-order terms. The most obvious suggestions about the nature of the higher-order relationship between guessing that p and knowing that p are these:

- (a) In answering that p at t , S is guessing that p only if S believes that he does not know that p at t .
- (b) In answering that p at t , S is guessing that p only if S does not believe that he knows that p at t .

Unfortunately, it follows from points already made that neither (a) nor (b) identifies a genuine necessary condition of guessing that p . (a) is refuted by the observation that it is possible for you to guess that p despite the fact you know that p , and are aware, when you make your guess that p , that you know the answer to the question, despite your inability to recall it. In such a scenario, I think, you may have guessed that p despite the fact that you did not believe that you did not know that p . You did not believe that you did not know that p because you knew that you knew the answer to the question and, from your present perspective, p may well be the answer. Since p may well be the answer, from your present perspective, and you know that you know the answer, it may well be the case from your present perspective that you know that p , so you don't believe that you don't know that p . Thus, (a) fails to identify a genuine higher-order necessary condition of guessing that p . (b) is refuted by iterated application of the point that it is possible for you to know that p despite

the fact that you are unable to recall that p . Suppose that you know p and also believe yourself to know p , but that neither your knowledge nor your belief is presently accessible. In appropriate circumstances such a condition is one in which you could be led to guess that p , in line with the kind of example described above, despite the fact that you believe yourself to know that p . Consequently, (b) also fails to identify a genuine higher-order necessary condition of guessing that p .

How should we respond to the failure of (a) and (b) to correctly articulate necessary conditions of guessing that p ? One suggestion is that what is necessary for guessing p is not the presence (or absence) of higher-order attitudes about whether one is in a certain epistemic *state of mind*—viz. knowing that p —but rather the presence (or absence) of higher-order attitudes about the epistemic status of a *particular mental event*, viz. one's act of answering that p . Once again, there are two obvious ways of implementing this proposal, corresponding to the higher-order conditions (a) and (b) we have already considered:

- (c) In answering that p at t , S is guessing that p at t only if S believes that he is not answering that p knowledgeably at t .
- (d) In answering that p at t , S is guessing that p at t only if S does not believe that he is answering that p knowledgeably at t .

Does either of these conditions identify a genuine necessary condition of guessing that p ? Condition (c) seems too strong. Someone who is asked to name the year of the Russian Revolution might have a rather weak apparent memory that it took place in 1915, and consequently be unsure about whether he is answering knowledgeably in answering that it occurred in 1915. If (c) is right, then he is not guessing in answering 1915, because he does not positively believe that he is not answering

knowledgeably. Yet, it seems to me natural to say that he may still be making a guess about the answer. This might lead us to reject (c) in favour of (d), but (d) seems to face the opposite problem, that of being too weak to exclude certain cases as instances of guessing. For example, consider a case in which a subject unreflectively asserts that p in response to the question of whether p . There seems to be no good reason to think that such a subject *must* have a higher-order belief to the effect that he is answering that p knowledgeably. But if he lacks this belief then, for all (d) says, his assertion may constitute a guess that p . If we want the higher-order condition linking guessing with knowing to exclude such cases, we need a stronger higher-order condition than (d).

The following condition steers a middle course between (c) and (d):

- (e) In answering that p at t , S is guessing that p at t only if either (i) S believes that he is not answering that p knowledgeably at t , or (ii) S suspends judgement about whether he is answering that p knowledgeably at t .

This condition does not preclude the possibility that the subject is guessing in the Russian Revolution case, because a subject who is unsure about whether he can remember the answer to a question suspends judgement about whether his answer to the question is knowledgeably. But the condition is also strong enough to explain why cases of unreflective assertion are not instances of guessing, because an unreflective subject neither believes himself not to be answering knowledgeably nor suspends judgement about whether he is answering knowledgeably.³¹

³¹ For a discussion of the nature of suspended judgement that includes the point that neither believing p nor believing *not-p* is insufficient for suspending judgement about whether p , see Friedman (2013).

How can the proponent of the higher-order account of the relationship between guessing p and knowing p account for the fact that believing p is typically incompatible with guessing p ? Recall that (e) says that S is guessing that p only if either (i) S believes that he is not answering that p knowledgeably, or (ii) S suspends judgement about whether he is answering that p knowledgeably. The proponent of (e) can accordingly explain why a subject who believes that p typically won't be guessing that p in answering that p by appealing to the idea that a subject who believes that p will typically *neither* believe that he is not answering that p knowledgeably *nor* suspend judgement about whether he is answering that p knowledgeably. For, typically, someone who believes that p will either answer that p unreflectively, i.e. without considering whether he is answering knowledgeably, or will believe himself to be answering knowledgeably. Thus, the typical incompatibility of believing p with guessing p is a consequence of (e). An advantage of this way of explaining why believing p is typically incompatible with guessing p is that it explains why having a hunch, feeling or suspicion that p is compatible with guessing p , because someone who only has a hunch, feeling or suspicion that p will often satisfy condition (e) in answering that p . Indeed, even a subject who feels confident that p may admit, on reflection, that he does not (or at least may not) know that p . That we normally allow this possibility is suggested by our reaction to the political pundit case discussed above.

It follows from this account of the difference between judging that p and guessing that p that there is a constitutive connection between whether S is judging that p , in affirming that p , and S's higher-order attitudes concerning the epistemic status of his affirmation that p . In particular, when S judges that p , S neither believes himself not to be affirming that p knowledgeably nor suspends judgement about whether he is

affirming that p knowledgeably. It should be noted, however, that this claim about the relationship between judging and the subject's higher-order epistemic attitudes is purely negative: all that is required for a subject's affirmation that p to constitute a judging that p is the *absence* of certain higher-order epistemic attitudes. I have not argued that there is any positive respect in which a subject takes himself to know that p when he judges that p . Whether a further positive claim might be made here is an issue I will be considering in the next chapter.

5. Summary

In this chapter I have considered the conditions under which a subject is guessing that p in answering that p to some question. Some of the claims I have made are purely negative. Contrary to Sorensen, I have argued that a subject may be guessing that p even if believes that some alternative answer to the relevant question is more probable, and that he may be guessing that p even if he believes and, indeed, knows, that p . I have also argued that satisfying Sorensen's conditions is not sufficient for a subject to guess that p . Contrary to Owens, I have argued that it is not necessary for a subject to guess that p that he guesses with the intention of guessing that p only if p is true. In cases of indifferent guessing, a subject's answers to questions are subject to the standard of correctness constitute of guessing in virtue of features of the situation that are independent of the aim with which he answers those questions—e.g. in virtue of the fact that he is participating in a quiz, or a guessing game. Although I have not attempted to provide an alternative set of necessary and sufficient conditions for guessing that p , I have emphasised two necessary conditions of guessing that seem to me to express particularly important features of this kind of mental act. The first is that truth is the constitutive standard of success for guesses.

The second is that there is a higher-order connection between guessing that p and knowing that p : when S guesses that p , S either (a) believes himself not to be answering that p knowledgeably or (b) at least suspends judgement about whether he is answering that p knowledgeably. Such higher-order attitudes are absent when S judges that p . In the next chapter, I will be considering the significance of this conception of judgement for the question of the conditions under which a subject will acquire a belief about whether p as a consequence of engaging in enquiry.

Chapter 4

How to Answer the Activity Question

The previous two chapters were dedicated to questions about the nature of belief, judgement, and other mental events, such as guessing and supposing. Although these questions seem to me interesting in themselves, a further reason for pursuing these issues is that they may be relevant to responding to the problematic reasoning that I presented in Chapter 1. They may, therefore, be relevant to answering the activity question—‘Why do we seek knowledge, and not merely true beliefs, given that true belief is just as true, and just as useful, as knowledge?’—for insofar as there is an interesting challenge to the rationality of seeking knowledge embodied in that question, it seems to amount to the challenge expressed by the problematic reasoning. If this is right, then issues about the nature of belief, judgement, and other mental acts may be significant in a further way that is not usually appreciated. I argued in the Introduction that the activity question raises a more basic issue about the value of knowledge than the comparative state question (‘Why is someone who knows that p in a better position, other things being equal, than someone who merely truly believes that p ?’) that is traditionally pursued in philosophical discussions of the value of knowledge, and that tackling the comparative state question is only one way (and not a very promising way) of tackling the activity question. If it is possible to answer the activity question by drawing on considerations about the nature of belief and judgement, without presupposing an answer to the comparative state question, then the obvious motivation for pressing the comparative state question lapses. Ultimately, therefore, issues about the nature of belief, judgement, and other mental acts may be integral to the resolution of philosophical puzzles about the value of knowledge. My aim in this chapter is advance the case for this claim by showing

how the proposals I have advanced about belief and judgement can be used to undermine the problematic reasoning.

Here is the plan. § 1 recapitulates the problematic reasoning and the main critical points I made about it in Chapter 1. § 2 argues for the thesis that the enquirer will not have arrived at so much as an outright belief about whether p until he takes himself to have found out whether p . I indicate how the claims I have made about belief and judgement can be used to construct an argument for this thesis, and explain why certain examples in epistemology and philosophy of mind that might be thought to refute it actually fail to do so. In § 3 I explain how the thesis can be used to undermine the problematic reasoning. Since it is not obvious what other challenge someone might have in mind in pressing the activity question (aside from the challenge raised by Williams, which I have already dismissed), it is reasonable to conclude, at least provisionally, that the activity question can be resolved without appealing to an answer to the comparative state question. I also briefly consider an alternative version of the activity question. Although I have focussed on the issue of why it is not rationally sub-optimal for an enquirer to seek knowledge of whether p , given the possibility of merely seeking a true belief about whether p , one might also ask why it is not rationally sub-optimal for an enquirer to seek knowledge of whether p , given the possibility of merely attempting to arrive at an accurate assumption about whether p for practical purposes. § 4 briefly indicates how the claims I have made about the acquisition of belief might be taken to bear on the impossibility of acquiring beliefs at will.

1. The Problematic Reasoning Again

In Chapter 1 I introduced the problematic reasoning, which poses a challenge to the rationality of seeking knowledge. The reasoning went as follows:

1. Having a true belief about whether p is just as good for the enquirer as knowing whether p , because true belief is just as true, and just as useful, as knowledge.
2. Having a true belief about whether p is necessary but insufficient for knowing whether p .
3. Therefore, it is more difficult for the enquirer to arrive at knowledge of whether p than it is for the enquirer to arrive at a true belief about whether p .
4. If objective X is just as good as objective Y , but it is more difficult for the agent to achieve objective Y than objective X , then it is rationally sub-optimal for the agent to adopt objective Y as his aim; he should, at most, adopt objective X .
5. Given (1), (3) and (4), it is rationally sub-optimal for an enquirer to adopt the aim of arriving at knowledge of whether p ; he should, at most, adopt the aim of arriving at a true belief about whether p .

I suggested that this reasoning is one source of motivation for pressing the activity question. I argued that we cannot adequately respond to the reasoning by denying premise (1) or the validity of the inference from (2) to (3). The problem is that such denials provide at best a partial response to the reasoning; in response to these objections, the proponent of the reasoning can restrict the premises so that they avoid the objections but nevertheless entail a restricted form of the conclusion—e.g. that,

in circumstances in which the persistence of true belief does not matter, it is rationally sub-optimal for an enquirer to adopt the aim of arriving at knowledge of whether p . Neither can we adequately respond to the reasoning by rejecting premise (2), the claim that truly believing that p is necessary but insufficient for knowing that p . Merely denying that true belief is necessary for knowledge is of little help, because in that case it can be argued that having a true belief about whether p or knowing whether p is necessary but insufficient for knowing whether p , and therefore that it is rationally sub-optimal to seek knowledge given that one could set oneself the disjunctive objective instead. Furthermore, Sartwell's attempts to cast doubt on the insufficiency of true belief for knowledge are unconvincing. If we wish to defend the rationality of seeking knowledge across a broader range of cases, we will have to respond to the reasoning by rejecting the maxim of rationality expressed by premise (4). The issue, of course, is how to do this in a non-dogmatic way.

In Chapter 1, I pointed out that the cases in which we appeal to principle (4) (or something like it) to justify the claim that the agent should pursue only some more modest objective are ones in which certain further conditions are met. One of these conditions, for example, is that the agent must be capable of recognizing that there are ways in which he can attempt to achieve the more modest objective that are not ways in which he can attempt to achieve the more ambitious one. It is not obvious that these further conditions are met in the case of true belief and knowledge. For example, one might think that an enquirer cannot even attempt to arrive at a true belief about whether p using some method unless he regards that method as a potential way of finding out whether p . In fact, as I shall argue, this claim is incorrect, but highlighting it helps us to see how we might challenge the fourth premise of the problematic reasoning.

In the next section, I shall argue for a claim that is similar to the claim that an enquirer cannot even attempt to arrive at a true belief about whether p using some method unless he regards that method as a potential way of finding out whether p . The claim I shall be arguing for is that an enquirer who is seeking the truth about whether p will not have arrived at so much as a belief about whether p by engaging in enquiry until he takes himself to have found out whether p . In § 3 I will explain how it follows from this claim that conditions whose satisfaction we take for granted in appealing to (4) cannot be satisfied in the case of true belief and knowledge. It follows that it is illegitimate to appeal to that principle in this case. The problematic reasoning fails.

2. Thesis (T)

Chapters 2 and 3 advanced two proposals about the nature of belief and judgement that are significant in the present context. The first was that having the outright belief that p involves having a disposition to judge that p . I went on to argue that the problematic reasoning should be understood to concern outright belief. (From now on, I will use ‘belief’ to refer to the state of outright belief.) The second was that judging that p is to be distinguished from guessing that p , at least in part, in higher-order epistemic terms: when S guesses that p , S believes himself not to be affirming that p knowledgeably, or at least suspends judgement about whether he is affirming that p knowledgeably, but neither of these higher-order epistemic attitudes is present when S judges that p . What I want to do now is to indicate how these proposals can be used to defend the claim that an enquirer who is seeking the truth about whether p will not have arrived at so much as a belief about whether p by engaging in enquiry until he takes himself to have found out whether p . For the sake of convenience, I shall refer to this claim as ‘(T)’.

I want to begin by considering the conditions under which an enquirer judges that p . The proposal I made about how we are to distinguish judging from guessing in the previous chapter was, in a certain respect, purely negative: for an affirmation that p to qualify as a judging that p , the subject does not have to take himself to be affirming that p knowledgeably: all that is required is the *absence* of certain higher-order attitudes concerning the epistemic status of his affirmation. It would be overly hasty to infer from this, however, that an *enquirer* doesn't take himself to know that p when he judges that p . A subject who has the capacity to engage in enquiry is a subject who has the capacity to seek the truth about the answers to questions in a variety of different ways. Recognizing that he does not know the answer to a certain question, he can, for example, seek to find out the answer to that question, but he can also hazard a guess about the answer to that question, or aim to arrive at an accurate working assumption about the answer. It follows that a subject who has the capacity to engage in enquiry is a subject who has the capacity for higher-order beliefs about his epistemic condition, and also for higher-order beliefs about the epistemic statuses of his own cognitive mental acts. Arguably, when such a subject affirms that p , the absence of the belief that he is not affirming that p knowledgeably, and also of the attitude of suspended judgement about whether he is affirming that p knowledgeably, suffices for the subject to take himself to be affirming that p in a knowledgeable way. In other words, the absence of certain *negative* higher-order epistemic attitudes is sufficient for an enquirer to manifest a *positive* higher-order attitude concerning his epistemic standing with respect to the proposition he has affirmed.

More would need to be said to provide a full defence of this claim, but for now I just want to consider its implications for (T). Believing that p involves a disposition to judge that p , and when an enquirer judges that p he takes himself to be affirming that

p knowledgeably. Under what conditions will an enquirer enter into a state, as a consequence of engaging in enquiry, that involves a disposition to affirm that *p* in such a way that he takes himself to be affirming that *p* knowledgeably? Whether the enquirer acquires such a state will depend, I think, on what he takes himself to have done as a consequence of engaging in his activity: in particular, he will not acquire such a state unless he takes himself to have found out whether *p*. If the enquirer does not take himself to have succeeded in finding out whether *p*—and to have succeeded in finding out that *p*, in particular—then he will not regard himself as knowing that *p*, and hence he will not be disposed to affirm that *p* in such a way that he takes himself to be affirming that *p* knowledgeably. But having such a disposition is a necessary condition of having the outright belief that *p*. Clearly, parallel reasoning could also be applied to the acquisition of the belief that *not-p* in enquiry. Given that acquiring a belief about whether *p* is a matter of acquiring either the belief that *p* or the belief that *not-p*, it follows that the enquirer will not have arrived at so much as a belief about whether *p*, as a consequence of engaging in enquiry, unless he takes himself to have found out whether *p*. But this just is thesis (T).

At this point, I imagine, some will want to object that there are obvious counterexamples to (T), and therefore that the grounds I have presented for accepting (T) must be flawed. My response to this objection is that simply asserting, in the abstract, that counterexamples to (T) exist, is unpersuasive; the objector needs to describe a specific case in which a subject who is considering the question of whether *p* acquires a belief about whether *p* without taking himself to have found out whether *p*. Obviously, I cannot anticipate every case that might be presented here. I shall merely consider two examples that might be offered as demonstrations of this possibility, both of which are familiar from debates in epistemology and philosophy

of mind. The first case is that of a subject who ‘believes’ that his ticket will not win the jackpot in a lottery. The second case is one of the examples presented by Ginet (2001) in support of his contention that it is possible to acquire beliefs ‘at will’. Unsurprisingly, I shall be arguing that neither example refutes (T).

Suppose that *A* has bought a ticket for the UK National Lottery. Calculating that the odds of all six of his numbers being drawn are approximately one in 14 million, *A* comes to believe that his ticket won’t win the jackpot, but *A* does not take himself to have found out that his ticket won’t win; like most of us, *A* thinks that he can’t come to know that his ticket won’t win just on the basis of the fact that it is extremely unlikely to do so. If this is right, (T) is false: an enquirer *can* arrive at a belief about whether *p* without taking himself to have found out whether *p*, for *A* can arrive at a belief about whether his ticket will win the jackpot without taking himself to have found out whether his ticket will win the jackpot. To arrive at a belief about whether *p* by engaging in enquiry, an enquirer at most has to take himself to have evidence that it is highly probable that *p*.

It is undeniable that there is a sense in which *A* has arrived at a belief about whether his ticket will win. He takes himself to have calculated that the odds of his ticket winning the jackpot are roughly 1 in 14 million, and therefore believes that these are the odds of his ticket winning, and there is a clear sense in which the belief that the odds of his ticket winning the jackpot are roughly 1 in 14 million is a belief about whether his ticket will win. However, in the sense in which I am using the expression here, a subject has a belief about whether *p* just in case he either believes that *p* or believes that *not-p*. It is much less clear that *A* believes, *simpliciter*, that his ticket won’t win the jackpot, and therefore has a belief about whether his ticket will win in this sense. Here it will be no doubt be pointed out that it is quite natural to say

that *A* believes that his ticket won't win the jackpot, and not only that *A* believes that it is overwhelmingly likely that his ticket won't win the jackpot. However, sentences of the form 'S believes/thinks that *p*' are often used as shorthand to report the fact that S believes that *p* is more likely than not, or even (more minimally) that S believes that the most probable answer to a certain question is that *p*. Imagine, for example, that a football pundit is asked to predict who will win the Premier League next season. She suggests that Chelsea are most likely to win, given their performance last season and the pre-eminence of their manager, although it is possible that they will be overwhelmed by a resurgent Manchester City. It is natural to describe the pundit as believing/thinking that Chelsea will win, although strictly speaking her words indicate only that she believes that Chelsea are more likely than any other team to win the Premier League next season. The same might be held to be true of *A*. It is natural to say that *A* believes that his ticket won't win the jackpot because *A* believes that it is overwhelmingly likely that his ticket won't win, and we often use simple belief ascriptions to convey more complex facts about the subject's take on the balance of probabilities. A more careful characterisation of *A*'s beliefs would attribute him only the belief that his ticket is very unlikely to win the jackpot. If this is right, the example does not refute (T).

In 'Deciding to Believe' Carl Ginet provides four examples 'of the sorts of cases that seem to me good candidates for being described as someone's deciding to believe something' (2001: 64). His aim in doing so is to justify the claim, rejected as a metaphysically impossibility by Williams in his paper of the same name, that 'a person might come to believe something simply by deciding to do so' (2001: 63). Ginet's cases are relevant here because the examples he describes are ones in which the subject has what he regards as inconclusive evidence concerning the answer to

some question, and are thus ones in which the subject does not take himself to be in a position to know the answer to the relevant question. If the subject can nevertheless form a belief about the answer to the question just by deciding to do so, they will be cases in which a subject acquires a belief about the answer to a question without taking himself to have found out the answer to the question. This, of course, is exactly what is required for a case to constitute a counterexample to (T). But are Ginet's examples compelling? Here is the fourth case that Ginet describes:

We have started on a trip by car, and 50 miles from home my wife asks me if I locked the front door. I seem to remember that I did, but I don't have a clear, detailed, confident memory impression of locking that door (and I am aware that my unclear, unconfident memory impressions have sometimes been mistaken). But, given the great inconvenience of turning back to make sure and the undesirability of worrying about it while continuing on, I decide to continue on and believe that I did lock it. (2001: 64)

Many people, I think, are intuitively suspicious of the idea that this belief ascription is to be taken literally. Strictly speaking, they want to say, I decide only to *assume* that I did lock the door, in a sense of 'assume' that does not entail believing that I did lock it. It is a familiar fact, after all, that one can use the term 'belief', in a loose way, to refer to other ways of accepting a proposition. Thus, I might ask someone to suppose for the sake of argument that the witness is telling the truth by saying, 'Believe, for a moment, that she is telling the truth'. The mere fact that it is quite natural to describe me as 'deciding to believe' that I locked the door is not a good reason to think that I genuinely do decide to believe that I locked the door in this case. Furthermore, I have a good reason to deny that this is a case in which I form the outright belief that I locked the door. I argued that belief of this kind involves a

disposition to judge that p , and that one of the things that distinguishes judgement, in subjects like myself who have the capacity for higher-order epistemic attitudes, is that it involves the subject taking himself to know the proposition that he has judged to be so. In the case described by Ginet, however, I clearly do not take myself to know that I locked the door; I only seem to remember doing so, and I am aware that ‘my unclear, unconfident memory impressions have sometimes been mistaken’. Thus, if the current account of belief is correct, Ginet’s example is not one in which I acquire the belief that I locked the door.

What, then, should we say about the nature of the attitude that I decide to adopt in this case? A plausible proposal is that, strictly speaking, what I decide to do is assume for practical purposes that I locked the door. My attitude resembles that of belief, because I am disposed to act as though I locked the door. However, it is not a belief because I regard the constraint of treating as true the proposition that I locked the door as a self-imposed constraint. I am treating this proposition as true partly for practical reasons—it would be very inconvenient to turn back to make sure I locked it, and I do not wish to worry about it—and I am aware that I am doing so.

This completes my defence of (T). I have done two things. Firstly, I have indicated how the account of belief and judgement I proposed earlier might be used to construct an argument in favour of (T). Secondly, I have argued that (T) is not refuted by familiar examples from epistemology and philosophy of mind. I want now to consider how (T) bears on the problematic reasoning. In particular, can we use (T) to show that conditions that have to be included in the antecedent of any version of (4) that is a plausible consequence of the principle that rational agents conserve their resources cannot be satisfied in the case of true belief and knowledge?

3. Why Principle (4) Doesn't Apply in the Case of True Belief and Knowledge

In Chapter 1, I pointed out that we appeal to principle (4) (or something like it) to justify the idea that an agent ought to adopt a more modest objective only in cases in which certain conditions are met. One such condition is that the agent must be capable of recognizing that there are ways in which he can attempt to achieve the more modest objective that are not ways in which he can attempt to achieve the more ambitious objective. This condition is clearly met in the archery example that I used to illustrate the intuitive appeal of (4). If the only way in which the agent can attempt to achieve a more modest objective is by using methods that he regards as ways of achieving a more ambitious one, he cannot make life easier for himself by adopting only the more modest aim. I observed in Chapter 1 that it is not clear that this condition is met in the case of true belief and knowledge. As I put it there, there appears to be something in the thought that I can use a certain procedure to attempt to arrive at a true belief about whether p only if I regard that procedure as a potential way of finding out whether p . Is the truth of this thought a consequence of the conception of belief and judgement I have advocated? In particular, is it a consequence of thesis (T)?

To begin with, it should be pointed out that (T) is clearly distinct from the claim that an enquirer can attempt to arrive at a true belief about whether p using some procedure only if he regards that procedure as a potential way of finding out whether p . What (T) says is that an enquirer who is seeking the truth about whether p will not have arrived at so much as a belief about whether p by engaging in enquiry until he takes himself to have found out whether p . As such, (T) makes no direct claim about the kinds of procedure an enquirer can use in attempting to arrive at a true belief about whether p ; it makes a claim only about the enquirer must take himself to have

done if he is to have arrived at so much as a belief about whether p , and thus, *a fortiori*, a true belief about whether p . Still, this claim about what the enquirer must take himself to have done if he is to have succeeded in arriving at a true belief about whether p might be thought to entail (in conjunction with further assumptions) that an enquirer can attempt to arrive at a true belief about whether p using some procedure only if he regards that procedure as a potential way of finding out whether p . Here is one line of argument that might be proposed at this point. Suppose that the enquirer recognizes that (T) is true. He recognizes, that is to say, that he will not have arrived at so much as an outright belief about whether p by engaging in enquiry until he takes himself to have found out or determined whether p . (Such recognition needn't take the form of explicit knowledge that (T) is true—if it did, this claim would obviously be implausible.) Given that this is so, what kinds of procedure can the enquirer adopt with the aim of arriving at a true belief about whether p ? It might be thought that the only kind of procedure he can adopt are procedures that he regards as potential ways of finding out or determining whether p . If he does not regard a procedure in this way, then he will take there to be no chance of his arriving at so much as a belief about whether p by putting it into operation, because there is no chance that he will take himself to have found out whether p by doing so. Provided, then, that the enquirer recognizes that (T) is true, it might be argued that we can also establish the claim that an enquirer can attempt to arrive at a true belief about whether p using some procedure only if he regards it as a potential way of finding out whether p . And this is just to say that one of the conditions that must be met for (4) to be applicable isn't satisfied in the case of true belief and knowledge.

There is, however, a flaw in this argument. It relies on the assumption that if an enquirer does not regard a procedure as a potential way of finding out whether p , he

will take there to be no chance of him arriving at a belief about whether p using that procedure, given his recognition that he will not have arrived at a belief until he takes himself to have found out whether p . However, it appears to be possible, in principle, for a subject to induce beliefs in himself using procedures that he does not regard as ways of finding things out. Williams identifies reasons for allowing this possibility in the following passage, from 'Deciding to Believe':

[W]e all know that there are causal factors, unconnected with truth, which can produce belief: hypnotism, drugs, all sorts of things could bring it about that I believe that p . Suppose a man wanted to believe that p and knew that if he went to a hypnotist or a man who gave him certain drugs he would end up believing that p . Why could he not use this more roundabout method, granted that he cannot get himself into a state of believing just by lifting himself up by his own shoe straps [i.e. simply by deciding to hold the relevant belief]; why could he not bring it about that he believes that p by going to the hypnotist, the drug man or whatever? (1973: 149)

Given that I know that there are causal factors, unconnected with truth, that can induce belief, it seems to be possible, at least in principle, for me to intentionally manipulate such factors to induce a particular belief in myself. This is consistent with (T), provided that causal factors unconnected with truth can bring it about that I take myself to know that a certain proposition is true. Thus, even if an enquirer does not regard a procedure (e.g. seeing the hypnotist) as a way in which he might find out whether p , he may nevertheless know that, were he to put that procedure into operation, he would take himself to have found out whether p .

It appears, therefore, that there are two kinds of procedure by which an enquirer can attempt to arrive at a belief about the answer to some question. The first are procedures that the enquirer regards as potential ways of finding out the answer to that question. The second kind are procedures such that, although the enquirer does not regard them as ways of finding out the answer to the question, he nevertheless believes that, were he to put them into operation, he would take himself to have found out the answer to the question. In effect, what has been argued so far is that the possibility of procedures of the first type that issue in true belief but not in knowledge cannot be used to show that it is rationally sub-optimal for the enquirer to adopt the aim of acquiring knowledge of whether p , given that he could merely adopt the aim of acquiring a true belief about whether p . It remains to be seen, however, whether the possibility of procedures of the second type can be used to show this.

Here it might be objected that although an enquirer could use a procedure of the second type to arrive at a belief about whether p , he could not use such a procedure with the aim of arriving at a *true* belief about whether p . If this is right, then the only procedures an enquirer can employ with the aim of arriving at a true belief about whether p are procedures that he regards as ways of finding out whether p ; thus, he cannot avail himself of more efficient methods by adopting only the more modest aim of true belief. However, what is the basis for the allegation that an enquirer could not use a procedure that he does not regard as a way of finding out whether p with the aim of arriving at a true belief about whether p ? In order to ϕ with the aim of ψ -ing, it might be argued, the agent must believe that he will ψ by ϕ -ing, but the agent cannot believe that he will arrive at a true belief about whether p by going to see the brain surgeon, or the drug man; at most, he can believe that there is a chance that he will acquire a true belief about whether p in this way. In general, however, it

is false that in order to ϕ with the aim of ψ -ing the agent must believe that he will ψ by ϕ -ing. Guessing provides a clear illustration of this. When a subject guesses that p with the aim of guessing truly, he does not believe that he will guess truly by guessing that p ; he merely desires to guess truly, and believes that there is a chance of him doing so by guessing that p . Thus, it seems that, in general, in order to ϕ with the aim of ψ -ing, the agent must at most believe that there is *a chance* that he will ψ by ϕ -ing. A subject who goes to see the brain surgeon can surely satisfy this more minimal doxastic constraint with respect to the objective of acquiring a true belief about whether p . Hence, it seems consistent with the general doxastic constraints that apply to cases in which a subject ϕ s with the aim of ψ -ing that an enquirer could pay a surgeon to induce in him the belief that p with the aim of acquiring a true belief about whether p . Perhaps it will be responded that special doxastic constraints apply in the case in which the agent's objective is to arrive at a true belief that rule out the possibility of using such a procedure with this aim, but then the onus is on the objector to specify what these constraints are, and why their application is limited to a subset of cases that includes that of an enquirer whose aim is to arrive at a true belief about whether p .

We must concede that it is possible, in principle, for an enquirer to set out to deceive himself into thinking that he has found out whether p with the aim of acquiring a true belief about whether p . Thus, it is not the case that the only way in which an enquirer can attempt to arrive at a true belief about whether p is by using a procedure that he regards as a way of finding out whether p . But, even conceding that it is possible for an enquirer to pursue a true belief about whether p using such a procedure, it remains implausible that this can be used to show that, in cases in which we are pretheoretically disposed to judge that the enquirer can adopt the aim of finding out

whether p without being subject to rational criticism, it is actually rationally sub-optimal for him to do so.

An initial point here is that, as far as I know, the hypnotists, drug men and brain surgeons who you can pay to make you believe that p exist only in the writings of philosophers. If this is right, then in practice the only way in which ordinary human enquirers can arrive at beliefs about the answers to questions is by using procedures that they regard as ways of finding out the answers to those questions. The problematic reasoning therefore fails to show that, in any actual case, it is rationally sub-optimal for an enquirer to adopt the aim of finding out whether p . Leaving matters here, however, might be taken to imply that seeking knowledge is not a rationally sub-optimal strategy for us only because of contingent limitations we have as human enquirers. In fact, I think, it is possible to argue for a stronger conclusion: even if we were able to induce particular beliefs in ourselves just, as it were, by the flick of a switch, I do not think it would follow that, in cases in which we now judge enquirers to be behaving entirely properly in seeking knowledge, they would be acting in a rationally sub-optimal manner.

One point here is that, at least from my own point of view, I am far more likely to end up with a *true* belief about whether p by employing a procedure that I regard as a way of finding out whether p than I am by intentionally manipulating causal factors, unconnected with the truth, to induce in myself a belief about whether p . This reminds us of the basic point that, in selecting means to their ends, agents should be sensitive to considerations about how likely they are to achieve their end by a given means, as well as considerations about how efficient those means are. But this is not the only point that is relevant. Also significant is the fact that the alternative to using a method that I regard as a way of finding out whether p is to *deceive* myself into

thinking that I have found out whether p . It seems reasonable to assume that enquirers, as subjects who engage in the pursuit of truth, have a strong aversion to being deceived. This point is strengthened by the following observation. Although I can arrive at a belief about whether p by deceiving myself, so that I take myself to know whether p , such a belief will not be stable in the face of future evidence unless it is 'backed up', as it were, by further beliefs about how I know whether p . So to arrive at a stable belief about whether p , I will have to deceive myself, not only about whether I know whether p , but also about other related matters, e.g., by implanting in myself beliefs about how I know whether p . The cumulative effect of these observations is to make this alternative way of attempting to arrive at a true belief about whether p deeply unattractive. Not only is there a high chance of error, but I will also have to deceive myself about my epistemic standing with respect to a certain proposition, where this may involve not only taking myself to know it when I don't, but also holding false beliefs about how I know it.

But what if it is important for me to have the truth about whether p , and there is no way in which I can find out whether p , or it would at least be very difficult for me to do so? Isn't it the case that, in these circumstances, I should adopt the aim of arriving at a true belief about whether p , and deceive myself into taking myself to know whether p ? The first thing to say about this is that, in these circumstances, it may well be the case that I shouldn't adopt the aim of finding out whether p . If I know that I cannot find out whether p , then I cannot even decide to try to find out whether p . If I know that it is very difficult for me to find out whether p then, even if it is important for me to have the truth about whether p , it may still be better for me to prioritise my other objectives and not attempt to find out whether p . Thus, the point that in such circumstances I ought at most to adopt some aim that falls short of

finding out whether p hardly calls into question the rationality of ordinary enquiry. We would not pretheoretically take it to be rational for an enquirer to seek knowledge in circumstances of this kind.

It must also be noted, secondly, that if in these circumstances I should not adopt the aim of finding out whether p , then I should not adopt the aim of arriving at so much as a true belief about whether p . At most, I should adopt the aim of arriving at an accurate working assumption about whether p . If I shouldn't adopt the aim of finding out whether p , this must be because the procedures by which I can attempt to find out whether p involve difficulties too great to justify their use. In principle, there are other procedures by which I can attempt to arrive at a true belief, as we saw, but I am unlikely to arrive at true beliefs by such procedures, and they involve self-deception. Even if such procedures were available to me, it would surely be better for me to adopt the aim of arriving at an accurate working assumption about whether p instead. Accurate working assumptions appear to be just as good, for the purposes of successful action, as true beliefs. Echoing Plato's point, someone who has an accurate working assumption about the way to Larissa will get there just as well as someone who has a true belief about the way. I can arrive at working assumptions about the answers to questions, without using methods that I regard as ways of finding out the answers to those questions, without resorting to self-deception. So by adopting the more modest aim, I avoid a significant cost of pursuing the aim of true belief in the only way available. Furthermore, I am no more likely to arrive a *true* belief about whether p by such means than I am to arrive at an *accurate* working assumption. Paying a hypnotist to induce in me a belief about whether p is no better than making a pure guess about whether p ; I am just as likely to end up with the truth by simply deciding to adopt, at random, a certain working assumption about whether

p. So it would seem that, when I shouldn't adopt the aim of finding out the answer to some question, I should at most adopt the aim of arriving at an accurate working assumption about the answer.

The point that an enquirer can merely adopt the aim of arriving at an accurate working assumption, or estimate, about the answer to a question may raise a further worry about the rationality of pursuing knowledge. Why should an enquirer who wants the truth about whether *p* ever want to know whether *p*, rather than merely to arrive at an accurate working assumption about whether *p*? If the enquirer simply adopts a working assumption about whether *p* at random, there is a high chance that the answer to the question he adopts as his assumption will be incorrect. It is surely sensible, therefore, for him to attempt to estimate which of *p* and *not-p* is more likely to be correct. Attempting to estimate whether *p*, in this sense, is a matter of working out which of *p* and *not-p* is more likely to be true. It is, therefore, to engage in a form of epistemic activity. Consequently, it is a mistake to think that merely attempting to arrive at accurate working assumptions represents a serious alternative to engaging in forms of epistemic activity; the need to minimise the risk of error brings with it a motivation to engage in just such activity. Now, in many cases it will be no easier for the enquirer to arrive at an accurate estimate about whether *p* than it would be for him to find out whether *p*. For example, it is no easier for me to determine whether it is more likely than not that Jane Austen was born in 1783 than it is for me to find out whether she was born in 1783. Furthermore, when it is important for practical purposes that the enquirer acts on an accurate assumption about whether *p*, the greater convenience that comes from lazily estimating whether *p* may be more than outweighed by the higher chance of error. It is only when knowledge is hard or impossible to come by, or the costs of error not too great, that the enquirer should

merely attempt to estimate and adopt accurate working assumptions, rather than find things out.

To summarize, I have been arguing in this section that it follows from (T) that the fourth premise of the problematic reasoning does not apply in the case of true belief and knowledge. The fourth premise of the reasoning states that if objective X is just as good as objective Y , but it is more difficult for the agent to achieve objective Y than objective X , then it is rationally sub-optimal for the agent to adopt objective Y as his aim. In the face of the problematic reasoning, I have conceded that, at least in circumstances in which the stability of true belief does not matter, true belief is just as good as knowledge. I have also conceded that there is a sense in which true belief may be easier to come by than knowledge: sometimes, there is a procedure available to the enquirer such that (a) that procedure issues in a true belief about, but not knowledge of, whether p , and (b) that procedure is less taxing on the enquirer's resources than any procedure he could use that would issue in knowledge of whether p . It seems to follow, given (4), that where the stability of true belief is unimportant, and there is some procedure available to an enquirer that satisfies conditions (a) and (b), the enquirer should at most adopt the aim of arriving at a true belief about whether p . However, I also noted, in Chapter 1, that when we appeal to principle (4) to justify the contention that an agent ought to adopt only some more modest aim, we presuppose that further conditions are met. Particularly relevant to the present issue is, I think, the following presupposed condition: the agent must be capable of recognizing that there are ways in which he can attempt to achieve the more modest objective that are not ways in which he can attempt to achieve the more ambitious objective. Is this condition satisfied in the case of true belief and knowledge? It might be thought to follow from (T) that it can't be. According to (T), the enquirer

will not have arrived at so much as a belief about whether p until he takes himself to have found out whether p . However, I argued that, even given (T), it is possible for an enquirer to seek a true belief about whether p using a procedure that he does not regard as a way of finding out whether p : he can attempt to deceive himself into taking himself to have found out whether p . Once it is appreciated, though, that this is the only alternative to using a procedure that he regards as a way of finding out whether p , even supposing such procedures to be available, it is obvious why it would not be rational for enquirers to use them. They are unlikely to issue in true beliefs, and they involve self-deception. Even supposing, then, that the easiest way for an enquirer to arrive at a belief about whether p is by taking a pill or seeing a hypnotist, such methods are far less effective (i.e. more likely to issue in false beliefs, at least from the enquirer's perspective), and involve unpleasant side effects (i.e. false beliefs about what I know, and how I know it). The idea, then, that an enquirer can avail himself of superior methods by merely adopting the aim of true belief is a chimera. This, I think, is the fundamental point that needs to be made in response to the problematic reasoning. Although it is legitimate to appeal to principle (4) in some cases, like the archery example, it is not legitimate in the case of true belief and knowledge.

The claim that an enquirer will not have arrived at so much as a belief about whether p until he takes himself to have found out whether p is relevant to how we should respond to the problematic reasoning. It is also relevant, I think, to the question of why it is impossible to acquire beliefs at will that Williams raises in 'Deciding to Believe'. In the next section, I want to briefly indicate how (T) might be taken to be relevant to this issue.

4. Deciding to Believe

What is the relationship between belief and decision? In ‘Deciding to Believe’, Williams contends that it is impossible to acquire beliefs simply by deciding to do so, or ‘at will’, and that this is not merely a contingent limitation of human beings but a necessary restriction of the will of any possible subject. As we saw earlier, the claim that it is impossible to form beliefs at will has not gone unchallenged. Ginet argues that ordinary human subjects form beliefs at will in deciding to act, or not to act, in certain ways. Others, whilst conceding that ordinary human subjects do not have capacity to form beliefs at will, nevertheless argue that we should regard this as only a contingent limitation of our wills.³² There is no reason to deny that other believers might, in principle, have this ability. According to this view, our inability to form beliefs at will is similar to our inability to blush at will, which plausibly is a merely contingent limitation of our wills.³³ This view is motivated, in part, by the conviction that the explanation of the impossibility of forming beliefs at will that Williams offers in ‘Deciding to Believe’, based on the idea that belief aims at the truth, is inadequate.³⁴ However, once we have thesis (T) on the table, an alternative kind of explanation becomes available. Briefly put, the impossibility of forming beliefs at will, just by deciding to do so, is a consequence of the subject’s awareness that he cannot know things at will.

Let me spell this out a little. Suppose, in accordance with (T), that the subject will not have arrived at the belief that *p* until he takes himself to have found out that *p*. Suppose, furthermore, that the subject knows that he cannot find out that *p* simply by

³² See, for example, Alston (1988).

³³ Provided, at least, that what we mean by ‘blushing’ is simply going red in the face, and not showing shyness, embarrassment or shame by going red in the face.

³⁴ For a good critical discussion of Williams’ argument, see Winters (1979).

deciding to do so. If such a subject were to form the belief that p at will, in full consciousness, then he would know that he had just done so. But given that he knows that he cannot find out that p at will, he would not regard himself as having just found out that p . But if he does not regard himself as having found out that p , then he will not have acquired so much as the belief that p . Neither will he *regard* himself as having acquired the belief that p , provided we assume that the subject recognizes the connection between believing and taking oneself to know. It follows that the subject is unable to acquire beliefs at will, in full consciousness, and also that he recognizes that he is unable to do so.

Naturally, this explanation of the impossibility of forming beliefs at will raises further questions. Many have the intuition that it is impossible for S to form the belief that p at will even if S knows that, were he to believe that p , then it would be true that p . An example that is often given as an illustration of this possibility is that someone trustworthy tells S that she will give him a million pounds if S believes that she will give S a million pounds. Does the above explanation apply in cases of self-fulfilling belief? One might think that in such circumstances the subject can know that p at will. After all, if he believes that p then it will be true that p , and the subject knows this. Thus, the subject might know that he had acquired the belief that p at will and nevertheless regard himself as knowing that p . One might also wonder, more generally, about how the subject knows that he cannot acquire knowledge at will. For example, is this to be explained in terms of the subject's grasp of some necessary condition for knowing that p , and his appreciation that this condition would be violated by a belief that p that was acquired at will? I will not attempt to resolve these questions here; my intention is just to highlight some further issues that are raised by the present explanation of the impossibility of forming beliefs at will.

5. Concluding Remarks

My central aim in this thesis has been to suggest an alternative approach to philosophical issues about the value of knowledge, one that brings issues about enquiry to the fore. The position I have advocated makes two key moves. Firstly, it argues that the comparative state question that usually receives attention in philosophical discussions of the value of knowledge is underwritten by a more basic activity question that concerns enquiry. Secondly, it argues that there is an important epistemic dimension to belief: an enquirer who has the outright belief that p is disposed to judge that p , and judging that p involves taking oneself to know that p . I have used this conception of belief to argue for a significant constraint on the acquisition of belief in enquiry, thesis (T). Once the idea that an enquirer will not have arrived at so much as a belief about whether p until he takes himself to know whether p is in play, it is intelligible why merely seeking true beliefs does not represent a genuine alternative, from the enquirer's perspective, to seeking knowledge. I have also indicated how (T) might be used to explain the impossibility of acquiring beliefs at will.

One upshot of my view is that it is no longer clear that the comparative state question that usually receives attention in philosophical discussions of the value of knowledge is worthy of serious attention. If we can explain why enquirers seek knowledge without appealing to the idea that knowledge is a superior state, it is no longer clear why there should be a general presumption that we will be able to identify respects in which knowledge is superior to (e.g.) mere true belief. At the very least, philosophical discussion of the comparative state question ought to proceed more critically, with an eye on the broader purpose such discussions serve.

Bibliography

- Alston, W. (1988) 'The Deontological Conception of Epistemic Justification'. *Philosophical Perspectives*, 2 Epistemology: 257–299.
- Armstrong, D. (1969) 'Does Knowledge Entail Belief?'. *Proceedings of the Aristotelian Society*, 70: 21–36.
- Aristotle (1995) *The Complete Works of Aristotle*. Edited by J. Barnes. Sixth printing with corrections. Princeton: Princeton University Press.
- Bratman, M. (1987) *Intention, Plans, and Practical Reasoning*. Cambridge MA: Harvard University Press.
- Bratman, M. (1992) 'Practical Reasoning and Acceptance in a Context'. *Mind*, 101 (401): 1–16.
- Cassam, Q. (2010) 'Judging, Believing and Thinking'. *Philosophical Issues*, 20 (1): 80–95.
- Cohen, J. (1992) *An Essay on Belief and Acceptance*. Oxford: Oxford University Press.
- Craig, E. (1990) *Knowledge and the State of Nature*. Oxford: Oxford University Press.
- Dummett, M. (1973) *Frege: Philosophy of Language*. London: Duckworth.
- Gettier, E. (1963) 'Is Justified True Belief Knowledge?'. *Analysis*, 23: 121–123.
- Ginet, C. (2001) 'Deciding to Believe'. In: M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford: Oxford University Press: 63–76.
- Ichikawa, J. J. and Steup, M. (2013) 'The Analysis of Knowledge', *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), E. N. Zalta (ed.), <<http://plato.stanford.edu/archives/fall2013/entries/knowledge-analysis/>>.
- Kripke, S. (2011) 'On Two Paradoxes of Knowledge'. In: S. Kripke, *Philosophical Troubles*. Oxford: Oxford University Press: 27–51.
- Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lackey, J. (1999) 'Testimonial Knowledge and Transmission'. *The Philosophical Quarterly* 49: 471–490.
- O'Brien, L. (2005) 'Imagination and the Motivational View of Belief'. *Analysis* 65: 55–62.

- Owens, D. (2003) 'Does Belief Have an Aim?'. *Philosophical Studies*, 115: 283–305.
- Peacocke, C. (1999) *Being Known*. Oxford: Oxford University Press.
- Plato (2005) *Meno and Other Dialogues*. Translated by R. Waterfield. Oxford: Oxford University Press.
- Pritchard, D. (2010) 'Knowledge and Understanding'. In: D. Pritchard, A. Millar and A. Haddock, *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press: 3–88.
- Radford, C. (1966) 'Knowledge—By Examples'. *Analysis*, 27 (1): 1–11.
- Sartwell, C. (1991) 'Knowledge Is Merely True Belief'. *American Philosophical Quarterly*, 28 (2): 157–165.
- Sorensen, R. (1984) 'Unique Alternative Guessing'. *Logique et Analyse*, 27: 77–86.
- Soteriou, M. (2013) *The Mind's Construction: The Ontology of Mind and Mental Action*. Oxford: Oxford University Press.
- Shah, N. and Velleman, D. (2005) 'Doxastic Deliberation'. *The Philosophical Review*, 114 (4): 497–534.
- Steglich-Petersen, A. (2006) 'No Norm Needed: On the Aim of Belief'. *Philosophical Quarterly*, 56 (225): 499–516.
- Velleman, D. (2000) 'On the Aim of Belief'. In: D. Velleman, *The Possibility of Practical Reason*. Oxford: Oxford University Press: 244 – 281.
- Williams, B. (1973) 'Deciding to Believe'. In: B. Williams, *Problems of the Self*. Cambridge: Cambridge University Press: 136–151.
- Williams, B. (2005) *Descartes: The Project of Pure Enquiry*. London: Routledge.
- Williams, B. (2006) 'Knowledge and Reasons'. In: B. Williams, *Philosophy as a Humanistic Discipline*. Princeton: Princeton University Press: 47–56.
- Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Winters, B. (1979) 'Believing at Will'. *Journal of Philosophy*, 76 (5): 243–256.
- Zagzebski, L. (2003) 'The Search for the Source of Epistemic Good'. *Metaphilosophy*, 34 (1–2): 12–28.