



Thèse

2017

Public access

This version of the publication is provided by the author(s) and made available in accordance with the copyright holder(s).

Knowledge, certainty and practical factors

Vollet, Jacques-Henri

How to cite

VOLLET, Jacques-Henri. Knowledge, certainty and practical factors. Doctoral Thesis, 2017. doi: 10.13097/archive-ouverte/unige:95005

This publication URL: <https://archive-ouverte.unige.ch/unige:95005>

Publication DOI: [10.13097/archive-ouverte/unige:95005](https://doi.org/10.13097/archive-ouverte/unige:95005)

© This document is protected by copyright. Please refer to copyright holder(s) for terms of use.

Last deposit update in Archive ouverte UNIGE on 15.03.2023 02:47

KNOWLEDGE, CERTAINTY AND PRACTICAL
FACTORS

JACQUES-HENRI VOLLET

THÈSE DE DOCTORAT ÈS LETTRES

UNIVERSITÉ DE GENÈVE

Directeur de thèse:

Professeur PASCAL ENGEL, Université de GENÈVE et EHESS.

Président du jury:

Professeur MARCEL WEBER, Université de GENÈVE.

2017

ACKNOWLEDGMENTS

I discovered contemporary epistemology in 2010 when I undertook a Certificate of Specialization in Contemporary Philosophy at the University of Geneva. In 2011, I started my Ph.D. with Pascal Engel. He was the supervisor of my *mémoire de maîtrise* a few years before at the Sorbonne. I am extremely grateful to Pascal Engel for having offered me the opportunity to go further, and for his constant support during all these years. Without his encouragements, trust and expert guidance, I doubt that this project would have got off the ground.

I also owe a great debt to Mikkel Gerken and Jessica Brown for extensive feedback on drafts of several chapters during my stay in Edinburgh and St Andrews. I benefitted enormously from their comments.

This research has been supported by many institutions: the University of Geneva gave me a one-year scholarship from 2011 to 2012; from 2012 to 2015, I was a candoc within a project directed by Pascal Engel and Fabrice Teroni and financed by the Swiss National Science Foundation; I also received a scholarship from the Swiss National Science Fondation to study at the universities of Edinburgh and St Andrews from 2015 to 2016; finally, I received le Prix Philibert Collart from la Société Académique de Genève in 2016. I thank all these institutions for the opportunities they offered.

I also would like to thank in particular Miloud Belkoniene, Yves Bouchard, Evan Butts, Pablo Carnino, Ariel Cecchi, Jean-Marie Chevalier, Michel Croce, Julien Deonna, Julien Dutant, Davide Fassio, Jie Gao, Benoit Gaultier, Sandy Goldberg, Santiago Echeverri, Paul Egré, Jeremy Fantl, Jean-Baptiste Guillon, Jaakko Hirvela, Mikal Hladky, Kristopher Kelp, Federico Lauria, Arturs Logins, Matthew McGrath, Robin McKenna, Anne Meylan, Paolo Natali, Marie Robert, Mélanie Sarzano, Kegan Shaw, Martin Smith, Mona Simion, Fabrice Teroni, Claudine Tiercelin, Emile Thalabard, Marie van Loon, Marcel Weber and Pierre Willaime for helpful discussions and support. A special thank to Julien Dutant and Davide Fassio from whom I have learnt so much.

Last but not least, I wish to thank my family for their confidence and love. I dedicate this thesis to Laetitia who accompanied me during all these years.

Contents

1	Introduction	7
1.1	Orthodox views and shifty views of knowledge	7
1.2	Four arguments in favour of shifty views	8
1.3	Summary of the chapters	12
I	Justification of belief and practical factors	16
2	Intellectualism, pragmatism and pragmatic encroachment	17
2.1	Introduction	17
2.2	Truth-relevant factors, intellectualism and the norm of belief	18
2.3	Practical factors, pragmatism and pragmatic encroachment	21
2.3.1	Intellectualism, pragmatism and pragmatic encroachment	21
2.3.2	Pragmatic encroachment, reflective control and inconclusive evidence	24
2.4	Practical factors, intellectualism and belief formation	28
2.5	Conclusion	35
3	The threshold problem	36
3.1	Introduction	36
3.2	Inconclusive evidence and arbitrariness	37
3.3	Intellectualist approaches to the threshold problem	44
3.3.1	A fallibilist approach to the threshold problem	44
3.3.2	An infallibilist approach to the threshold problem	49
3.4	Inconclusive evidence and pragmatic factors	58
3.4.1	Practical factors and suspension of judgement	59
3.4.2	The connection between justified belief and rational action	62
3.5	Conclusion	65

II	‘Knowledge’ ascriptions	66
4	Linguistic intuitions and the rule of relevance	67
4.1	Introduction	67
4.2	The bank cases and the rule of relevance	68
4.3	The account	71
4.4	Objection 1: The rule of relevance cannot be put to work	74
4.5	Objection 2: The rule of relevance is insufficient, and hence, the account <i>ad hoc</i>	76
4.6	Objection 3: Rejection of the epistemic closure on knowledge or more <i>ad hocness</i>	79
4.7	Objection 4: The conversational context is underdescribed	82
4.8	Objection 5: Two unexplained facts	86
4.9	Conclusion	87
5	Linguistic intuitions and belief	88
5.1	Introduction	88
5.2	The relevance of practical factors to our beliefs	89
5.3	The confusion between the epistemic standards of action and those of belief	91
5.3.1	Bach’s account	91
5.3.2	Two problems: low stakes assessors and the kind of underlying error	94
5.4	High stakes and desire for truth	96
5.4.1	Need-for-closure and epistemic anxiety	96
5.4.2	Epistemic anxiety and ‘knowledge’ ascriptions	99
5.4.3	Four problems for Nagel’s account	102
5.5	Conclusion	112
6	Linguistic intuitions and the certainty norm of assertion	114
6.1	Introduction	114
6.2	The certainty norm of assertion	115
6.2.1	The epistemic norm of assertion	115
6.2.2	Subjective and epistemic certainty	116
6.2.3	Knowledge and certainty	118
6.2.4	Against scepticism about certainty: the meaning of ‘certain’	124
6.3	Arguments in favour of the certainty account of assertion	136
6.3.1	Moore-paradoxical sentences	136
6.3.2	Challenges	138
6.3.3	Unified explanation of Moorean sentences and appropriate challenges	142
6.3.4	Parenthetical uses	143
6.3.5	Six further arguments	144

6.4	Defence of the certainty account	147
6.4.1	Defence against Pritchard (2008)	147
6.4.2	Defence against Turri (2010)	149
6.5	The certainty norm of assertion and ‘knowledge’ ascriptions	152
6.5.1	Why ‘knowledge’ ascriptions shift with practical factors: the role of the certainty norm of assertion	152
6.5.2	Why a mixed approach outperforms the simple doxastic approach	156
6.5.3	Why a mixed approach outperforms the Gricean approach	157
6.6	Conclusion	160
 III Knowledge and rational action		161
7	The epistemic evaluation of actions	162
7.1	Introduction	162
7.2	The evaluation of action	163
7.2.1	The normative and the descriptive	163
7.2.2	Normative reasons and what you ought to do	164
7.2.3	Deontic and evaluative normativity	165
7.2.4	Deontic norms	166
7.2.5	Objective and subjective reasons	167
7.2.6	Subjective reasons, motivating reasons and rationality	172
7.2.7	Justification, rationality, excuse	173
7.3	The equivocality of the epistemic evaluation of action	180
7.3.1	Ways in which we can conceive of the epistemic evaluation of action	180
7.3.2	Unreflective and reflective rationality	185
7.3.3	Possibilities decisive for action, subjective reasons to act, and considerations warranted enough to justify the action	189
7.4	The main accounts	191
7.4.1	The character and the source of the epistemic norm of action	192
7.4.2	The epistemic condition of the norm, context-sensitivity and pluralism.	193
7.5	Conclusion	198
8	The necessity of knowledge for action, and the warrant account	199
8.1	Introduction	199
8.2	Arguments for the knowledge norm of action.	200
8.2.1	The linguistic argument	200
8.2.2	The bad reasoning argument	201

8.2.3	The principled arguments	202
8.3	The linguistic argument and the warrant account	205
8.3.1	The significance of the linguistic argument	205
8.3.2	The linguistic considerations and the warrant account	208
8.4	Arguments against the necessity of knowledge and the excuse manoeuvre	214
8.4.1	The epistemic cases and the excuse manoeuvre	214
8.4.2	The practical cases and the excuse manoeuvre	225
8.4.3	Criticism for not acting on p (weaker)	229
8.5	Conclusion	229
9	The context-sensitivity of the norm of action	231
9.1	Introduction	231
9.2	The pragmatic encroachment paradox and the context-insensitivity of the epistemic norm of action	232
9.2.1	The pragmatic encroachment paradox	232
9.2.2	The equivalence thesis	236
9.3	Insensitivism about the epistemic norm of action	237
9.3.1	The two norms approach	237
9.3.2	The insufficient reason approach	252
9.4	The non-maximalist argument in favour of the sensitivity of the epistemic norm of action	257
9.5	Conclusion	258
10	The insufficiency of knowledge for action, and the certainty account	260
10.1	Introduction	260
10.2	Against the sufficiency of knowledge	261
10.2.1	Cases in which knowledge is insufficient	261
10.2.2	Maximal infallibilism and the sufficiency of knowledge	265
10.2.3	Epistemic modals infallibilism and the sufficiency of knowledge	267
10.2.4	Against epistemic modals infallibilism: Rysiew and Dougherty's proposals	276
10.2.5	Against epistemic modals infallibilism: the certainty norm of assertion	283
10.3	The certainty account	286
10.3.1	Eight desiderata for an account of the epistemic norm of action	286
10.3.2	The eight desiderata and the main accounts: summary	289
10.3.3	The eight desiderata and the certainty norm	290
10.4	Conclusion	292
	Bibliography	293

Chapter 1

Introduction

1.1 Orthodox views and shifty views of knowledge

Epistemology is the theory of knowledge and rational belief (or theoretical rationality). Its aim is to investigate the conditions under which a belief can count as a piece of knowledge, and what it takes to have a justified belief. There has been a lot of disagreement on what these conditions should be, but epistemologists have generally accepted that these conditions must have to do with the truth of the belief. It is typically said, for example, that a subject knows something only if this subject has a reasonable true belief, where ‘reasonable’ means that the subject has sufficient reason to think that her belief *is true*. This is intuitive enough. For example, Paul can know that it is raining if he believes that it is raining because he sees that it is raining. Seeing that p seems to be a sufficient reason to believe that *p is true*. But John cannot know that it is raining on the basis of his dream that it is raining. Dreaming that p does not seem to be a sufficient reason to believe that *p is true*. It has also been proposed that the true belief should not be true by mere luck, or that it should be such that it could not be false, etc. In general, all these proposals aim at identifying a condition which connects the belief to the truth in a specific way. That this is the right approach has been taken for granted by virtually all epistemologists until recently, and it is not exaggerated to call the view according to which knowledge and justification of belief are only a matter of truth-relevant considerations ‘the orthodox view’.

Over the last thirty years, however, this orthodox view has been put into question. It has been suggested that besides truth, we should appeal to practical factors to explain what ‘knowledge’ and ‘justified belief’ mean. Some philosophers have defended a view called ‘epistemic contextualism’ (or ‘attributor contextualism’). This is the view according to which the meaning of ‘know’ depends on the context of utterance. According to these philosophers, how much justification you need to be truly ascribed ‘knowledge that p ’ depends on the epistemic standards which are operative in the

conversation where this sentence is uttered. And these epistemic standards depend on the practical situation. In particular, if it is really important to be right about p – the cost of error, or the stakes, are high – these standards are high. If it does not really matter whether p is true, these standards are lower.

Other philosophers have proposed the view according to which although the meaning of ‘know’ does not vary with the context, knowledge is affected by practical stakes. This view has received different names: ‘anti-intellectualism’, ‘impurism’, ‘subject contextualism’, ‘interest-relative invariance’, or ‘subject-sensitive invariantism’. The core idea is that if it is not rational for someone to take a certain proposition for granted in her action because, due to the cost of error, she should double-check, then this person does not know this proposition.

While different, epistemic contextualism and anti-intellectualism have in common that practical factors can affect whether a ‘knowledge’ ascription (i.e. an utterance of the form ‘ S knows that p ’, where ‘ S ’ is the subject to whom knowledge is ascribed, and ‘ p ’ stands for the target proposition) is true. That is, they have it that practical factors can affect whether someone can truly say ‘ S knows that p ’. In other words, whether someone can truly say ‘ S knows that p ’ varies with someone’s practical situation. For epistemic contextualists, the relevant practical situation is the one associated to the conversational context of the ascriber. For anti-intellectualists, the relevant practical situation is the one associated to the subject to whom ‘knowledge’ is ascribed. Since, on these views, the truth of ‘knowledge’ ascriptions is supposed to shift with someone’s practical context, we can call these views ‘shifty views (of ‘knowledge’ or knowledge)’.

In contrast to shifty views, the orthodox view is characterized by two commitments. First, in contrast to epistemic contextualism, the orthodox view is *invariantist* about the word ‘know’. It denies that the semantic content of this word changes across contexts of utterances. Second, in contrast to anti-intellectualism, the orthodox view is *intellectualist*. It maintains that whether or not a subject knows, or is justified in believing, does not *constitutively* depend on the subject’s practical situation.

In this thesis, my aim is to defend the orthodox account of knowledge against shifty views. My general strategy will be to undermine the main arguments which have been proposed in favour of shifty views, and to show that we can account for most of the data which have been advanced without compromising the orthodox view.

1.2 Four arguments in favour of shifty views

Shifty views are motivated by at least four main considerations. I will now briefly introduce them, but for reasons that will emerge, my thesis will concentrate on three of them only.

A first influential argument which has been proposed in favour of shifty views, but that I will leave aside, is that shifty views could meet the sceptical challenge whereas orthodox views could

not. Scepticism is the view according to which almost all our ordinary ‘knowledge’ ascriptions are false. On this view, we do not in fact know what we ordinarily take ourselves to know, say, that we have hands, that there are cars, etc. Often, the sceptic appeals to the ignorance argument, according to which there always are possibilities of error you supposedly cannot know that they are false (e.g., that you are a brain in a vat, that you are dreaming, that you are hallucinating, etc.). And, it is argued, if you do not know that these possibilities are false, you do not know, say, that you have hands.

Advocates of shifty views have argued that they can solve the sceptical problem, for they can explain the strength of the ignorance argument while avoiding scepticism. Basically, the idea is that when the sceptic mentions a possibility of error, the epistemic standards of the context rise and ‘knowledge’ ascriptions become false.¹

In this thesis, I will leave this argument aside, for it is far from clear that shifty views actually have an advantage over orthodox views regarding the sceptical challenge. First, even if shifty views can provide an interesting explanation of the strength of the ignorance argument in a non-sceptical framework, they do not really answer the deep sceptical problem. They assume that we can know or truly ascribe ‘knowledge’ in some contexts, when the epistemic standards are low. But the sceptic’s line of thought also challenges the claim that low epistemic standards can be met. The fundamental sceptical challenge is whether we can have any justification at all for believing something, and this challenge is not addressed by shifty views.² Secondly, orthodox philosophers have proposed alternative explanations why the ignorance argument is *prima facie* appealing.³

Still, philosophers generally reply to the sceptical challenge, and to the ignorance argument, by granting that knowledge does not require an absolutely maximal epistemic position or justification. But this reply raises a new problem: How are we to determine without arbitrariness how strong one’s epistemic position should be to know? This is the threshold problem.

The threshold problem is supposed to provide another argument in favour of shifty accounts of knowledge. If we assume that practical factors partially determine the threshold of justification that the subject must satisfy to know (or to be truly said ‘knowing’), it is argued, we can determine without arbitrariness this threshold. The practical situation and the rationality of acting on the target proposition given this practical situation would provide us with non-arbitrary epistemic standards for knowledge, or for true ‘knowledge’ ascriptions.

The examination of the threshold problem, and the question whether this problem can be solved by shifty views but not by orthodox views of knowledge, will constitute the first part of my thesis. I will argue that shifty views does not give a satisfactory answer to this problem.

A third kind of argument in favour of the claim that knowledge (or ‘knowledge’) is sensitive

¹See, in particular, Cohen (1988), Cohen (1999), and DeRose (1995) for a contextualist solution. See Stanley (2005b, 129-130) for an anti-intellectualist adaptation of this reply.

²See Klein (2000) and Stanley (2005b, 128-130).

³See, for example, Sosa (2000).

to practical factors invokes our practice of ‘knowledge’ attribution. Consider, for instance, the following cases:⁴

Bank Case A (Low Stakes). My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays. I reply, No, I know it’ll be open. I was just there two weeks ago on Saturday. It’s open until noon.

Bank Case B (High Stakes). My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, Banks do change their hours. Do you know that the bank will be open tomorrow? Remaining as confident as I was before that the bank will be open then, still, I reply, Well, no. I’d better go in and make sure.

In both of these cases Keith believes with the same degree of confidence something true – that the bank is open the following day. His belief is based on the same evidence in the same way – on memory and inductive reasoning. Still, intuitively, Keith’s ‘knowledge’ attribution is appropriate in case A and Keith’s ‘knowledge’ denial is appropriate in case B. What might explain this?

Arguably, the difference has to do with the practical situation. In case B – but not in case A – it is very important for Keith not to be wrong (“they have written a very large cheque”), and he must respond to a salient possibility of error (“Banks do change their hour”). Presumably, our intuitions track this difference. But this is an unexpected fact if the orthodox view of knowledge is right.

These cases suggest that two subjects may satisfy the same epistemic standards with respect to a proposition while not being on a par as far as ‘knowledge’ ascriptions are concerned. It could be that the word ‘know’ does not mean the same thing in both cases. This is the contextualist hypothesis. Or it might be that ‘know’ means the same thing in both cases, but the epistemic standards that one must satisfy for knowledge vary with the practical situation. This is the anti-intellectualist view. In

⁴DeRose (2009).

both of these explanations, the truth-value of a ‘knowledge’ ascription depends on practical factors. But the advocate of the orthodox view must explain these linguistic intuitions differently.

I will devote the second part of my thesis to the consideration of these shifty ‘knowledge’ ascriptions. I will propose a new explanation of these shifty ‘knowledge’ ascriptions which is compatible with orthodox views of knowledge. This explanation appeals in particular to the idea that assertions are governed by a context-sensitive certainty norm.

Finally, a fourth argument which has been proposed in favour of shifty views appeals to a supposed normative connection between knowledge and rational action, or to the idea that knowledge is the epistemic norm of rational action. The very idea of epistemic norm of action is intuitive. When you perform actions, in general, you act on things you take to be true. For instance, if you turn left to go to the restaurant, presumably, you rely on the belief that the restaurant is on the left. When a surgeon uses a scalpel to operate, he assumes that it is safe. When a judge sentences someone, he takes for granted that this person is guilty. Sometimes, we perform actions without considering the truth. For instance, I can turn left to go to the restaurant having no idea where the restaurant is. Also, we sometimes perform actions on the basis of falsehoods. For instance, it may be that the scalpel’s surgeon is in fact unsafe and the person sentenced not guilty. In such cases, often, we can be criticised for what we do. My partner can appropriately blame me by saying: “Why on earth did you turn left if you had no idea where the restaurant is?”. If the surgeon does not want to infect the patient and the judge does not want to sentence an innocent person, and they have no reason, respectively, to think that the scalpel is safe and the person guilty, they are clearly irrational. If they are wrong as to whether the scalpel is safe and the person guilty, they may have an excuse. Perhaps, given the situation, no one could be expected to have right beliefs. But the fact that actions are assessed in these ways suggests that there is an epistemic norm governing action.

In brief, when evaluating someone’s action, we consider whether this action was performed for good reasons. And this requires that the action was properly motivated. To see whether the action is properly motivated, we do not only look at the result of the action, or whether the action conforms to the norms of morality, prudence, legality, etc. We also consider the agent’s epistemic position with respect to the relevant facts, and whether a relevant epistemic requirement is satisfied.

Now, it seems that the relevant epistemic requirement for action varies with the practical context. The argument for this relies in particular on normative intuitions. Let us return to the bank cases. Most of us are, I suppose, inclined to say that Keith should not come back without checking whether the bank has changed its hours in case B, while it is permissible for him to come back without checking in case A. This is easily explained if the epistemic requirement on rational action varies with the stakes. In case A, Keith satisfies the epistemic condition, which is why he need not check, while in case B, he does not satisfy this epistemic condition, which is why he need to check before acting. This variation can be explained if the norm is context-sensitive.

But suppose that knowledge is the epistemic condition that we must satisfy for rational action. If this epistemic condition varies with the practical situation, then knowledge varies with the practical situation, and the orthodox account of knowledge is wrong. Consider the bank cases. According to the orthodox view, Keith knows in case A if and only if he knows in case B. Accordingly either the orthodox philosopher must deny that Keith knows in case A, or she must say that Keith knows in case B. If she denies that Keith knows in case A, then scepticism threatens. We do not know many things we ordinarily take ourselves to know. If she says that Keith knows in case B, then it seems that she is forced to deny that knowledge is the epistemic norm of practical rationality, for although Keith is supposed to know that the bank is open, he must still check before acting.

The third part of my thesis will be devoted to this issue. I will argue that knowledge is always necessary for rational action, but not always sufficient. I will then defend a context-sensitive certainty norm for rational action. By the rejection of the claim that knowledge is always sufficient for rational action, the fourth argument in favour of shifty views of knowledge will be undermined.

1.3 Summary of the chapters

Chapter 2 Intellectualism, pragmatism and pragmatic encroachment

In this chapter, I present the intellectualist approach underlying the orthodox account of knowledge, and its main motivation, namely, that there is a fundamental connection between belief and truth which is normative. I consider the idea according to which practical factors are also relevant to the justification of our beliefs, and I distinguish the notion of justification of belief and that of *epistemic* justification of belief. On this basis, I distinguish the thesis according to which there are practical reasons to believe and the pragmatic encroachment thesis underlying anti-intellectualism. I explain that pragmatism about reasons to believe is in principle compatible with intellectualism about *epistemic* justification of belief and the orthodox view of knowledge. Finally, I distinguish pragmatism and pragmatic encroachment from the thesis according to which there are practical reasons to *form* beliefs. I argue that the thesis that there are practical reasons to *form* beliefs is plausible and that there are no categorical epistemic obligations to *form* beliefs. However, I stress that this does not impugn the orthodox view of knowledge.

Chapter 3 The threshold problem

In this chapter, I consider the claim according to which embracing pragmatic encroachment on knowledge – a shifty view of knowledge – provides us with a solution to the threshold problem. I explain that appealing to practical factors does not solve the threshold problem, for either the practical condition which is supposed to be met for epistemic justification and knowledge is too weak, or a similar threshold problem arises for when this practical condition is met. I conclude

that the threshold problem does not favour pragmatic encroachment over orthodox accounts of knowledge and epistemic justification to believe.

Chapter 4 Linguistic intuitions and the rule of relevance

In this chapter, I start considering our shifty ‘knowledge’ attributions. I discuss a popular explanation of these shifty ‘knowledge’ attributions which is compatible with an orthodox view of knowledge. This explanation appeals to the notion of pragmatic appropriateness and to Grice’s conversational rule of relevance. I raise a certain number of objections to this approach. In particular, I show that it is *ad hoc* and that it cannot explain two significant conversational facts. I conclude that this approach is insufficient to explain our shifty ‘knowledge’ attributions in an orthodox framework.

Chapter 5 Linguistic intuitions and belief

In this chapter, I consider another way in which an orthodox philosopher might want to explain our shifty ‘knowledge’ attributions. This approach aims at explaining these variations by invoking intuitions concerning whether the subject or the ascriber (justifiably) believes the target proposition. I consider two ways of fleshing out this view. The first one has it that practical factors may raise our threshold for confidently believing because we are prone to confuse the epistemic standards that we must satisfy to be in a good enough epistemic position to act on p with the epistemic standards that we must satisfy to be in a good enough epistemic position to believe p . I raise two objections against this view. The second way of fleshing out this approach relies on the claim that practical factors affect our belief-formation process and our willingness to *form* a belief. While I agree that practical factors can affect our willingness to *form* beliefs, I put forth four objections showing that this approach is insufficient to explain all the cases.

Chap. 6 Linguistic intuitions and the certainty norm of assertion

In this chapter, I defend a mixed account regarding the variation of ‘knowledge’ ascriptions and denials. This account appeals to the claim that practical factors affect our willingness to *form* a belief and to the claim that certainty is the norm of assertion. I provide arguments for the idea that assertions are governed by a context-sensitive certainty norm, and I respond to some objections which have been raised against this view. Next, I show that if we adopt this certainty account of assertion, we can propose a mixed account of our shifting intuitions concerning the appropriateness of ‘knowledge’ ascriptions. As this mixed approach is compatible with the orthodox account of knowledge, I conclude that our shifty ‘knowledge’ ascriptions can be explained in an orthodox framework.

Chapter 7 The epistemic evaluation of actions

In this chapter, I lay the groundwork for discussing the argument against the orthodox account of knowledge from the supposed normative connection between knowledge and action. I start by presenting the idea of evaluation of action in general. I identify three levels of evaluation of action, two of which involving a consideration of the agent's epistemic position. I show how these levels of evaluation may be related to the notions of justification, rationality and excuse. Next, I consider the notion of epistemic evaluation of action. I show that this notion is equivocal because there are different things that can be evaluated by considering the subject's epistemic position. Finally, I present the main accounts of the epistemic norm of action and I identify the different ways in which the orthodox philosopher might wish to reply to the argument from the supposed normative connection between knowledge and action.

Chapter 8 The necessity of knowledge for action, and the warrant account

Often, orthodox philosophers deny that there is a normative connection between knowledge and action. They endorse a warrant account of action. In this chapter, I show that the warrant account is not a fully satisfactory alternative to the knowledge account, and I defend the claim that knowledge is necessary for rational action. In particular, against recent criticisms, I defend the appeal to the distinction between excuse and justification to account for cases in which knowledge does not seem necessary. In the course of doing it, I also stress that the warrant account does not do justice to our intuitions underlying the principled argument for the necessity of knowledge and the argument from bad reasonings. I conclude that we should adopt the claim that knowledge is always necessary for rational action. Although this does not threaten the orthodox account of knowledge, this leads to reject the warrant account.

Chapter 9 The context-sensitivity of the norm of action

In this chapter, I argue that the epistemic norm of action is sensitive to the agent's practical context, and hence, that an orthodox epistemologist should reject the claim that knowledge is the epistemic norm of action. I consider, but reject, two attempts to explain the cases in which our epistemic evaluations of actions shift with the practical context in a framework in which the epistemic norm of action does not shift. These attempts are partially based on the claim that there are epistemic assessments which do not bear on whether the epistemic norm of action is satisfied. While I agree with that claim, I stress that, given the natural way of understanding these cases, conceding this point is not sufficient to reject the context-sensitivity of the epistemic norm of action. Finally, I offer a second argument in favour of the sensitivity of the epistemic norm of action. According to this argument, insensitivism about the epistemic norm of action leads to endorse a maximal-standard account of the epistemic norm of action, which is an implausible account.

Chapter 10 The insufficiency of knowledge for action, and the certainty account

In the last chapter, I argue, first, against the sufficiency of knowledge on grounds independent from the denial of pragmatic encroachment. Second, I defend a context-sensitive certainty norm for action. To argue against the sufficiency of knowledge, I consider cases in which more than knowledge is intuitively required, and I argue that advocates of the sufficiency of knowledge offer no convincing explanation. Next, I reject an argument for the sufficiency of knowledge based on the idea that if you know that p , then not- p is epistemically impossible for you. I reject the main arguments in favour of this kind of infallibilism about knowledge, and hence, the main argument in favour of the sufficiency of knowledge. To defend the certainty norm for action, I argue that this account satisfies all the desiderata that an account of the epistemic norm of action should satisfy. I emphasize that a certainty norm of action does not threaten the orthodox view of knowledge, which thereby undermines the third argument against the orthodox view of knowledge.

Part I

Justification of belief and practical factors

Chapter 2

Intellectualism, pragmatism and pragmatic encroachment

2.1 Introduction

Orthodox epistemologists are intellectualist. According to them, whether a belief is justified, or amounts to knowledge, is fully determined by truth-relevant factors only. These factors are considerations that connect in some way the subject's belief to the truth. In this chapter, I present the intellectualist approach and its main motivation, namely, that there is a fundamental connection between belief and truth which is normative. Next, I consider the idea according to which practical factors are also relevant to the justification of our beliefs. I distinguish the notion of justification of belief and that of *epistemic* justification of belief, and on this basis, I distinguish two theses: first, there is the pragmatist thesis according to which practical considerations are relevant to justification of our beliefs because, besides epistemic reasons to believe, there are practical reasons to believe; second, there is the pragmatic encroachment thesis, according to which practical considerations are relevant to *epistemic* justification of our beliefs, because practical considerations determine whether an epistemic reason is a *sufficient* epistemic reason to believe. I explain that pragmatism about reasons to believe is in principle compatible with intellectualism about *epistemic* justification of belief and knowledge (the orthodox view), and that pragmatic encroachment is compatible with the denial of pragmatism thus understood.

Finally, I distinguish pragmatism and pragmatic encroachment from a third thesis. This is the thesis according to which practical factors are relevant to (epistemic) justification of *forming* beliefs. I show that this third thesis is plausible because, I argue, there are no categorical epistemic obligations to form beliefs. This is so either because there are no epistemic obligations to believe

anything, or because these epistemic obligations are conditional on the subject's practical situation. Even if we grant that practical factors affect the (epistemic) obligation of forming beliefs, I stress, this does not impugn intellectualism about epistemic justification of belief and about knowledge. I conclude that the first and third theses do not constitute threats for the orthodox philosopher.

2.2 Truth-relevant factors, intellectualism and the norm of belief

According to intellectualism – the orthodox view about knowledge and justification of belief in epistemology – only truth-relevant factors are relevant to whether a belief is justified, or amounts to knowledge. Truth-relevant factors are factors that connect in some way the belief to the truth, for example, the evidence, the reliability of the belief-forming process, and so on. In other words, these are considerations which affect the likelihood that the belief is true.¹

Intellectualism is the default view. As Reed (2014, 95) writes:

Traditionally, this point of agreement [i.e. intellectualism] was so widely and deeply shared that epistemologists never really thought about it. It has been given a name . . . only recently and only by the relatively small number of philosophers who have argued against it.

Opponents to intellectualism agree. They often note that the denial of intellectualism (sometimes called 'impurism') is "surprising".² And they sometimes go as far as conceding that the denial of intellectualism seems "mad" or "very hard to swallow".³

Intellectualism can be motivated by noting that there is a normative relation between belief and truth. You have a belief that *p* when you have a certain psychological attitude toward a certain proposition (or content) *p*. This propositional attitude can be distinguished from others – such as accepting that *p*, desiring that *p*, hoping that *p*, etc. – on the basis of the fact that when you enjoy this attitude toward the content in question, you take this content, in a distinctive way, to be true. For example, when you desire, imagine, or hope that it is raining, you do not take the proposition that it is raining as true, or when you accept a proposition as true, you do not take it as true whatever the context, but only for the purposes at hand.

There is also a connection between truth and other non-doxastic attitudes. For example, to desire that *p* is to desire that *p* is true, and to imagine that *p* is to imagine that *p* is true. So, it is useful to say more on how we should understand the connection between belief and truth.

¹ According to DeRose (2009, 24) a factor is truth-relevant if and only if "it affects how likely it is that a belief is true, either from the point of view of the subject or from a more objective vantage point". Stanley (2005b, 1) refers to "truth-conducive" factors. A factor is truth-conducive if and only if "its existence makes our belief more likely to be true objectively or from the subject's point of view".

² See, e.g., Fantl and McGrath (2014, 88) and Sripada and Stanley (2012, 4).

³ See, respectively, Fantl and McGrath (2009, 27-28, 66) and Fantl and McGrath (2011, 560).

Intuitively, if you believe that p and p is false, your belief is wrong or incorrect. In contrast, if you desire, hope, or accept, that p and p is false, this does not imply that there is something wrong in your desire, hope, or acceptance. Consider the belief-forming process. If you believe that p , but realise that p is false, normally, you change your belief. This is not the case, for example, with desires and acceptances. The belief-forming process is affected by the question whether p . In particular, an answer to the question whether to believe that p is settled by an answer to the question whether p .⁴ For example, if you ask yourself whether you should believe that it is raining, presumably, you will ask yourself whether it is raining. A positive answer to the question “Is it raining?” provides a positive answer to the question “Should I believe that it is raining?”. This fact is easily explained if truth is the correctness condition of beliefs, and if the relation between belief and truth is hereby normative: Whether or not you ought to believe that p depends on whether p is true or false.⁵

This special relation between belief and truth is expressed by many philosophers by the claim that belief aims at truth.⁶ In particular, the idea that there is a normative connection between belief and truth is often *explained* by the idea that belief has a constitutive and individuating property of truth-directedness. This property is supposed to be constitutive of belief in the sense that a mental state lacking this property could not be a belief (i.e., this property is essential to belief), and this property is supposed to be individuating in that it would be sufficient for distinguishing belief from the other types of mental states (guesses, imaginings, desires, hopes, etc.).

There are different ways of fleshing out the idea that belief aims at truth.⁷ To illustrate, consider Velleman’s account. Velleman (2000) argues that the claim that belief aims at truth is a descriptive claim. On this view, a belief is a mental state of accepting (or regarding as true) a proposition that it is regulated by the intention (or the sub-intentional aim) of getting the truth-value of the proposition right. This distinguishes belief from other mental states of acceptance (such as assuming or imagining) where the aim is not necessarily to get the truth-value of the proposition right. Now, suppose that there is such an aim of belief. Only a true belief is a belief which has reached its goal. This explains why we can say that a true belief is right and a false belief is wrong. This also explains why only truth-relevant factors are relevant to whether you ought to believe that p , that is, on this teleological approach, relevant to the question whether such a belief would reach its goal.

A different view is developed by Shah (2003). Shah argues that the claim that belief aims at truth must be understood metaphorically as the claim that our *concept* of belief is the concept of a mental state that is correct if and only if its content is true. This normativist account can explain two facts which, taken together, are at odd with Velleman’s descriptive and teleological approach. First, wishful thinkings are mental states that are not entirely formed through mechanisms aiming

⁴See Shah (2003).

⁵For discussion, see Engel (2004).

⁶The expression ‘aim of belief’ goes back to Williams (1973).

⁷See Fassio (2015) for an overview.

at truth. For example, if you think of yourself as a good philosopher because the thought that you are not a good philosopher is depressing, then this thought is not formed with the aim of having the truth. Still, such thoughts can count as genuine beliefs. So, the teleologist should accept that truth is not the only factor regulating beliefs. Sometimes, she should accept, beliefs are formed with the aim of releasing the pain. But if so, this teleologist cannot explain a second fact, called ‘transparency’. According to transparency, doxastic deliberation – reasoning guided by the question whether to believe that p – is explicitly regulated *only* by considerations relevant to the question whether p is true. For example, suppose you ask yourself whether you should believe that you are a good philosopher. The consideration that believing that you are a good philosopher releases pain is impotent. You cannot form the belief that you are a good philosopher on this basis. Indeed, this consideration does not indicate that such a belief is true. Now, if the teleologist wants to explain transparency, she must deny that beliefs can be formed with an aim different from truth. For, otherwise, how can she explain that we cannot form beliefs explicitly on the basis of truth-irrelevant factors? But if she denies that, she must deny that wishful thinkings are beliefs.

In contrast, according to Shah’s account, doxastic deliberation is just one kind of belief-forming process. Doxastic deliberation is a belief-forming process in which the concept of belief is employed by the subject. If the concept of belief is the concept of a mental state that is correct if and only if its content is true, we can explain why – in accordance with transparency – *only* the factors relevant to whether p is true are considered as relevant in this kind of belief-forming process. Indeed, only considerations indicating that p is true can move the subject to believe that p because only these considerations can show that the belief that p is correct.

There are other possible accounts of the aim of beliefs and of the special connection between belief and truth.⁸ But the thought that a belief is correct if and only if it is true underwrites an intellectualist account of the norm of belief. On this approach, the question whether there is a justification for a certain belief amounts to the question whether there are considerations that sufficiently indicate that this belief is right, that is, true.

Intellectualist philosophers have derived different types of norms for beliefs from this fundamental normative connection between belief and truth, depending in particular on the kind of truth-relevant factors that they think are relevant to the justification of our beliefs (the evidence, the counterfactual relations between the belief and the truth, the reliability of the belief-forming process, etc.).

If we call ‘ S ’s epistemic position with respect to p ’ S ’s profile with respect to all the factors relevant to the truth of p , we can state intellectualism about justified belief as follows:

⁸Shah and Velleman (2005) defend a mixed view about the aim of belief, according to which the concept of belief is the concept of a mental state that is (at least in part) regulated by the aim of truth – mental states entirely insensitive to evidence do not seem to count as genuine beliefs – and that is correct if and only if it is true – only truth-relevant considerations count in doxastic deliberation. Engel (2004) defends a normativist account of the aim of belief and the view that knowledge, and not truth, is the aim of belief.

(Intellectualism-JB) If S1 and S2 are in the same epistemic position with respect to p, S1 is justified in believing that p if and only if S2 is justified in believing that p.

My aim is to defend the orthodox view of knowledge in general, and hence, intellectualism about *knowledge*. In the remaining of this chapter, however, and in the following, I shall consider a line of argument which has been raised explicitly against a specific version of (Intellectualism-JB), namely, evidentialism – defined as the following thesis:

(Evidentialism-JB) If S1 and S2 have the same evidence with respect to p, S1 is justified in believing that p if and only if S2 is justified in believing that p.⁹

According to (Evidentialism-JB), evidence alone is relevant to the justification of our beliefs. Other truth-relevant factors, such as the reliability of the belief-forming process, are irrelevant to justification of belief. Now, according to the argument proposed against (Evidentialism-JB) that I shall consider, truth-irrelevant factors are essential to the justification of our beliefs. If this argument is correct, I shall explain, we should reject not only (Evidentialism-JB), but also (Intellectualism-JB).

If (Intellectualism-JB) is false, we might think, intellectualism in general is false. However, before considering the argument against (Evidentialism-JB), I will show that we must distinguish intellectualism about justification of belief (Intellectualism-JB) and intellectualism about *epistemic* justification of belief. Hence, we should also distinguish intellectualism about justification of belief and intellectualism about knowledge. In sum, we can in principle reject (Intellectualism-JB) and (Evidentialism-JB) while maintaining intellectualism about *epistemic* justification of belief and knowledge.

2.3 Practical factors, pragmatism and pragmatic encroachment

2.3.1 Intellectualism, pragmatism and pragmatic encroachment

According to intellectualist philosophers embracing (Intellectualism-JB), the only factors that are relevant to whether our beliefs are justified are truth-relevant factors. However, we might think that even if truth-relevant factors are relevant to the justification of our beliefs, they are not the *only* relevant factors. There are two ways of denying (Intellectualism-JB).

First, following pragmatism about reasons to believe, we might think that there are practical reasons to believe. According to Pascal, for example, you have a reason to believe in God if having this belief has the highest expected utility, whatever the strength of your evidence in favour or against this belief.¹⁰ According to William James, you can have a reason to believe p, even if the

⁹See Conee and Feldman (2004).

¹⁰See Pascal (1670, §233).

evidence is insufficient to show that p is true, provided that believing p is beneficial.¹¹ If there are practical reasons to believe, (Intellectualism-JB) is false. S1 and S2 can have the same epistemic position with respect to p even if S1's belief is justified and S2's belief is unjustified, because S1, but not S2, has a practical reason justifying her belief.

As such, pragmatism about reasons to believe does not threaten the intellectualist account of *epistemic* justification of belief and of knowledge, though. On this pragmatist view, the distinction between practical reasons to believe – practical considerations favouring a certain belief – and epistemic reasons to believe – truth-relevant considerations connecting the belief to the truth – can be preserved.¹² As a result, the distinction between pragmatic and epistemic justification of beliefs can in principle be preserved.

If it is granted that knowledge and epistemic justification of belief do not require pragmatic justification of belief, this version of pragmatism is compatible with the following claims:

(Intellectualism-EJ) If S1 and S2 are in the same epistemic position with respect to p , S1 is epistemically justified in believing p if and only if S2 is epistemically justified in believing p .

(Intellectualism-K) If S1 and S2 are in the same epistemic position with respect to p , S1 has knowledge-level justification for p if and only if S2 has knowledge-level justification for p .

There are reasons to reject pragmatism about reasons to believe. Practical reasons seem to be the wrong kind of reasons for beliefs. These are reasons for action and, arguably, beliefs are not actions. We cannot deny that there are obviously false beliefs which are more useful than others, and that it is in some sense better to have them. But if we reject pragmatism about reasons to believe, we can account for this by saying that even if there are practical reasons to form these beliefs, these reasons are not reasons to believe.¹³ If it is beneficial to believe that p but nothing shows that p is true, you have a reason to form the belief that p , but this does not imply that you have a reason to believe that p . You just have a reason to form an irrational belief.

The second way of denying (Intellectualism-JB) is to deny (Intellectualism-EJ) and (Intellectualism-K). Thus, against the claim that truth-relevant factors alone are relevant to the justification of our beliefs, it has been argued more recently that these factors are insufficient to determine whether a certain belief is *epistemically* justified and amounts to knowledge. This is the thesis of pragmatic encroachment.¹⁴ It is important to note that this thesis is completely independent from pragmatism

¹¹See James (1897).

¹²A stronger form of pragmatism, defended by James, identifies truth and utility. This form of pragmatism cannot distinguish between epistemic and practical reasons to believe.

¹³See Thomson (2008); Skorupski (2010); Gibbons (2013).

¹⁴See Owens (2000); Hawthorne (2004); Stanley (2005b); Fantl and McGrath (2002, 2007, 2009). The expression 'pragmatic encroachment' has been forged by Kvanvig on his blog Certain Doubts.

about reasons to believe.¹⁵ Pragmatism is a thesis directly about the falsity of (Intellectualism-JB), whereas pragmatic encroachment is primarily a thesis about the falsity of (Intellectualism-EJ) and (Intellectualism-K).

As my aim is to defend the orthodox view of knowledge, and hence, (Intellectualism-K), I shall put aside pragmatism about reasons to believe. I shall focus on the pragmatic encroachment thesis. I shall consider an argument that has been proposed, with the aim of arguing against the idea that we have reflective control over our beliefs, against (Evidentialism-JB), and hence, against (Intellectualism-JB). As made clear above, not every argument against (Intellectualism-JB) is an argument against (Intellectualism-EJ) and (Intellectualism-K). Again, consider arguments in favour of pragmatism about reasons to believe. If practical reasons to believe do not affect epistemic justification of belief and knowledge, then, pragmatism about reasons to believe, while incompatible with (Intellectualism-JB), is compatible with (Intellectualism-EJ) and (Intellectualism-K).

Still, the argument that I shall consider in this chapter and the following targets (Evidentialism-JB) by attacking the evidentialist versions of (Intellectualism-K) and (Intellectualism-EJ):

(Evidentialism-EJ) If S1 and S2 have the same evidence for p, S1 is *epistemically* justified in believing p if and only if S2 is *epistemically* justified in believing p.

(Evidentialism-K) If S1 and S2 have the same evidence for p, S1 has knowledge-level justification for p if and only if S2 has knowledge-level justification for p.

According to this argument, when the evidence is inconclusive, that is, when the evidence is compatible with the falsity of the belief, pragmatic considerations are required to justify a belief *epistemically speaking*, and to make it a good candidate for knowledge. If we reject pragmatism, we assume that there is no difference between *epistemic* justification of belief and justification of belief. Thus, according to this argument, even if we reject pragmatism, pragmatic considerations are required to justify a belief. I will call this argument the ‘inconclusive evidence argument’ and the problem that it raises for the epistemic intellectualist the ‘threshold problem’.

To clarify the meaning of the inconclusive evidence argument, and the related threshold problem, it will be important to distinguish two things that we might think are normatively affected by practical factors. First, we might think that practical factors affect the epistemic justification of the *formation* of beliefs. Second, we might think that practical factors affect the epistemic justification of the belief. Drawing on this distinction, I shall reconstruct in the remaining of this chapter two versions of the inconclusive evidence argument. I shall show that intellectualism about epistemic justification of belief and knowledge is not threatened by reasons to think that practical factors affect the (epistemic) justification of the formation of beliefs. I shall consider the version of the argument according to which practical factors affect the epistemic justification of belief in the next chapter.

¹⁵See Marusic (2011) for an overview.

2.3.2 Pragmatic encroachment, reflective control and inconclusive evidence

In *Reason without freedom*, Owens argues against the claim that we have reflective control over our beliefs. To have reflective control over our beliefs is to have the capacity to motivate a belief that *p* by a judgement that we have sufficient reason to believe that *p*, or a judgement that we ought to believe that *p*. To argue against the idea that we have reflective control over our beliefs, Owens relies on the claim that the considerations which are relevant to whether we ought to believe that *p* cannot, on reflection, influence our motivation to believe that *p*. He concludes that a mere reflective judgement that we ought to believe that *p* is impotent to motivate a belief in *p*, and hence, that we do not have reflective control over our belief.

What are, according to Owens, the considerations which are relevant to whether we ought to believe that *p*? Firstly, Owens refers to a traditional truth-relevant factor: the evidence. Basically, *E* is evidence for *p* if *E* indicates that *p* is true. Owens endorses an internalist account of justified beliefs. Hence, he assumes that *unpossessed* evidence – evidence that the subject does not have – is irrelevant to whether the subject ought to believe that *p*. Now, can reflection on her evidence motivate a rational subject to form a belief? According to Owens, if the evidence is not conclusive, that is, if it is compatible with the falsity of the proposition, then reflection on this evidence cannot motivate the subject to believe. Thus, he writes:

To believe that *p*, I must be under the impression that I have a conclusive reason to think *p* true: reflection on inconclusive evidence is quite inadequate. (Owens, 2000, 35)

Still, as he notes, it is clear that we are justified in believing many things which are not supported by conclusive evidence. For example, my belief that I have hands seems justified. Still, I do not have conclusive evidence for this belief. I might be a brain in a vat deceived so that it merely seems to me that I have hands. If a subject can be justified in believing that *p* on the basis of inconclusive evidence, but she cannot be moved to form a belief by a reflective judgement on her justifying but inconclusive evidence, she does not have reflective control over her beliefs.

Now, it must be said under which conditions a rational subject is motivated to believe that *p* on the basis of inconclusive evidence that *p*. Thus, secondly, Owens appeals to non-evidential factors, or “pragmatic consideration”:

evidence alone cannot motivate the formation of a rational belief. When we assess how sensible a belief would be, we must ask ourselves whether there is sufficient evidence to justify that belief, and it is impossible to determine how much evidence we should require to convince us of the point without considering the importance of the issue, the cognitive resources available to resolve it, and so forth. These considerations are non-evidential and yet essential to the justification of belief. Therefore evidentialism must be rejected. (Owens, 2000, 23)

Or, in other terms,

purely evidential considerations under determine what we ought to believe until they receive pragmatic supplementation. (Owens, 2000, 26-27)

We might think that the (alleged) relevance of pragmatic factors to the justification of our beliefs can explain why a mere reflective judgement on inconclusive evidence is impotent to motivate a subject to believe, without compromising the idea that we have reflective control over our beliefs. If pragmatic considerations are necessary to determine what we ought to believe, then a reflective judgement on mere inconclusive evidence is bound to be motivationally inert: this judgement under determines the judgement about what we ought to believe.

However, Owens notes that a reflective judgement on pragmatic considerations is also impotent to motivate a belief:

the whole idea of reflective control over belief is under threat once we acknowledge that such non-evidential considerations are needed to motivate rational belief. I can get myself to make a decision or perform an action by reflecting that time is running out, but I can't convince myself of *p*, even if *p* is favoured by the evidence, simply by judging that the time for a decision has arrived. Though pragmatic considerations are needed to rationalise belief, our rationality alone won't ensure that reflection on them can compel belief. Belief cannot be justified without considering pragmatic constraints, but reflection on these constraints is not a source of rational motivation for belief. It looks as though (...) the power of reflection will not secure us control over belief once we allow that non-evidential factors are needed to motivate belief. (Owens, 2000, 23-24)

In sum, Owens defends the following claims:

- (1) Inconclusive evidence alone cannot motivate the formation of a rational belief.
- (2) A reflective judgement about inconclusive evidence cannot move one to form a belief.
- (3) Pragmatic considerations are necessary to determine what we ought to believe.
- (4) A reflective judgement on pragmatic considerations (even in addition to a reflective judgement about inconclusive evidence) cannot move one to form a belief.

It is useful to consider briefly what these claims are supposed to show and how they interact.

First, Owens aims at arguing against evidentialism, understood as “the thesis that what justifies belief in *p* is just evidence in *p*'s favour[, nothing else is relevant to belief justification” (Owens, 2000, 23), or the thesis that “evidential considerations alone are relevant to the fixation of rational belief” (Owens, 2000, 24). So Owens argues against what I have called (Evidentialism-JB):

(Evidentialism-JB) If *S*₁ and *S*₂ have the same evidence with respect to *p*, *S*₁ is justified in believing that *p* if and only if *S*₂ is justified in believing that *p*.

Evidentialism is often defined, in opposition to pragmatism, as the claim that only evidential considerations are reasons to believe. In contrast, pragmatism has it that a belief can be justified by practical considerations, for example, by the fact that believing would be useful. In arguing against what he calls ‘evidentialism’, however, Owens does not argue in favour of pragmatism. He agrees that practical considerations cannot be reasons to believe. Nevertheless, he maintains that practical considerations are essential to the justification of our beliefs, but not because they are reasons to believe. Rather, according to him, they are essential because they determine whether a reason to believe (i.e., a piece of evidence) is a *sufficient* reason to believe. In other words, pragmatic considerations would be essential to make a reason to believe a reason sufficiently strong to *justify* the belief *from an epistemic point of view*. Thus, Owens argues in fact against evidentialism thus understood:

(Evidentialism-EJ) If S1 and S2 have the same evidence for p, S1 is *epistemically* justified in believing p if and only if S2 is *epistemically* justified in believing p.

(Evidentialism-K) If S1 and S2 have the same evidence for p, S1 has knowledge-level justification for p if and only if S2 has knowledge-level justification for p.

Given that Owens rejects pragmatism, he rejects the claim that we should distinguish between epistemic justification of belief and justification of belief. Hence, as he rejects (Evidentialism-EJ), he rejects (Evidentialism-JB). Owens also thinks that the only kind of truth-relevant factors which are relevant to justification of belief is the subject’s possessed evidence. Hence, since he rejects (Evidentialism-JB), he rejects (Intellectualism-JB).

In brief, if we define evidentialism in opposition to pragmatism, as the claim that there are no practical reasons to believe, Owens’ view is still evidentialist. But if, following Owens, we define evidentialism as the claim that evidential considerations alone are essential to the justification of our beliefs, Owens’ view is not evidentialist.

Now, assuming that inconclusive evidence can justify a belief, Owens argues against evidentialism (understood in terms of (Evidentialism-JB) and (Evidentialism-EJ)) by arguing in favour of claims (1) and (3) above:

my point against evidentialism can be made simply by asking: how are you going to tell us, in purely evidential terms, what level of evidence is needed to justify belief? Unless this question can be answered, evidentialism (internalist and externalist) must be abandoned. (Owens, 2000, 26)

Second, Owens aims at arguing against the idea that we have reflective control over our beliefs. Indeed, if (3) and (4) are true, then we do not have reflective control over our beliefs. Note, however, that even if we reject (3) and (1) by endorsing evidentialism, a similar argument can be run, on the basis of (2), against the idea that we have reflective control over our beliefs. Indeed, if inconclusive

evidence can justify a belief but a reflective judgement on inconclusive evidence cannot move one to form a belief, then a reflective judgement that we ought to believe p is motivationally inert. We do not have reflective control over our beliefs.

Finally, note that the argument in favour of (3) must be independent of (2). Indeed, if we reject the claim that we have reflective control over our beliefs, evidentialism is compatible with (2) and (2) does not support (3). On the assumption that we do not have reflective control over our beliefs, we cannot support the claim that pragmatic factors are required to justify a belief (when the evidence is inconclusive) by noting that reflection on inconclusive evidence cannot move one to form a belief. On the other hand, if we embrace the claim that we have reflective control over our beliefs, we must grant that evidentialism is incompatible with (2). Still, given (4), the falsity of evidentialism supports (3) in no way. Given that reflection on pragmatic factors does not move one to form a belief either, if we have reflective control over our beliefs, we cannot think that reflection on inconclusive evidence does not move one to form a belief *because* pragmatic factors are required for this inconclusive evidence to justify a belief.

My aim is not to decide between the advocate of reflective control and her opponent, but merely to defend intellectualism about knowledge. So, I am interested in whether (3) is true. Given that there is no clear argument in favour of (3) from (2), I will set (2) aside (although, I am inclined to think that (2) is false). Still, (3) may be partially supported by (1). It is not implausible to think that inconclusive evidence alone cannot motivate the formation of a rational belief because pragmatic considerations are essential to the justification of our beliefs. But if (3) is true, it seems that the orthodox account of knowledge, is wrong. Knowledge and epistemic justification to believe are not merely a matter of truth-relevant factors.

The underlying reasoning seems to be the following:

ARGUMENT FROM INCONCLUSIVE EVIDENCE

1. S can have inconclusive evidence that p and a rational belief that p .
2. Inconclusive evidence alone cannot motivate the formation of a rational belief.
3. If inconclusive evidence alone cannot motivate the formation of a rational belief, pragmatic considerations are needed to rationalise belief.
4. If pragmatic considerations are needed to rationalise belief, then evidentialism is false.
5. Therefore, evidentialism is false.

This argument is somewhat unclear. We might want to distinguish between rationalising the formation of a belief and rationalising a belief. Thus, the argument could be seen as focused on the belief-forming process:

ARGUMENT FROM INCONCLUSIVE EVIDENCE (FORMATION)

1. S can have inconclusive evidence that p and have rationally formed a (rational) belief that p.
2. Inconclusive evidence alone cannot rationalise the formation of a (rational) belief.
3. If inconclusive evidence alone cannot rationalise the formation of a (rational) belief, pragmatic considerations are needed to rationalise the formation of a (rational) belief.
4. If pragmatic considerations are needed to rationalise the formation of a (rational) belief, then evidentialism is false.
5. Therefore, intellectualism (and hence, evidentialism) is false.

Alternatively, the argument could be seen as focused on the result of the belief-forming process:

ARGUMENT FROM INCONCLUSIVE EVIDENCE (RESULT)

1. S can have inconclusive evidence that p and a rational belief that p.
2. Inconclusive evidence alone cannot rationalise a belief.
3. If inconclusive evidence alone cannot rationalise a belief, pragmatic considerations are needed to rationalise belief.
4. If pragmatic considerations are needed to rationalise belief, then evidentialism is false.
5. Therefore, evidentialism is false.

In the remaining of this chapter, I shall focus on ARGUMENT FROM INCONCLUSIVE EVIDENCE (FORMATION). I shall argue that this argument fails because its fourth premise is wrong: evidentialism is not a claim about what rationalises the formation of a (rational) belief, but a claim about what can rationalise (or justify) a belief.

In contrast, ARGUMENT FROM INCONCLUSIVE EVIDENCE (RESULT) can be seen as a more serious argument against intellectualism. I shall devote the next chapter to this argument.

2.4 Practical factors, intellectualism and belief formation

Is the claim that practical considerations are relevant, and even required, to rationalise the formation of a belief incompatible with evidentialism, and, more generally, with intellectualism? It does not seem so.

As Owens notes, we may distinguish two issues:

that of when we ought to engage in theoretical inquiry (i.e. collect, deliberate about and retain evidence) and, second, that of whether we ought to form a belief on the basis of the evidence before us. The evidentialist can happily concede that practical considerations have a bearing on the former issue; the collection, assimilation and retention of evidence is an activity which we can choose to perform or not as we please. Inquiry and deliberation are rationally motivated by practical considerations. Nevertheless, the evidentialist insists, whether a belief gets formed or not once deliberation begins is (or should be) determined entirely by the balance of the evidence considered. Evidence is the only thing that should motivate the formation of a rational belief. (Owens, 2000, 27)

In this passage, Owens distinguishes the question whether you ought to inquire about p and the question whether, given the evidence that you have, you should believe that p . To illustrate, consider someone who must collect all the data about p . Of course, the mere fact that she must collect more evidence about p does not show that she ought not to believe that p . Perhaps she already has sufficient evidence to believe that p . Similarly, the mere fact that she must stop collecting evidence about p does not show that she ought to believe that p . Perhaps she still has no sufficient evidence to believe that p . Given that the practical reason to inquire or to stop inquiring about p has no direct bearing on whether we ought to believe that p , epistemic intellectualism and evidentialism are not compromised by the idea that practical factors can be relevant to whether we should inquire further or stop the inquiry.

Still, besides distinguishing a reason to inquire and a reason to believe, we must also distinguish between a reason to believe and a reason to form a (rational) belief. Is evidentialism committed to the claim that “evidence is the only thing that should motivate the formation of a rational belief”? This is far from clear.

We may distinguish three possible cases. First, you may have a practical reason to have, and thereby form, a *rational* belief as to whether p . In this case, if evidentialism is correct, the only way for you to form this belief is to look at the evidence you have in favour and against p . Second, you may have a practical reason to form a belief as to whether p , be it rational or not. In this case, you may rationally form a belief by rationally taking a pill which causes the belief that p or the belief that not- p , or you may look at your evidence. Suppose you take a pill. According to evidentialism, the resulting belief is irrational. Still, we may say that a practical factor rationalises the formation of your (irrational) belief, or, equivalently, that a practical factor rationalises the way of forming your irrational belief. Although you did not form your belief through a mechanism liable to rationalise your belief, it was not irrational to form the belief in this way. Third, you may have a practical reason to form the belief that p , even if your evidence is against p . Here too, taking a pill would be a rational way of forming the irrational belief that p .

There are practical reasons to form (rational or irrational) beliefs, through epistemically rational

or irrational mechanisms. There may also be no practical reason to form (rational or irrational) beliefs, or practical reasons not to form (rational or irrational) beliefs. All this is compatible with evidentialism, which is a claim about when a belief is rational, i.e., about what processes of belief formation can make a belief rational, and not a claim about when and how it is rational to form a (rational or irrational) belief.

This point is noted in particular by Engel (2009, 195-196) and Fantl and McGrath (2002). Thus, while attacking (Evidentialism-EJ), the claim that evidential considerations alone determine whether a certain belief is *epistemically* justified, Fantl and McGrath agree that this sort of evidentialism “is not defeated simply by noting that we often take (and ought to take) non-evidential considerations into account before forming or acting on a belief”. They suggest, for example, that an advocate of (Evidentialism-EJ) can appeal to Foley’s distinction between “responsible belief” and “epistemically rational belief” to explain why practical considerations are sometimes relevant when we assess the rationality of belief formation (and action).

Foley thinks that rationality is a goal oriented notion, and he argues that epistemic rationality is concerned with a “very specific goal, that of now having beliefs that are both accurate and comprehensive”.¹⁶ According to many epistemologists, this notion of epistemic rationality is the notion relevant to knowledge, the one which identifies with what is traditionally called ‘epistemic justification’.

In contrast, the notion of responsible belief takes other goals into account. Thus, in light of the mere epistemic goal, a belief may be judged “rational”, in the sense of “epistemically rational”, whereas if one takes all the goals into account, it may be judged “irresponsible”. (According to Foley, a belief in *p* is responsible if the subject has an epistemically rational belief that her procedures with respect to *p* has been acceptable given the limitation on her time and capacities and given all her goals).

Foley identifies the notions of responsible belief and justified belief. Indeed, following the pragmatist line of thought, Foley thinks that practical considerations can justify a belief. However, he makes it clear that this notion of justified belief is distinct from the notion of epistemically rational belief, and that it does not capture a condition on knowledge. Thus, although Foley is a pragmatist concerning justification to believe, he is not a pragmatist concerning knowledge and epistemic rationality.¹⁷ According to him, while the evidential standards for responsible belief can vary with the practical context, whether a belief is epistemically rational is only a matter of having sufficient evidence, where what counts as sufficient evidence is not determined by practical considerations, but by the mere epistemic goal of having now accurate and comprehensive beliefs.¹⁸

¹⁶See Foley (2005, 317). See also Foley (1987, 2002).

¹⁷It is important to note that, according to Foley, the notion of epistemically rational belief does not capture a condition on knowledge either. According to him, we must distinguish the externalist notion of warrant and the internalist notion of epistemic rationality, and the notion of warrant only is relevant to knowledge (see, e.g., Foley (2002)).

¹⁸For more on Foley’s view, see chapter 3 section 3.1.

Foley's distinction may help the *epistemic* evidentialist, who endorses (Evidentialism-EJ) and (Evidentialism-K). Indeed, as explained above, epistemic evidentialism is consistent with pragmatism about reasons to believe, since epistemic evidentialism is a thesis about *epistemic* justification of belief and knowledge, and not a thesis about overall justification of belief. However, it is difficult to see how Foley's distinction could help the evidentialist who maintains that evidence alone is relevant to the justification of belief, and who thereby endorses epistemic evidentialism but also denies pragmatism. According to this version of evidentialism, the notion of justification of belief is reducible to the notion of epistemic justification of belief. If we reject pragmatism about reasons to believe, we must give a different account of the distinction between responsible belief and epistemically rational belief.

A natural suggestion is to understand the question whether a certain belief is responsible as the question whether it is rational to form this (rational) belief, and the question whether this belief is epistemically rational as the question whether this belief is justified. On this approach, cases of epistemically irrational beliefs that are formed for practical reasons are cases of responsible beliefs that are irrational.

Is this approach consistent with evidentialism? Given the distinction between rational formation of beliefs and rational beliefs, the claim that pragmatic considerations are relevant, or necessary, to rationalise the formation of a belief can create a trouble for (epistemic) evidentialism only if the following claim is true:

(Rational formation of beliefs) S ought to form the belief that p if and only if the belief that p is (epistemically) justified for S.

Suppose that (Rational formation of beliefs) is false. Then, the fact that practical factors are relevant to where and when I should form a belief creates no trouble for (epistemic) evidentialism. Maybe the belief that p is (epistemically) justified for S, but, for practical reasons, it is not rational for S to form this (epistemically) justified belief. Alternatively, perhaps the belief that p is (epistemically) unjustified for S, but, for practical reasons, it is rational for S to form this epistemically unjustified belief.

Evidentialism and epistemic evidentialism as such are not committed to (Rational formation of beliefs).¹⁹ And it is in particular controversial that if the belief that p is (epistemically) justified

¹⁹It is unclear whether or not Owens agrees. On the one hand, he seems to agree when he writes:

to get the evidentialist view exactly right, we must distinguish the rationality of beliefs themselves from the rationality of actions intended to induce belief. It may be rational to induce a false belief in oneself if one can't live with the truth, but the belief thereby induced remains irrational: the practical rationality of the belief-inducing action does not affect the theoretical irrationality of the induced belief. For example, the evidentialist can perfectly well admit that it is reasonable for me to visit the hypnotist in order to ensure that I think well of my boss and thereby improve my promotion prospects, regardless of how little there is to be said in his favour. And because this action is reasonable, a rational person can get herself to visit the hypnotist simply by reflecting on how desirable it would be to have this opinion. What the evidentialist will deny is that the belief induced by the hypnotist is reasonable: my

for S, then S ought to form the belief that p. But if this latter claim is false, it is not surprising that practical factors are sometimes relevant, or necessary, to determine whether S ought to form the (epistemically) justified belief that p. However, again, this does not impugn (epistemic) evidentialism, which is a claim about when a belief is (epistemically) rational or justified, and not a claim about when forming a belief is rational.

Thus, suppose you endorse (Evidentialism-JB). Then, you deny that there are practical reasons to believe. But you may say that we sometimes have practical reasons not to form justified beliefs. Suppose you endorse (Evidentialism-EJ). Then, you leave open the possibility that there are practical reasons to believe. But you may say that we sometimes have practical reasons not to form epistemically justified beliefs. These are cases in which we ought not to form a certain epistemically justified belief because this belief is not, overall, justified, given the practical reasons we have not to believe (in Foley's terms, that would be cases of irresponsible but epistemically rational belief).

Against the claim according to which if the belief that p is epistemically justified for S, then S ought to form the belief that p, note that the contrary claim has some plausibility. Suppose that there is absolutely no practical reason to form an attitude as to whether p, say, because p is not worth being considered. Even if the evidence is sufficiently strong to justify a belief that p, it does not seem irrational for S not to form the belief that p. This does not mean that it is permissible for S to believe not-p or to withhold judgement about p. This just means that it is permissible for S to have no attitude toward p.

To illustrate this point, suppose that, in normal conditions, you are reflecting on the fact that it appears to you as if it is raining outside, but you do not form the belief that it is raining outside. Suppose that you do not form the belief that it is raining outside just because you are not interested in what is going on outside. You just want to consider your perceptual state. The fact that you are not motivated to form the belief that it is raining outside, due to your interest at that time, does not show that the belief that it is raining outside is not epistemically justified for you. You would

desire to please my boss can provide no justification for this belief. (Owens, 2000, 25)

On the other hand, he also writes the following against evidentialism:

The evidentialist may be right to insist that whether I believe p rather than not-p is something that should be fixed purely by the balance of evidence for and against p. But where and when I form a view as to whether p is true will be determined by my sense of how important the issue is, what the consequences of having a certain belief on the matter would be and how much of my limited cognitive resources I ought to devote to it before reaching a conclusion. I can't spend my life assimilating evidence for and against all the propositions which interest me; and, even if I could, I would have no way of retaining it all for future consultation. At some stage I must form a (revocable) view on them; once I have done so I can cease to deliberate about them, throw away the evidence both for and against, and think and act on the assumption that they are true (or false). (Owens, 2000, 27)

However, if we grant the distinction between whether it is rational to form a belief and whether the belief is rational, it is unclear why the claim that pragmatic factors are relevant to where and when I should form a belief creates any trouble for evidentialist. Although it is right that pragmatic considerations are relevant to the rational formation of beliefs, this does not concern whether you have enough evidence to rationalise the belief that p, but whether there is a practical reason for you to form such a belief. Even if at some stage you "must form a (revocable) view", we may say that the "must" is pragmatic and not epistemic.

be epistemically justified in believing this if you believed it on the basis of what appears to you. But, intuitively, the lack of motivation to form the belief does not show that you are irrational.

(Epistemic) evidentialism as such is not committed to (Rational formation of beliefs). Arguably, (epistemic) evidentialism as such is not even committed to this weaker claim:

(Epistemically rational formation of belief) S *epistemically* ought to form the belief that p if and only if the belief that p is (epistemically) justified for S.

If S epistemically ought to form the belief that p, then, it seems, p is epistemically justified for S. But it is unclear that if p is epistemically justified for S, then S epistemically ought to believe that p, and thereby, epistemically ought to form the belief that p.

Firstly, as Littlejohn (2012, 46) notes, there is not always an intuition that there is something wrong in not forming a belief even if the evidence is sufficient to justify it:

it is hard to see what wrong you could be guilty of if you do not bother to draw some of the obvious conclusions from your evidence. If you do not form beliefs concerning the obvious logical consequences of things you believe, what of it? Not forming these beliefs is not like not bothering to lift a finger to save a life that could easily be saved.

If there is no sense in which it is wrong not to form these epistemically justified beliefs, it is not *epistemically* wrong not to form these beliefs.

One might be tempted to reply that although it is *prima facie* epistemically required to form all the epistemically justified beliefs, we often have, all things considered, reason not to form all these beliefs. The problem with this reply is that no excuse and regret for not having formed a trivial or totally uninteresting but epistemically justified belief seem appropriate. Further, there seems to be absolutely no residual duty to form these beliefs.²⁰ For example, there is no intuition that we should form these beliefs if there is nothing else to do.

Secondly, the claim according to which you epistemically ought to form a belief that p if p is epistemically justified for you seems too strong. Principles of formation of rational beliefs, governing theoretical reasoning, cannot require to infer all the beliefs that are epistemically justified for us, not even those beliefs we have conclusive evidence for. Indeed, if ought implies can, it is clear that there is no sense in which one ought to form all these beliefs. Further, as Harman points out, one has “no reasons to clutter one’s mind with trivialities just because they follow from other things one believes.”²¹ The principles of formation of beliefs must be such that “they discourage a person from cluttering up either long-term memory or short-term processing capacities with trivialities”. One way to do this is to allow one to “accept a new belief *P* only if one has (or ought to have) an interest in whether *P* is true”.²² Thus, you may easily realise that A and B implies C but not infer

²⁰See Dancy (2004, 29) for the idea that outweighed reasons can explain residual duties and regrets.

²¹See Harman (2002).

²²See Harman (1986, 56; ch. 6).

C because you are not interested enough in inferring the conclusion. That does not mean that you are epistemically irrational, or that C is not epistemically justified for you.

From the fact that there are beliefs that are epistemically justified for us but that we are not epistemically obliged to form, we might want to conclude that there are no epistemic obligations to believe.²³ This conclusion may rely in part on the plausible principle that if it is epistemically required to believe p, then it is epistemically required to form the belief that p. If there are epistemically justified beliefs that it is not epistemically required to form, then these justified beliefs are merely epistemically permissible beliefs. And we may think that for any epistemically justified belief, there is a possible situation in which it is in no sense required to form this belief. If so, epistemic justification to believe merely states a permission to believe. Believing that p with insufficient evidence is always epistemically wrong. Hence, it is always epistemically irrational to form a belief on the basis of insufficient evidence. However, on this view, there is nothing *epistemically* wrong in having no attitude toward a certain proposition, even if there is sufficient evidence supporting this proposition.

Alternatively, from the fact that there are beliefs that are epistemically justified for us but that we are not epistemically obliged to form, we might want to conclude that although there are epistemic obligations to believe, and hence, epistemic obligations to form beliefs, these obligations are conditioned upon the practical situation.²⁴ And, indeed, it is difficult to deny that there are situations in which it would be epistemically irrational for you not to form a belief if the belief is epistemically justified for you. For instance, as Littlejohn notes, “if your evidence is sufficient for knowledge, you are aware of this evidence, you are not aware of any defeaters, and you are concerned with settling the question whether p, it seems irrational for you to withhold judgment or to judge that not-p” (Littlejohn, 2012, 46). The idea that there are epistemic obligations to believe, and hence, epistemic obligations to form epistemically justified beliefs, but that these obligations are conditioned upon the subject’s practical situation (her interests, what she is currently considering, etc.), may seem more plausible than the idea that there are no epistemic obligations to believe.

Be that as it may, suppose that there are no absolute epistemic obligations to believe, either because there are no epistemic obligations to believe or because epistemic obligations are conditioned, and hence, suppose that there are no absolute epistemic obligations to form beliefs. Then, the claim according to which (inconclusive) evidence alone does not motivate the formation of a (rational) belief, or the claim that practical factors are needed to justify the formation of a belief, creates no trouble for (epistemic) evidentialism. An (epistemic) evidentialist may say that appealing to pragmatic factors is required to explain why the subject is (epistemically) obliged to form a certain belief.

In sum, if we do not embrace the controversial claim that there are absolute (or categorical)

²³This is what Littlejohn (2012) concludes.

²⁴See Feldman (2000). See 7.2.4 for the distinction between absolute, conditional and conditioned norms.

epistemic obligations to believe the (epistemically) justified beliefs, and hence, absolute epistemic obligations to form the (epistemically) justified beliefs, it is not surprising that practical factors can be relevant, and even necessary, to rationalise the formation of (epistemically justified) beliefs, even if the evidence alone is sufficient to justify a belief (epistemically speaking). Evidentialism as such is not committed to the idea that there are absolute epistemic obligations to believe. Therefore, evidentialism is compatible with the thesis that practical factors are often (or always) necessary to rationalise the formation of (epistemically justified) belief.

2.5 Conclusion

In this chapter, I have presented and motivated intellectualism, the view according to which only truth-relevant factors determine whether a belief is justified. I have distinguished pragmatism about reasons to believe and pragmatic encroachment. I have explained that pragmatism about reasons to believe is compatible with a version of intellectualism (or evidentialism) regarding *epistemic* justification of belief and knowledge. Next, I have distinguished two issues: the first one is that of the justification of the formation beliefs, and the second one is that of the justification of beliefs. I have suggested that practical considerations are often relevant, or even necessary, to rationalise the formation of beliefs, but that this claim does not compromise the intellectualist account of knowledge and epistemic justification. The intellectualist account of knowledge and epistemic justification is an account of what can rationalise beliefs, not an account of what can rationalise the formation of (rational or irrational) beliefs. I now turn to the claim according to which practical considerations affect the epistemic justification of beliefs and knowledge.

Chapter 3

The threshold problem

3.1 Introduction

According to the pragmatic encroachment thesis, practical factors affect whether a belief is *epistemically* rational (or justified), or whether a subject is in a good enough epistemic position to know. This thesis is inconsistent with epistemic evidentialism. More generally, it is inconsistent with intellectualism. In the previous chapter, I have articulated, but put aside, what I have called the ‘inconclusive evidence argument (result)’ against evidentialism. In this chapter, I assess this argument and I consider the threshold problem on which it relies. The threshold problem is the problem of determining in a non-arbitrary way the degree of justification required to justify belief, or required to be in a position to know, being taken for granted that maximal justification is out of reach, and that justification and knowledge are attainable.

An argument which is sometimes put forward in favour of pragmatic encroachment is that pragmatic encroachment can provide a solution to the threshold problem. In brief, it is said, the champions of pragmatic encroachment can say that this threshold is determined in a non-arbitrary way by the subject’s practical situation, and what it is rational for her to do. Hence, pragmatic encroachers would have an advantage over orthodox epistemologists. I argue that this line of argument is illusory. I agree that the threshold problem is a genuine challenge for epistemologists, but I argue that appealing to practical factors does not solve the threshold problem. Indeed, either the practical condition which is supposed to be met for epistemic justification and knowledge is too weak, or a similar threshold problem arises for when this practical condition is met. I conclude that the threshold problem does not favour pragmatic encroachment over intellectualist accounts of knowledge and epistemic justification to believe.

3.2 Inconclusive evidence and arbitrariness

According to intellectualism, whether a belief is (epistemically) justified depends on truth-relevant factors only, such as the evidence. In the previous chapter, I have articulated and put aside an argument against the evidentialist version of intellectualism, from Owens (2000). The argument is the following:

ARGUMENT FROM INCONCLUSIVE EVIDENCE (RESULT)

1. S can have inconclusive evidence that p and a rational belief that p.
2. Inconclusive evidence alone cannot rationalise a belief.
3. If inconclusive evidence alone cannot rationalise a belief, then pragmatic considerations are needed to rationalise a belief.
4. If pragmatic considerations are needed to rationalise a belief, then evidentialism is false.
5. Therefore, evidentialism is false.

It is now time to assess this argument.

The first premise is intuitive. It seems, for example, that I can have a rational belief that I have hands even if my evidence is compatible with the falsity of this proposition. This is the fallibilist view of justification. The second premise is less intuitive. After all, it seems that I can justify the belief that I have hands by saying that it seems to me that I have hands. But this evidence is not conclusive. Finally, to support the third premise, it must be shown that appealing to pragmatic factors can help to explain why a certain belief can be rational even if it is based on inconclusive evidence.

By appealing to lottery considerations, I will explain in the remaining of this section why we may think that the second premise of this argument is plausible. I will look at an attempt to reject this second premise in a fallibilist framework in section 3.1, but I will show that this proposal fails. I will also consider in section 3.2 an attempt to avoid the problem by rejecting the first premise of this argument and by embracing infallibilism. I will argue that this move is insufficient to solve other versions of the problem.

Still, in the section 4 of this chapter, I will argue that there is no reason to endorse the third premise just to solve the threshold problem, for embracing the claim that practical factors play a role in epistemic justification of beliefs does not solve the problem either. I shall then invite the conclusion that the threshold problem does not favour pragmatic encroachment over intellectualism (or evidentialism).

Let us consider why we might think that the second premise, according to which inconclusive evidence cannot justify a belief, is true. While attacking evidentialism, Owens proposes the following argument:

how are you going to tell us, in purely evidential terms, what level of evidence is needed to justify belief? Unless this question can be answered, evidentialism (internalist and externalist) must be abandoned. (Owens, 2000, 26)

The underlying line of reasoning seems to be the following. We cannot tell in purely evidential terms what level of evidence is needed to justify belief. The best explanation why we cannot tell in purely evidential terms what level of evidence is needed to justify belief is that (inconclusive) evidential considerations alone cannot justify a belief. This explanation contradicts evidentialism. Therefore, we should abandon evidentialism.

If this is the reasoning, it is disputable though. We might think that there are other explanations why we cannot tell in purely evidential terms what level of evidence is needed to justify belief. In particular, we might think that the reason why we cannot tell in purely evidential terms what level of evidence is required to justify a belief has to do with our cognitive limitations, rather than with the falsity of evidentialism. Further, we might object that, in fact, we *can* tell in purely evidential terms what level of evidence is required to justify belief. I shall consider these two objections in turn, and show that they fail to address the problem that Owens has in mind.

First, consider the suggestion according to which the reason why we cannot tell in purely evidential terms what level of evidence is required to justify a belief has to do with our cognitive limitations. Is this claim compatible with evidentialism? Consider internalist accounts of justification of beliefs. In general, internalism about justification requires that everything that is relevant to justification must be in some sense accessible to the subject.¹ Since the norm of belief (and whether it is satisfied) is relevant to whether a belief is justified, an internalist must grant that the norm must be in some sense accessible. That is, how much evidence is required to justify a belief must be in some sense accessible. Still, it is necessary to distinguish between access and reflective access (or motivation and reflective motivation). Perhaps the reason why we cannot tell in purely evidential terms what the required threshold is has to do with the fact that we have no reflective access to what this threshold is. We can have access to things we have no reflective access to, and we can be guided by norms without being able to tell what they are. If so, even if we cannot tell what is the degree of evidence sufficient for justification, this does not compromise evidentialism.

To illustrate this point, suppose that, as some epistemologists think, we should understand the normative regulation of beliefs with the idea of procedural knowledge, or know-how, and that to follow norms is to have internalised plans, habits and conditioned reflexes (Pollock and Cruz, 1999, 127 sq). On this view, “our epistemic norms are just the norms that describe this procedural knowledge, and rational cognition is cognition in compliance with the norms” (Pollock and Cruz, 1999, 129). If this is correct, we can be guided by norms, and have justified beliefs, without being

¹We can thus define externalism as the negation of internalism. According to externalism, not every factor relevant to justification need to be internal, or accessible. See Goldman (1999) for a criticism of internalism. For a defence of internalism, see Conee and Feldman (2001, 2004); Wedgwood (2002).

able to tell what these norms are:

Think, for example, of an expert golfer who knows how to swing a golf club. Nevertheless, he does not always get his stroke right. It is noteworthy, and it will be important later, that when he does not get his stroke right he is often able to tell that by something akin to introspection. When he does it wrong it “feels wrong”. The ability to tell in this way whether one is doing something right is particularly important for those skills governing performances (like golf swings) that take place over more than just an instant of time, because it enables us to correct or fine tune our performance as we go along. (Pollock and Cruz, 1999, 128).

Similarly, “[b]y virtue of knowing how to reason we know how to tell right reasoning when we see it, and that provides us with our data”. But this does not imply that we can explicitly articulate the norms we are following when we reason correctly. As Pollock and Cruz write:

because our automatic processing system operates in a non-intellectual way without any conscious monitoring, it need not be obvious to us what makes a particular belief justified even when it is evident to us that it is justified. Our data consists in the fact that various beliefs are justified – not why they are justified. (Pollock and Cruz, 1999, 160)

In sum, the idea that we cannot say what amount of evidence is required to justify a belief is consistent with a non-reflective form of evidentialist internalism.

Of course, there are types of evidentialist internalism which require reflective access. For instance, some philosophers might think that in order to be justified, we must have current reflective access to the norm and to the fact that we satisfy it. But it is well known that it is an implausible account of epistemic justification, because it leads to an infinite regress.²

Other philosophers advocate for a weaker view, according to which a belief is justified only if the norm is reflectively accessible. For instance, Chisholm writes:

the internalist assumes that, merely by reflecting upon his own conscious state, he can formulate a set of epistemic principles that will enable him to find out, with respect to

²See, for instance, Pollock and Cruz (1999, 125) for a criticism of the “intellectualist model” of how epistemic norms govern the reasoning. The intellectualist model states that epistemic norms regulate the reasoning through the consideration of explicit articulated norms, like those that one can find, for instance, in driving manuals and etiquette books. If we had to make appeal to explicit judgements about epistemic norms in order to reason in each particular case, then there would be a regress. To judge that it is dark on the basis of a sensation of darkness, we would have to appeal explicitly to the norm that, for instance, a sensation of darkness (prima facie) justifies the belief that it is dark. But then, we would have to appeal to explicit norms justifying the appliance of this norm to the case. That is, one would have to justify the belief that there is such an epistemic norm (namely, that it is permissible to believe p on the basis of the sensation that p), and second, the belief that such sensation is a sensation of darkness. In order to justify these beliefs, we would have to appeal to further norms, and norms for applying these norms, and so on. See also Littlejohn (2012, 15-16).

any possible belief he has, whether he is justified in having that belief. The epistemic principles that he formulates are principles that one may come upon and apply merely by sitting in one's armchair, so to speak, and without calling for any outside assistance. In a word, one need only consider one's own state of mind. (Chisholm, 1988, 286)

Unless we endorse this view, it is unclear why the (supposed) fact that we cannot tell (in purely evidential terms) what the epistemic norm is would be problematic for evidentialism. But there are many other internalist views. Perhaps, for example, it is sufficient to be in a position to tell (always, or most of the time) whether a belief is justified, just in the same way as the golfer who can feel that he is doing right or wrong.

Now, consider externalism about justification. It is clear that externalism does not require that the norm, and whether it is satisfied, should be reflectively accessible. Accordingly, the same objection to Owens' argument can be made in an externalist framework.

To sum up: many internalist views and all externalist views of epistemic justification maintain that the epistemic norm as such need not be reflectively accessible. If so, it might be said, it is unclear why the (alleged) fact that we cannot tell what level of evidence is needed to justify a belief would support the rejection of evidentialism in general.

Let us turn to the second objection to Owens' argument. According to this objection, we can tell in purely evidential terms what level of evidence, or justification, is required to justify a belief. Indeed, there are plausible proposals appealing to purely evidential factors. Consider the case of perceptual beliefs. A foundationalist may argue that there is a norm according to which if it seems to me that p , I have a *prima facie* reason to believe that p , and that if this reason is not defeated, I have sufficient reason to believe that p . But this epistemic norm does not appeal to non-evidential considerations. That it seems to me that p is evidence that p . Defeaters are evidence that p is false, or considerations that cancel the application of the epistemic norm by showing that the supposed evidence is not good. For instance, if I have some evidence that the wall is illuminated by a red light, that it seems to me that the wall is red is not a sufficient reason to believe that the wall is red. Therefore, as far as perception is concerned at least, it seems in principle possible to tell what level of evidence is sufficient for justification, even when the evidence is not entailing. One might think that this strategy can be extended to other sources of justification (memory, testimony, etc.).

Thus, the two objections one might want to put forth against Owens' argument are the following. First, it is unclear why the fact (if it is a fact) that we cannot tell what level of evidence is required to justify a belief speaks against evidentialism in general. Second, it is unclear that, in fact, we cannot tell in purely evidential terms what level of evidence is needed.

These two objections do not really address the problem that Owens has in mind, though. Of course, philosophers have proposed epistemic rules appealing merely to evidential considerations. The problem is that we need to show that these rules are not arbitrary. In addition, the problem is not that we have cognitive limitations, but that there can be in principle no purely evidential

reason to set the threshold at such or such level.

To illustrate this point, consider the rule according to which if it seems to me that p and there are no defeaters, then I have sufficient justification to believe that p . We may ask why the fact that it seems to me that p is a (*prima facie*) reason to believe that p . After all, as the sceptic notes, that it seems to me that p is compatible with the falsity of p .

We can also state the problem in probabilistic terms. According to probabilistic approaches of epistemic justification, there are degrees of justification, and these degrees are measured by a probability function. The degree to which a proposition p is supported by evidence E is given by $\Pr(p|E)$, and E is evidence for a proposition p only if the probability of p given E is greater than the prior probability of p . E is a better support for p than E^* if the probability of p given E is greater than the probability of p given E^* . How we should interpret the notion of evidential probability – and epistemic probability (if these are different notions) – is a vexed issue and I will not take a stand on this.³ Nevertheless, it useful to illustrate – in very rough lines – how this kind of account

³There is a clear distinction between evidential (or epistemic) probability and physical or objective probability. A law of nature has an objective probability of one, but it can be very uncertain for us, given our evidence, and hence, it can have evidential probability less than one for us. Further, we should distinguish evidential (or epistemic) probability and the actual credence (degree of belief) a subject would have. Indeed, actual subjective credences are often irrational. One might then try to define evidential (or epistemic) probability in terms of *rational* credences. Now, this idea can be understood in one of two ways. According to probabilism, credences are rational if and only if they are probabilistically coherent, in the sense that they conform to the axioms of probability calculus. It can be shown that credences, understood as betting quotients, that do not satisfy the axioms of the probability calculus, are subject to Dutch books. A Dutch book is a set of bets in which the subject will lose no matter what happens. Some philosophers think that this shows that probabilistically incoherent degrees of belief are irrational. Nevertheless, there are at least two reasons to think that this notion of rational degree of belief (i.e., the notion of probabilistically coherent credences) does not capture the notion of evidential (or epistemic) justification. The first one is that the rationality in question is practical, and it is unclear that we can explain epistemic rationality in terms of practical rationality. The second one is that this rational constraint on degrees of belief does not determine a unique set of degrees of belief (there are many ways of being probabilistically coherent). Still, as far as epistemic rationality is concerned, it does not seem that different degrees of belief can be rational for two subjects if they have the same evidence. The second way of understanding the idea that evidential (or epistemic) probability is a rational degree of belief would be to say that evidential (or epistemic) probability is an *epistemically* rational degree of belief. But then, if epistemic rationality is a question of evidential (or epistemic) probability, this is not illuminating (see in particular Pollock and Cruz (1999, 92sq); they conclude – against probabilistic accounts of justification – that the notion of epistemic probability cannot help to analyse the notion of epistemic justification). One might then be tempted to think that evidential probability is a *sui generis* kind of probability, that captures a kind of logical relation between propositions, namely, that of partial entailment, where the probability is a measure of the strength of the entailment – the objective inductive relation between propositions. However, among other difficulties, it is unclear what the inductive relation between propositions is, and how it could be known. Further, it is not clear how the notion of degree of entailment or confirmation – assuming that this is equivalent to the notion of evidential probability – has to be related to the notion of epistemic probability – conceived as the degree of belief that one is justified in having. It does not seem that one is always justified in having the degree of belief in p equal to the degree of support that one's evidence has for p . For instance, if evidence E entails p , then the evidential probability for p is one. However, the entailment may be so complex that one may be not justified in having a full belief in p (i.e., credence one in p). Also, any logically true proposition will have an evidential probability of one (a logically true proposition is entailed by any proposition). But it is clear that one is not justified in having credence one in a proposition just because it is logically true. One might then think that evidential (or epistemic probability) is the degree of belief that a perfectly rational being would have in the proposition. But this does not work either. A perfectly rational being would have – qua perfectly rational being – evidence that an imperfect rational being could not have. Then, the degree of belief he would have in some hypothesis could not be equivalent to a degree of belief it would be rational for an imperfectly rational being to have (see Williamson (2000, 109)). According to Williamson, evidential probabilities

might work.

Assume that evidential probability can be interpreted as a kind of statistical probability having to do with *observed* frequencies. It is intuitive, for instance, that the observation that a raven is black is some evidence in favour of the fact that the next raven will be black (or that all ravens are black). The degree of justification for believing that the next raven is black may then depend on how many observed ravens have been black so far. Suppose that until now, all observed ravens have been black. The evidential probability for the proposition that all ravens are black – conditional on the proposition that all ravens observed so far are black – may be close to one. Suppose that 90 out of 100 observed ravens have been black so far. One may say that you are justified in believing on this basis that the next raven is black to a degree very close to 0.9. And so on.

Now, suppose that only conclusive evidence confers a probability of degree one to a proposition, but that a proposition can be justified even if its evidential probability is less than one. We should ask: If not one, what degree of evidential probability is necessary and sufficient for epistemic justification? The problem is that any degree of probability that falls short of one seems arbitrary. In order to see this, consider the lottery paradox.⁴

Suppose you know that there is a fair lottery with 100 tickets and only one winning ticket. Then, the evidential probability for you of an arbitrary ticket that it is a loser ticket is .99. Suppose that the degree of probability sufficient to be justified in believing a proposition is .99. According to this threshold, you are justified in believing of an arbitrary ticket that it is a loser ticket. But how could you be justified in believing this? If you are justified in believing this, then, by parity reasoning, you are justified in believing for each ticket that it is a loser ticket. But if we grant that, it seems that you are justified in believing that no ticket is winning. Of course, by assumption, you are not justified in believing that no ticket is winning. Therefore, you are not justified in believing of an arbitrary ticket that it is a loser ticket on the basis of the fact that there is .99 chance that it is a loser ticket. So, .99 cannot be the threshold for being justified in believing something. We can reason in the same way no matter how high the probability is, provided that it is inferior to one. So, it seems, no degree of probability inferior to one can justify a belief.⁵

To avoid the (absurd) conclusion that you are justified in believing that no ticket will win, while maintaining that justification to believe does not require a maximal degree of evidential probability, we might be tempted to reject the appropriateness of parity reasoning. That is, we might be tempted to reject the idea that it is appropriate to reason in the same way for each ticket. But then, something arbitrary follows. Either we grant that we are justified in believing that the last ticket we consider is the winning ticket. This is absurd because the belief that the last ticket is the winning one is grounded in something arbitrary: we had no reason to start with such or such

“measure something like the intrinsic plausibility of hypotheses prior to investigation” (Williamson, 2000, 111).

⁴See Kyburg (1961).

⁵We can also extend this reasoning to any kind of propositions, and not only to propositions explicitly about lotteries.

ticket, and we could have started by considering any other ticket, thereby coming up with a different result. Or we grant we should exclude another ticket. But once again, there is no non-arbitrary way to decide which one we should exclude. They are all on a par and we should consider them all together. We have no ground for exclusion.

In general, the threshold problem seems to arise from the fact that we cannot rationally motivate the exclusion of a portion of the probability space on the mere consideration of its size. Given that any probability space can be divided into equal parts, excluding any portion on the basis of its size is arbitrary. Parity reasoning is always appropriate. As a result, to avoid arbitrariness, we must concede either that justification to believe requires a maximal degree of evidential probability, or – without appealing to evidential factors – we must explain why some portion of the probability space can appropriately be excluded. In the latter case, we must concede that epistemic norms are not purely evidential.

In a nutshell, the claim that we cannot tell in purely evidential terms what level of evidence is needed to justify belief is not a claim about our cognitive limitations. It is a claim about what can, in principle, explain why, and when, inconclusive evidence can be sufficient to justify a belief. This is the threshold problem for evidentialism. Appealing to a distinction between access and reflective access, and to defeating accounts of justification, does not help to answer this question.

I have presented the threshold problem as a problem for the evidentialist who grants the fallibilist claim that we can have a justified belief that p and inconclusive evidence that p (this is the first premise of ARGUMENT FROM INCONCLUSIVE EVIDENCE (RESULT)). However, it is important to note that this problem can be generalised to non-evidentialist intellectualist and fallibilist accounts of epistemic justification of belief and of knowledge. Suppose that justification is understood in term of reliability. We may ask what degree of reliability is required for *sufficient* reliability. Suppose that justification is understood in terms of safety. We may ask how far in the possible worlds we should go to see whether the belief is *sufficiently* safe. And so on.

We might think that this problem does not affect the infallibilist approaches to epistemic justification of belief. However, as it will become clear, unless we opt for a sceptical form of infallibilism, in which epistemic justification of belief and knowledge require absolute epistemic certainty, the threshold problem arises for the infallibilists too.

Before considering whether pragmatic encroachment can help to solve this problem, it is necessary to consider what an intellectualist might want to say about it. I shall look at two approaches that may seem promising, the first one fallibilist, the second one infallibilist. Unfortunately, I will argue, these approaches are unsatisfactory.⁶ This will lead me to turn to the anti-intellectualist

⁶The possibility that other approaches could avoid this problem should not be excluded. For example, following Dretske (1971), we might think that in order to have a sufficient reason to believe that p , you must have a conclusive reason R to believe that p , where R is a conclusive reason to believe that p if and only if, if p were not true, R would not be the case. On this view, no threshold problem seems to arise. More recently, Smith (2016) has proposed an account of justification of belief according to which a belief based on E is justified if and only if, given E , normally, p is true, and if not- p , then a special explanation why not- p is needed. This approach seems to avoid the threshold

approach. If the anti-intellectualist approach can solve the threshold problem, assuming the failure of intellectualism on this score, there is a strong reason to abandon intellectualism. In the end, however, it will emerge that appealing to practical factors does not solve the threshold problem either. It follows that the threshold problem remains, and that it does not favour anti-intellectualism over intellectualism.

3.3 Intellectualist approaches to the threshold problem

3.3.1 A fallibilist approach to the threshold problem

Many philosophers conceive of epistemic justification in a fallibilist framework, in which a subject S can be justified in believing p on the basis of evidence E , even if E does not entail p . If we adopt a probabilistic approach, this means that a proposition can be justified for a subject even if the evidential probability of this proposition for this subject is less than one. We then have to explain how much justification, or what degree of evidential probability, we need for sufficient justification.

As we have seen, the challenge is to explain why, even if the justification for p is not entailing, or even if the evidential probability of p is not maximal, a subject can rationally reject the possibility that $\text{not-}p$. The lottery paradox is based on the idea that it seems arbitrary to believe of an arbitrary ticket that it will not win because we have to leave the possibility that it will win open. Indeed, if we may reject the possibility that such an arbitrary ticket is winning, we may do so for each ticket, which is absurd. Lottery considerations show that it is always arbitrary to reject the possibility that $\text{not-}p$ on the basis of inconclusive evidence that p .

Still, suppose that believing that p does not amount to rejecting the possibility that $\text{not-}p$.⁷ If so, it unclear why being justified in believing that p would require being justified in rejecting the possibility that $\text{not-}p$. Suppose that from the fact that believing p is justified, it does not follow that rejecting the possibility that $\text{not-}p$ is justified. Then, it cannot be argued that any threshold inferior to one is arbitrary *because* adopting such a threshold would lead to arbitrarily reject a possibility in which $\text{not-}p$.

In this regard, we can appeal to the distinction between outright belief and partial belief (or credence). Suppose we think that justification primarily concerns *degrees of belief*. For example, suppose that the evidential probability that the next raven will be black is 0.9. We may say that it is rational to have the degree of belief reflecting the evidential probability of this proposition, say, .9. Here, there is no threshold problem.

problem. Still, my aim is just to show that the threshold problem does not favour pragmatic encroachment over the orthodox accounts of knowledge. If there are intellectualist approaches which can solve or avoid this problem, this is a further reason to think that the threshold problem should not lead us to abandon intellectualism.

⁷See chapter 5 section 4.3 problem 4 and chapter 6 section 2.3 for defences of the view that outright belief does not entail subjective certainty and is compatible with a certain degree of epistemic anxiety.

Then, we can approach the threshold problem for justification of *belief* in at least two different ways. First, following some philosophers, we might want to argue that the notion of belief should be eliminated for the benefit of the notion of degrees of belief.⁸ If so, the threshold problem is directly avoided. However, this option is not open to the advocate of the orthodox account of knowledge. If we eliminate the notion of belief, arguably, we thereby eliminate the notion of knowledge that this philosopher has in mind.

Second, suppose that we distinguish belief and degree of belief, but that we do not eliminate the notion of belief. The challenge of determining the level of evidence sufficient to justify a belief becomes the challenge of determining the degree of belief which counts as a belief, at least if we adopt the plausible thesis that a belief is a degree of belief sufficiently high.⁹ This influential thesis is often called the ‘Lockean thesis’. Let me distinguish its two components.

First, there is the idea that, following Locke’s view, you must proportion your degree of belief to your evidence:

The mind, if it will proceed rationally, ought to examine all the grounds of probability, and see how they make more or less, for or against any probable proposition, before it assents to or dissents from it, and upon a due balancing the whole, reject or receive it, with a more or less firm assent, proportionally to the preponderance of the greater grounds of probability on one side or the other.¹⁰

The idea is that propositions are supported by some degree of evidential probability and a justified degree of belief reflects this degree of evidential probability. At this stage, there is no threshold problem. The rationality of the degree of belief is entirely a function of the evidential probability. Still, unless we can tell whether this degree of belief amounts to a belief, we cannot tell whether this amount of evidence justifies a belief.

Second, there is the idea that a certain degree of belief is sufficient to count as a belief. Thus, the challenge for the Lockean approach is to say what degree of belief is sufficient to count as a belief. Can we tell in a non-arbitrary way what minimum degree of belief counts as a belief?

We can easily tell for some degrees of belief that they do not count as beliefs, and hence, for some evidential degrees that they are insufficient to justify a belief. It seems clear that if your degree of belief in p is greater than your degree of belief in $\text{not-}p$, you do not believe $\text{not-}p$. This is reflected in the fact that if the evidence supports p rather than $\text{not-}p$, then it is not rational to believe $\text{not-}p$. Evidence justifying a degree of belief less than 0.5 cannot count as sufficient to justify a belief. Also, if you have the same degree of belief in p as in $\text{not-}p$, you do not believe p nor $\text{not-}p$. This is reflected in the fact that when the evidence equally supports p and $\text{not-}p$, then it is rational to withhold judgement. So, to count as justifying a belief, the evidence must justify a

⁸See Jeffrey (1970).

⁹See Foley (1993, ch. 4).

¹⁰See Locke and Nidditch (1979, IV, xv, 5).

degree of belief higher than 0.5. Finally, we might think that you can believe that p without being certain that p . For instance, you can believe that the next raven will be black, even if you are not absolutely certain that it will be black. You can believe that this lottery ticket is a loser ticket, although you are not certain that it is a losing ticket. This is reflected on the fact that we can be justified in believing things we have no conclusive evidence for. If so, the threshold for believing must be inferior to one.

These negative constraints are insufficient to tell us what degree of belief counts as a belief, though. The mere fact that you have a greater degree of belief in p than you have in $\text{not-}p$ does not show that you believe that p . But, we may think, if the Lockean thesis is true and epistemic justification primarily concerns degrees of belief, the idea that there is no way of identifying the amount of evidence required to justify a belief only shows that there is no way of identifying what degree of belief is sufficient to have a belief. But there might be good explanations why this is so which have nothing to do with the falsity of evidentialism, or more generally, with the falsity of intellectualism.

Following a suggestion from Foley, we may argue that a belief-state is not something that correspond to a precise degree of belief. Arguably, the threshold for believing is vague. This does not prevent us from appealing to stipulation when it comes to discussing issues of rational belief in epistemology. For instance, if it seems clear that .99 is sufficient, but unclear whether .9 is sufficient, when discussing about belief, we might just stipulate that the degree of belief corresponding to a belief is .99. If so, the fact that there is some arbitrariness in the way in which the threshold for believing is fixed should not be a problem for intellectualists.

However, there is a deeper problem. Some might think that .9 is clearly sufficient, while doubting whether .85 is sufficient. Yet some others might think that .85 is clearly sufficient, and so on. As Foley points out, “there doesn’t seem to be a non-arbitrary way to identify even a vague threshold.” Still, does that imply that the threshold is totally arbitrary?

Foley’s suggestion is that the threshold can vary from believers to believers, depending on their own deep standards of rationality:

The threshold of belief is determined by how on reflection you would resolve a hypothetical decision problem. The problem is framed in terms of a purely intellectual aim, that of providing an accurate and comprehensive black-and white picture of the world. (...) The two risks [inaccuracies and reduced sharpness of the image of the world] can be balanced in a variety of (...) ways, but there is no one correct answer about how to do this balancing. The relevant balancing for you is one that conforms to your own deep standards, one that further reflection would not prompt you to retract. This is ultimately a matter of your intellectual character. Or perhaps more accurately, it is a matter of the kind of intellectual being that on reflection you would want yourself to be. (Foley, 1993, 199-200)

According to Foley's approach, although it is not possible to tell *in general* what level of evidence is required to justify a belief, this has nothing to do with the falsity of evidentialism. Rather, this has to do with the fact that an answer to this question depends upon an answer to the question what degree of belief counts as a belief. And even if there is no general non-arbitrary answer to this question, there are particular answers which are not grounded in pragmatic considerations. They depend on the kind of believer you are.

In sum, on the one hand, we may grant that there is no general truth regarding the level of evidence required to justify a belief. On the other hand, we may argue that this is not a problem for evidentialism, or intellectualism. For each believer, we can in principle tell what level of evidence is required to justify a belief.

To see why this fallibilist approach is compatible with intellectualism, it is useful to distinguish the Lockean thesis from pragmatic accounts of belief. According to an influent pragmatic account, whether a certain degree of belief amounts to a belief depends on whether this degree of belief is high enough for the subject to be disposed to act on the proposition in the circumstances.¹¹ On this pragmatic view of belief, whether a degree of belief counts as a belief depends on the practical situation. The Lockean thesis, in contrast, has it that the notion of belief captures a (possibly vague) threshold of confidence independently of the practical situation and invariant across contexts. Thus, according to the Lockean view, a subject can believe that *p* although, in the circumstances, she is not inclined to act on *p* because her degree of belief is too low.¹²

It is also useful to distinguish the Lockean notion of belief from a notion of acceptance sometimes used by philosophers. According to Harman (1986, 47), for example, "Belief in or full acceptance of *P* involves... [allowing] oneself to use *P* as part of one's starting point in further theoretical and practical thinking". In contrast, according to Foley's account, one can believe that *p* without committing oneself to use *p* as evidence in practical or epistemic inference. This is fully consistent with the idea that one can be justified in believing something on the basis of a non-maximal degree of evidential probability without arbitrariness, since being justified in believing that *p* in this way does not imply that one rejects, in fact, or that one is authorised to reject, the probability in which not-*p*.

According to this fallibilist approach, there is a simple explanation of the lottery cases. While a subject can be justified in believing, for each ticket, that it is a loser ticket, that does not imply that she is justified in rejecting the possibility that it is a winning ticket, and that she is thereby justified in believing that no ticket will win. One is not always authorised to use what we are justified in believing as premise in theoretical (and practical) reasoning. In the lottery case, the subject has evidence that one ticket is winning – which explains why, for each ticket, the probability that it

¹¹See, for example, Weatherston (2005). See also chapter 5 section 4.3 problem 4.

¹²We must also distinguish the Lockean thesis from contextualist views of 'belief', according to which what degree of belief a 'belief' attribution refers to depends on the ascriber's context (see, e.g., (Sturgeon, 2008)). On such a view, the threshold varies across conversational contexts.

is loser is not one. It follows that the subject is not authorised to use her beliefs concerning each ticket to conclude that no ticket will win. Further, while a subject can believe, for each ticket, that it is loser, that does not imply that she believes that each ticket is loser.

This solution implies the rejection of the principle of epistemic closure about justification (at least in its multiple premise version) and the rejection of the idea of closure of belief. According to the principle of epistemic closure about justification (in its multiple premise version), if you are justified in believing p , and justified in believing q , then you are justified in believing what you competently deduce from p and q . From “ p ”, and from “ q ”, you can competently deduce “ p and q ”. Now, from “ticket 1 is loser”, and from “ticket 2 is loser”, and from “ticket 3 is loser” etc., you can competently deduce “ticket one is loser and ticket two is loser and ticket 3 is loser...” (by a multiple premises deduction). From this, you can competently deduce “All tickets are loser” (by a single premise deduction). So, if we grant the principle of closure on justification (in its multiple premise version), on this approach, since it is granted that we are justified in believing that, for each ticket, it is a loser ticket, we should say that the belief that not ticket is winning is justified, which is obviously false. The rejection of this principle is to be expected if epistemic justification is understood in terms of evidential probability less than one. If p has evidential probability .9 and q has evidential probability .9, then the evidential probability of p and q is .81. If the threshold of justification is .9, then your belief in p is justified, and your belief in q is justified, but your belief in p and q is not justified.

According to the idea of closure of belief, if you believe p , and you believe q , then you believe p and q . But if this principle is correct, then, if you believe “ticket one is loser”, and you believe “ticket two is loser”, etc., then you believe “ticket one is loser and ticket two is loser, etc.” Hence you believe that no ticket is winning, which is obviously false. The rejection of belief closure is to be expected if beliefs are understood in terms of degrees of belief above a certain threshold. If your credence in p is .9 and your credence in q is .9, then, if you are probabilistically coherent, your credence in p and q is .81. If the threshold for belief is .9, then you believe that p , and you believe that q , but you do not believe p and q .

Let me sum up how an orthodox fallibilist might want to solve the threshold problem. By distinguishing degrees of belief and beliefs, it may be argued that justification concerns primarily degrees of belief. Next, we can adopt the Lockean thesis. Finally, we may say that the threshold for believing depends on the psychology of the believer (her own deep epistemic standards arising from the kind of believer she wants to be). Thus, there is no general correct answer to what degrees of belief should count as a belief. But this fact is not a problem for evidentialism or intellectualism. The level of evidence required to justify belief depends on what degree of belief, on careful reflection, you would count as a belief.

While interesting, this line of thought faces at least two objections, though. First, it is unclear that appealing to the psychology of the believer, and to what kind of intellectual being she

wants to be (on careful reflection), is sufficient to respond to the challenge. The challenge is to determine in a non-arbitrary way what evidential threshold applies to the believer, while avoiding anti-intellectualism. Suppose the threshold is fixed by the pure will of the believer. If so, pragmatic factors play no role, but the threshold it is arbitrarily fixed. Suppose there is a reason for the believer to endorse such or such threshold. Presumably, this reason is pragmatic. If so, the threshold is ultimately determined by practical factors.

Second, it is unclear that the notion of Lockean belief can be identified with the attitude required by knowledge. Intuitively, there is a sense in which it can be rational for you to believe of a lottery ticket that it is a loser ticket. This is the sense according to which you may, at the same time, rationally believe that you do not know that the ticket is a loser ticket. For example, you may say without incoherence “I believe that this ticket is loser, but I do not know. Perhaps it will win, after all”. However, if you rationally believe that you do not know that p , or rationally believe that you are not in a position to know that p , arguably, you should not have the attitude licensed by knowledge-level justification. Thus, when you say that you believe of a lottery ticket that it is a loser ticket, you do not thereby say that you have the attitude required by knowledge. Or when you say that you believe that p , you do not thereby represent yourself as knowing that p . This is possible because having the attitude required by knowledge and merely believing – in the Lockean sense – are different attitudes. Now, if a Lockean belief is too weak an attitude to count as the attitude required by knowledge (and licensed by knowledge-level justification), even if we think that the threshold problem for Lockean beliefs can be solved by appealing to the believer’s psychology, the threshold problem for knowledge (and knowledge level-justification) remains.

3.3.2 An infallibilist approach to the threshold problem

I have considered how a fallibilist, evidentialist and probabilistic account of justification might try to deal with the threshold problem by distinguishing beliefs and degrees of belief. This approach is insufficient to solve the problem, though. A natural suggestion at this stage is to reject fallibilism and the first premise of INCONCLUSIVE EVIDENCE (RESULT). We can suggest that conclusive evidence, or an evidential probability of one, is required for epistemic justification and knowledge. We hereby avoid the threshold problem. In the lottery case, for instance, as the evidential probability that a given ticket is winning is not 0, you are not justified in believing outright and cannot know that this ticket is a losing ticket. This is intuitive enough.

The problem, though, is that the claim that epistemic justification (and knowledge) requires an evidential probability of one seems to entail that most of our beliefs are unjustified. If you have sufficient epistemic justification for p only if the probability of p given your evidence is 1, then it follows that what you are justified in believing is entailed by your evidence.¹³ And, it is said, this

¹³Following Brown (2011, 149-160), we may want to distinguish entailment infallibilism, the thesis that the evidence upon which the subject’s belief is based must entail the believed proposition, and probability one infallibilism, the

leads to scepticism because in the majority of cases your evidence does not entail the truth of your belief.

The argument for this consists in putting forward “bad cases” where the subject is supposed to have the same evidence as the subject in the “good cases”, but where the believed proposition is false. For example, in the brain in a vat scenario, it is said, the subject has the same evidence – the same perceptual states – but the believed proposition is false. Therefore, the perceptual evidence is not entailing. If so, granting that knowledge requires entailing evidence, the subject does not know in the good case. This leads to sceptical conclusions, because we can imagine bad cases for almost any proposition.

A standard reply has been to reject the claim that the subject has the same evidence in the good and bad cases.¹⁴ One may argue that in the good case, p is true, and this somewhat makes a difference in terms of the subject’s evidence, and hence, on whether the subject’s evidence entails p or whether p has an evidential probability of one. For instance, if factive mental states can be evidence, the subject may believe p on the basis of the factive mental state “perceive that p ”, so that p is entailed by her evidence. In contrast, in the bad case, p is false, and hence, the subject’s evidential basis cannot include this factive mental state. So, we may assume that one has knowledge-level justification for p only if there is no possible case where one believes p with the same evidence and not- p is the case, while avoiding scepticism. All depends on how we think about evidence.¹⁵

In this externalist approach, what the bad case shows at most is that – even if one has entailing evidence for p , or p has an probability of one given the evidence – there is not always an accessible mark guaranteeing that we have the evidence in question.¹⁶ But to say that one cannot always tell what the evidence one has is does not amount to saying that one cannot tell in purely evidential terms what level of evidence is required to justify a belief. Indeed, one may argue, maximal

thesis that the subject’s total evidence must entail the believed proposition. Since nothing relies on this distinction in what follows, I put it aside.

¹⁴See, e.g., Williamson (2000, ch. 8).

¹⁵See also Fantl and McGrath (2009); Brown (2011, 160-161). Williamson identifies knowledge and evidence. It follows from this – assuming that the subject knows that p in the good case – that there is no possible bad case in which the evidence is the same and not- p is the case: “If the sceptic tries to stipulate that the bad case is a case in which one falsely believes p while having the same evidence as one has in a case in which by externalist standards one knows p , those externalists will reply that, so defined, the bad case is impossible.” (2000, 169) Still, it is important to note that Williamson does not defend the idea that the level of evidence required to justify a belief must be entailing or should raise the probability of the target proposition to one. The claim is only that what you know has a probability of one on your evidence – as a consequence of the $E=K$ thesis. But it is also granted that you can be justified in believing something you do not know, for instance just because it is false. As he writes: “The equation of one’s evidence with one’s knowledge does not imply any particular theory of how a given body of propositional evidence justifies a given belief” (Williamson, 2000, 9).

¹⁶In Williamson’s words, epistemic rationality is not transparent. A rational thinker is not always in a position to tell what rationality requires of him because he is not always in a position to know what the evidence he has is. The opposite views rely on a phenomenal conception of evidence, where p is evidence for S if and only if, if S ’s evidence has property p , then S is in a position to know that his evidence has property p . Williamson proposes an antiluminosity argument against this conception. Since my aim is not to defend externalism about evidence, I will leave this argument aside. See Cohen (2010) for a criticism of the argument, and Srinivasan (2015) for a defence.

evidential probability, or entailing evidence, is required to justify a belief.

However, there is a problem here. Under this externalist interpretation, an evidential probability of one does not seem to be the highest possible degree of justification. For one thing, any logically necessary proposition must be attributed an evidential probability of one, because it is entailed by any set of evidence whatsoever, but it does not seem that these propositions are always justified for us. For another, one might have entailing evidence and yet reasons to doubt that the evidence is entailing, or the evidence may be entailing because it is contradictory.¹⁷ In such cases, the evidential probability of the target proposition must be one even if, intuitively, it is not justified.

Some philosophers may disagree. For example, Brown (2011, 160) writes:

Even if p is logically necessary, it does not follow that its probability on one's evidence is one. For instance, one could have inductive or testimonial evidence for a true mathematical proposition where that evidence gives the relevant proposition a probability less than one. Further, if one's evidence is inconsistent, it does not follow that the probability of any proposition p on that evidence is equal to one.

Still, if evidential probability satisfies the axioms of the probability calculus, a logically true proposition must be assigned an evidential probability of one, and if p logically entails q , then $\text{pr}(q|p)$ must be one.¹⁸ These claims are at the basis of the well-known logical omniscience objection to probabilism. Probabilism is the view according to which rational degrees of beliefs (rational credences) are probabilities, and hence, are constrained by the axioms of the probability calculus. The objection is that if rational degrees of belief are probabilistically coherent, then you must have credence one in any logical truth, and hence, it is rational for you to bet on any logical truth whatever the stakes. Still, it is intuitively irrational for a subject to take such bets. Rationality does not require logical omniscience. Therefore, rational credences are not fully constrained by the axioms of the probability calculus.

Now, setting aside probabilism, if we think that evidential probability satisfies the axioms of the probability calculus, we must say that any logical truth has an evidential probability of one. But if it is not always rational to take a bet on a logical truth whatever the stakes, we should think that an evidential probability of one is not always sufficient to justify credence one in a proposition. And if maximal evidential probability is not always sufficient to justify credence one, it seems that we should accept that something else may be relevant to epistemic justification. If so, perhaps we should concede that we cannot tell in purely evidential terms what level of evidence is required to justify an outright belief, after all, because evidential considerations alone may not always be sufficient to justify an outright belief. Responding to the inconclusive evidence argument by endorsing probability one infallibilism would not be sufficient to solve the threshold problem.

¹⁷Of course, if the evidence is defined in terms of what is known, the evidence cannot be contradictory.

¹⁸See, e.g., Pollock and Cruz (1999, 105).

In other words, it seems that we should distinguish evidential probability one and sufficient justification to believe. Sometimes, since the evidential probability of a proposition is not transparent, an evidential probability of one is insufficient for justification. Therefore, requiring a probability of one does not solve the threshold problem.

To illustrate this point, we can consider the high-stakes bet argument directed against the claim that there are propositions which have an evidential probability of one for us. The natural reply to such an argument suggests that an evidential probability of one may be insufficient to justify an outright belief. Indeed, in a nutshell, this reply has it that we should distinguish outright belief and credence one, and that an evidential probability of one justifies an outright belief but does not justify a credence of one. If this reply is correct, an evidential probability of one is not the highest degree of justification, since this degree of justification can be insufficient to justify a credence of one. But if an evidential probability of one is not the highest degree of justification, then it is unclear why this degree of justification is sufficient to justify an outright belief.

Let me explain in more details. The high-stakes bet argument aims at showing that we are never absolutely certain (in the sense of having credence one) of anything. Here is the argument:

If one's credence in p is 1, one should be willing to accept a bet on which one gains a penny if p is true and is tortured horribly to death if p is false. Few proposition pass that test. Surely complex logical truths do not even though the probability axioms assign them probability 1. (Williamson, 2000, 213-214)

If we assume that evidential probability is actual credence, then it follows that no proposition has an evidential probability of one because one is never willing to take this kind of bet.¹⁹ As Williamson notes, we can reply by distinguishing evidential probability and subjective probabilities (or credences):

we should question the association between evidential probability 1 and absolute certainty. For subjective Bayesians, probability 1 is the highest possible degree of belief, which presumably is absolute certainty. (...) But since evidential probabilities are not actual (...) credences, why should evidential probability 1 entail absolute certainty? (Williamson, 2000, 213-214)

Still, it seems that a slightly different argument could be made, concerning whether one is *justified* in having absolute certainty. Indeed, we may think of evidential probability in terms of *rational* credence. And, arguably, if for any proposition, there can be a bet on it such that it is irrational to take it, this may show that for no proposition it is rational to have the highest credence – absolute

¹⁹A similar argument concerning knowledge can be found in Unger (1975). According to Unger, knowledge implies certainty, and one is never certain of anything. As a result no one knows anything. For more on the notion of certainty, see chapter 6.

certainty.²⁰

Following Williamson, we can also reply to this argument by rejecting the interpretation of evidential probability in terms of “counterfactual” credence, and argue that evidential probability is not rational credence either (credence we would have were we to be perfectly rational). According to this approach, there is no direct connection between rational credence in p and the evidential probability that p has for us. In particular, we may say that there is no direct connection between the fact that p has an evidential probability of one for S and the fact that S is justified in having credence one in p . Consequently, on this approach, a proposition can have an evidential probability of one for you, you can know it, even if it is irrational for you to have credence one in this proposition, and hence, irrational for you to take high-stakes bets.

As is plain, however, this reply to the high-stakes bet argument raises a new question: How is an evidential probability of one relevant to the question whether an outright belief is justified? Suppose, as this approach assumes, that we should disconnect evidential probability of one and rational credence of one. Then, it seems that we sever the direct connection between evidential probability of one and justification of outright belief.

We can state the problem as follows. The evidentialist arguing that we can tell in purely evidential terms what amount of evidence justifies an outright belief – namely, the amount of evidence conferring an evidential probability of one to the target proposition – seems to face a dilemma. Suppose she accepts that an outright belief is justified for S if and only if its evidential probability is one for S . Then, she seems to be led to deny that most of our outright beliefs are justified – as the high-stakes bet argument above shows. Suppose she wants to maintain that many propositions have an evidential probability of one. Then, in order to avoid the threat of the high-stakes bet argument, she must deny that an evidential probability of one warrants a credence of one. But this seems to lead to the idea that an evidential probability of one is insufficient for epistemic justification of outright beliefs.

So, does the idea that a degree of evidential probability is not a rational degree of credence sever the connection between evidential probability and epistemic justification of outright belief? For there to be a gap between the notion of an evidential probability of one and the notion of epistemic justification of outright belief, it should be granted that outright belief is the highest credence. Standardly, credences are measured by betting behaviours – they are betting quotients. If outright beliefs are not betting quotients, what is relevant to rational betting quotients can be irrelevant to epistemic justification of outright beliefs. And if an outright belief that p is justified for S if and only if the evidential probability of p for S is one, that does not imply that credence

²⁰In a similar way, Unger (1975) argues that if you know that p , then you have *the right* to be certain, and that you never have the right to be certain because certainty is the attitude such that “no new information, evidence, or experience will now be seriously considered by [me] to be at all relevant to any possible change in how certain [I] should be in the matter; no matter what new experience I may have, I will be no less certain but that p ” (Unger, 1975, 105). While it seems true that one never has the right to have such an attitude, it is unclear that this is the attitude required by knowledge.

one in p is rational if and only if the evidential probability of p for S is one. In other words, the mere fact that it is irrational to bet on a logical truth whatever the stakes does not show that it is irrational to believe outright logical truths. For, according to this approach, even if logical truths have an evidential probability of one for us, and hence, even if we have enough evidence to believe them outright, we are not thereby justified in having credence one in them.

Accordingly, depending on whether we accept the idea that outright belief is the highest credence, we may interpret the high stakes bet case above either as showing that epistemic justification of outright belief does not supervene on an evidential probability of one (only), or as showing that outright belief is not the highest credence. As the view under consideration maintains that epistemic justification of outright belief supervenes on an evidential probability of one, it should say that outright belief is distinct from the highest possible credence.

According to Williamson, in particular, we should distinguish the highest credence and outright belief in the following way:

Intuitively, one believes p outright when one is willing to use p as a premise in practical reasoning. Thus one may assign p a high subjective probability without believing p outright, if the corresponding premise in one's practical reasoning is just that p is highly probable on one's evidence, not p itself. Outright belief still comes in degrees, for one may be willing to use p as a premise in practical reasoning only when the stakes are sufficiently low. Nevertheless, one's degree of outright belief in p is not in general to be equated with one's subjective probability for p ; one's subjective probability can vary while one's degree of outright belief remains zero. Since using p as a premise in practical reasoning is relying on p , we can think of one's degree of outright belief in p as the degree to which one relies on p . (Williamson, 2000, 99)

In this passage, Williamson identifies the degree of outright belief in p to the strength with which we rely on p in practical reasoning. This strength is not what is measured by betting behaviour. Indeed, we may have a very high credence in p – in the sense of assigning a very high degree of probability to p , without being willing to use the proposition that p as a premise in practical reasoning in any way. Consider a lottery case with one million tickets and only one winning ticket. Arguably, the rational credence that a given arbitrary ticket is loser should be extremely high for us. But we are not thereby closer to use the proposition that this arbitrary ticket is loser as a premise in our practical reasoning – rather than the premise that it is immensely probable that it is loser. Intrinsically probabilistic evidence – no matter how high – does not seem to be sufficient evidence to justify an outright belief. If we should never rely on something we do not believe outright, credences are not degrees of outright belief.

But what are degrees of outright belief? To understand the notion of degrees of outright belief, consider Williamson's example. It invokes the proposition "I am feeling cold". The degree of

confidence (or the degree of outright belief) in “I am feeling cold” – as distinct from the betting quotient – varies gradually with the strength with which we feel cold. Similarly, take a perceptual case. The confidence in “This animal is a dog” gradually varies with the strength with which this animal is perceived as a dog. These degrees of confidence can be roughly measured by reactions to the question “Is p the case?”. The stronger the hesitation to respond, the weaker the degree of confidence.

In this respect, we must emphasize that Williamson’s account allows epistemic differences between propositions that have an evidential probability of one for us, and hence, between propositions that we should believe outright. Presumably, these epistemic differences are reflected in the degree to which we should believe outright these propositions. Thus Williamson writes:

whatever evidence is, we are not always in a position to know how much we have of it. (...) [T]herefore, every part of the evidence has probability 1 on the evidence, even though instances of this may not be obvious even to the logically omniscient agent. Some parts of one’s evidence may have other epistemic advantages over others: for instance, one may know of some parts of one’s evidence that they are part of one’s evidence while not knowing (or even being in a position to know) of other parts of one’s evidence that they are part of one’s evidence. Thus some propositions with evidential probability 1 will have epistemic advantages over other propositions with evidential probability 1. (...) Since such epistemic differences between propositions that share evidential probability 1 are inevitable, I am unmoved by their emergence on my account of evidential probability. One’s knowledge is not all on a par in all epistemic respects, even though all of it has probability 1 on one’s evidence. No single probability distribution can capture all those epistemic differences. (Williamson, 2009)²¹

In sum, Williamson distinguishes degrees of evidential probability and degrees of justification of outright belief (degrees of justification of outright belief can vary while the evidential probability of the target proposition remains one), and he relatedly distinguishes credences and degrees of outright belief.

Let me illustrate the difference between credences and degrees of outright belief by considering a lottery case again. Suppose that you know that there is either one winning ticket in the lottery (p) or none ($\text{not-}p$), and the probability of p is 0.5. Suppose you can take a bet on p . Suppose that the tickets are drawn one by one, and whenever a ticket is drawn and it is a losing ticket, you can take the bet on p again. Your betting quotient for this bet should decrease regularly as the number of drawn loser tickets increase. (If the chance that p is 0.5 given all the tickets, say 1000, it decreases to 0.499 after the first drawn, to 0.498 after the second one, etc.). In contrast, it seems

²¹See also Fantl and McGrath (2009, 186-187): “The proposition *I was born in the United States* can be part of your total evidence, and so have evidential probability 1. But you are not rational to gamble on it if the stakes are too high”.

that the degree of confidence that there is a winning ticket does not decrease in this way, until the last ticket has been drawn. Intuitively, it does not become clearer and clearer that there is no winning ticket as the losing tickets are subsequently drawn – as it becomes clearer and clearer that this animal is a dog as you come closer and closer to it. Until the last ticket has been drawn, given the evidence, there is a chance that there is a winning ticket, and you should not believe outright that no ticket is winning. In contrast, given that you can see – to some sufficient degree of clarity – that there is a dog, you have sufficient evidence that there is a dog, and you can outright believe this proposition.

Therefore, on the basis of the distinction between degrees of belief and degrees of outright belief, the evidentialist may say that an evidential probability of one is necessary and sufficient to justify an outright belief, without conceding that it is rational to take the high-stakes bet above. Indeed, the evidentialist may say that an evidential probability of one is necessary and sufficient for outright belief, although it is not sufficient for maximally high outright belief, or maximal credence.

If this is correct, the worry according to which an evidential probability of one is sometimes insufficient to justify an outright belief would be based in great part on a confusion between outright belief and credence one. There are cases in which p has an evidential probability of one for S even if S is not justified in having credence one that p . But, on the view under consideration, such cases are cases in which an outright belief that p is still justified for S .

In brief, an infallibilist might want to avoid the threshold problem by endorsing Williamson's probability one infallibilism. To avoid the threat of scepticism, she may say that an evidential probability of one does not justify credence one, but merely an outright belief, where an outright belief that p is defined in term of willingness to use p as a premise in practical reasoning (in particular when the stakes are low).

Still, at least two problems can be raised for this approach. They concern the fact that an evidential probability of one does not always seem sufficient to justify an outright belief.

First, it is clear that it would be irrational to have a credence of one in all logical truths. But, similarly, it would be irrational to have an outright belief in all logical truths. Recall that, on the present view, when you outright believe that p , you are willing to use p in a possible practical reasoning (at least if the stakes are low). Now, it is uncontroversial that there are logical truths which are so complex that it would be epistemically irrational to use them as premises in a practical reasoning (even if the associated stakes are low). By assumption, these propositions have an evidential probability of one for us. Still, they are not justified enough for us to believe them outright.

Second, why should we think that an evidential probability of one is sufficient to justify an outright belief? It is intuitive that an outright belief is justified only if the confidence involved in this belief is reliably based. In general, one is justified in relying on p on the basis of E only if E is a sufficiently reliable ground with respect to p , in the sense that it does not lead us to misplace our

confidence in p . For the sake of argument, let us assume, following Williamson (2000), that there is a necessary condition on reliable degrees of confidence of the following kind:

(R) If one has some positive degree of confidence that p , when in fact p , and at a very slightly later time, on a very similar basis, one has a very slightly lower and positive degree of confidence that p , when in fact not- p , then one's earlier degree of confidence is not reliable enough to be justified.

The intuitive idea behind this condition is that our degrees of confidence are reliable if they are based in such a way that we cannot have some positive degree of confidence in a false proposition. To illustrate, consider the case of perception. Suppose that you are very close to an animal, and can clearly see that it is a dog. Presumably, your degree of confidence is calibrated so that you will have a very high degree of confidence that this animal is a dog. Now, the more remote the animal, the more difficult to say whether this is a dog or a sheep. Suppose that at some distance d , you are confident that the animal is a dog, and at $d+1$ you also have some positive (but weaker) confidence that the animal is a dog, whereas it is a sheep. One might think that your perceptual system, or the way in which your confidence is calibrated, is unreliable. Indeed, it leads you to have some confidence in false propositions. This means that, at d , you were already too confident that the animal is a dog. If you were not, at $d+1$, you would have no confidence that the animal is a dog.

Now, suppose that an outright belief is justified only if the degree of confidence necessary for outright believing is reliably based, that is, only if one does not have an outright belief *just* because one is overconfident.²² The fact that the probability of a proposition is one given E provides sufficient justification for this proposition only if believing p outright on the basis of E guarantee that one does not have an outright belief just because one is overconfident. In other words, if one can show that whenever one outright believes p on the basis of E and the probability of p given E is one, the outright belief is reliably based, then one can show that an evidential probability of one is sufficient to justify a belief. Otherwise, it is unclear why an evidential probability of one is always sufficient for epistemic justification of outright belief.

However, it does not seem that having entailing evidence is sufficient for being in position to have a reliably based belief. Indeed, the fact that you have entailing evidence for p and that you outright believe that p on this basis does not guarantee that, at a very slightly later time, on a very similar basis, you will not have a positive degree of confidence that p , when in fact not- p . Suppose that E entails p , but there is a very similar evidence E^* that does not entail p , and your powers of discrimination are such that were you having E^* you would believe p on the basis of E^* even if not- p were the case. In such a case, your confidence in p is not reliably based in E .

²²One may be overconfident but still have a reliably based outright belief, if one is far from the case of error. Your overconfidence implies that at some point you will have an outright belief while you should merely have some weaker positive degree of confidence. At this point only your outright belief will be considered as unreliably based.

For instance, suppose that there is a deductive reasoning from E to p , but if the premises were very similar but slightly different, one would still place confidence in the conclusion even if it is false, because we would not discriminate the two premises. In such a case the conclusion is not reliably believed, although the premise entails the conclusion, and the conclusion has a probability of one given the evidence. Therefore, that E entails p does not guarantee that if one believes outright that p on the basis of E , one reliably outright believe that p . Therefore, that does not guarantee that one is justified in believing that p outright.

We might think that the fact that an evidential probability of one cannot guarantee that the subject can reliably form her belief on the basis of E is not problematic for the evidentialist who maintains that an evidential probability of one is necessary and sufficient for epistemic justification of belief. Indeed, the reliability condition does not concern whether a proposition is justified by the evidence – propositional or *ex ante* justification – but whether the subject is justified in believing this proposition – doxastic or *ex post* justification. By saying that an evidential probability of one is necessary and sufficient for knowledge-level justification, we do not say that in any case in which a subject believes a proposition that has an evidential probability of one for her, she is justified. Indeed, the subject may base her belief on wrong reasons, or she may base her belief on right reasons, but in an unreliable way. In such cases, while her belief is justified – the proposition has an evidential probability of one for her – she is not justified in believing it.

However, if we must distinguish propositional and doxastic justification, and evidential probability one infallibilism is only a claim about propositional justification, then a threshold problem remains for doxastic justification. For we may ask what degree of reliability is necessary and sufficient for doxastic justification. If so, evidential probability one infallibilism is insufficient to solve the general threshold problem.

Let us take stock. I have been arguing that the evidentialist might want to reply to ARGUMENT FROM INCONCLUSIVE EVIDENCE (RESULT) by assuming that the level of evidence required to justify a belief is the one conferring a probability of one to the target proposition. I have considered two arguments against this approach. First, there are propositions which have an evidential probability of one for us but that it would be intuitively irrational to believe outright, such as complex logical truths. Second, this approach is focused on propositional justification. Hence, it is insufficient, for there is also a threshold problem for doxastic justification.

3.4 Inconclusive evidence and pragmatic factors

So far, I have considered the threshold problem, and two approaches that an intellectualist might want to take. I have argued that these approaches are unsatisfactory. There might be other intellectualist proposals, but at this stage, we might want to try a totally different approach.

Basically, the threshold problem is that it often seems rational to believe something with less

than maximal justification, although if we restrict ourselves to epistemic factors, there is no obvious non-arbitrary way of fixing the threshold. We may then think that factors of a different kind explain why the threshold is fixed in such or such way. A natural suggestion is to say that these factors are practical factors. Thus, if we embrace the pragmatic encroachment thesis, we might have a way of solving the threshold problem. This suggestion is based on the idea that there is a connection between epistemic justification of belief, or knowledge, and the rationality of action. The idea is that a belief is epistemically justified for a subject *S* (if and) only if *S*'s evidence is sufficient to rationalise *S*'s action on the believed proposition. On this approach, whether a belief in *p* is justified by evidence *E* in a certain situation depends on whether, given *E*, an action based on *p* is justified in this situation.

This promises to solve the threshold problem because, given *S*'s practical situation and *S*'s evidence for *p*, it can be determined without arbitrariness whether *S* should act as if *p*. We may use, for example, the expected utility calculus. Provided that *p* is relevant to *S*'s action, in the sense that whether or not *p* is the case makes a difference to what is best for *S* to do, we can see whether the option that has the best expected utility for *S* is the option in which *p* is assumed. If what it is rational to do in fact is what it is rational to do given *p*, and what it is rational to do given *p* is what it is rational to do in fact, then it is rational to act as if *p*.²³ Hence, according to this approach, it can be determined without arbitrariness, by using the expected utility calculus, what the evidential threshold for epistemically justified beliefs is given a certain practical situation.

If this approach is correct, we should abandon evidentialism and, more generally, intellectualism. Thus, it is important to see whether this approach can give a satisfactory answer to the threshold problem. First, I consider two cases presented by Owens, and I argue that they are not decisive to show that practical factors affect the threshold for justified beliefs. I leave for the next chapters the discussion of other cases that may seem more compelling. Second, and more importantly, I argue that whatever we think of these kinds of cases, the main way in which they have been explained – by appealing to a connection between justified belief and rational action – fails to provide a satisfactory solution to the threshold problem.

3.4.1 Practical factors and suspension of judgement

Owens put forth cases purporting to show that pragmatic factors can directly affect whether one is justified in believing *p*. Consider the following cases:

(NON-EXPERT / EXPERT) I am wondering whether to purchase a house this year and that decision depends largely on whether I think prices will rise. I carefully read the property pages listen to the pundits on television, determine that the house market will remain flat and plan an expensive holiday instead of a house purchase. Just before I

²³See Fantl and McGrath (2002).

form this view, a newspaper article by a respected economist appears which purports to show that house prices will rise. I don't read the article, and I don't let it worry me. Having insufficient time to try to assimilate and weigh his testimony against that of other experts, I stick to my view that prices will not rise. And this, we may suppose, is a perfectly reasonable line for me to take and one that I can explicitly avow to myself or to others. Were I an economist specialising in the housing market and about to give a paper on future price movements to a conference of estate agents, I probably won't be entitled to have a belief on the matter that did not take account of this article. (Owens, 2000, 26)

(KILLED BOY) Parents of a boy reported killed in action suspend their judgement on the matter beyond the point at which a disinterested bystander could reasonably conclude that their son was dead. Since it matters so much to them to be right about whether their son is dead, are they not entitled to ask for more evidence than the disinterested bystander before coming to a conclusion? (Owens, 2000, 30-31)

In the first case, the non-expert is supposed to be justified in believing that the house market will remain flat. Hence, in his situation, the evidence is supposed to be sufficient to justify his belief. In contrast, while he has the same evidence, the expert is supposed to be rational in suspending judgement.

In the second case, the parents are supposed to be rational in suspending judgement, while the disinterested bystander is taken to be justified in believing the target proposition. According to Owens, this shows that pragmatic considerations determine whether a belief is justified.

Other cases have been proposed in the literature. Most of the time, they consist of pairs of cases featuring two subjects having the same epistemic position with respect to a proposition, but facing different practical decisions with different stakes. In the low stakes situation, the subject is supposed to be in a good enough epistemic position to believe the proposition and to act on it, whereas in the high stakes situation, the subject is supposed not to be in a good enough epistemic position to believe the proposition and to act on it. These cases will be discussed more in details in the following chapters. Here, let me just note that Owens' cases are not decisive.

Consider (EXPERT / NON-EXPERT) and suppose that pragmatic factors affect the justification of the belief. If so, were the expert to believe that the house market will remain flat, his belief would be unjustified. But it is unclear that such a belief would be unjustified. For instance, if asked about his view on the matter, he could say that although he should read this article, he believes that the house market will remain flat. For sure, it may be irrational for him to give a conference or write a paper without reading this article (at least if he has enough time to read it). But we may explain this by the fact that, as an expert, this is his duty to read all the relevant papers before writing articles or giving conferences. It is unclear that this is relevant to what he is entitled to

believe.

We might be tempted to reply that although the expert is entitled to have a certain degree of belief on the house market, he is not entitled to have an outright belief. In particular, he is not entitled to claim to know. But, when we have a justified belief that p , we might think, we are entitled to claim to know that p .²⁴ The problem with this reply, though, is that under this understanding of the notion of belief, it also seems that the non-expert is unjustified in believing that the market will remain flat. It does not seem that it is reasonable for the non-expert to claim to know, for there is strong available counter evidence.

Let me now consider (KILLED BOY). About this case, Owens writes:

Such suspension of belief is not just understandable but epistemically unexceptionable. The parents fall into irrationality only when their need and interests determine what conclusion they arrive at, rather than whether they draw a conclusion at all. (Owens, 2000, 31)

While we may agree that it is not irrational for the parents to suspend judgement, that does not show that the importance of the matter raises the level of evidence required for them to have a justified belief. Is the suspension of judgement motivated by the fact that there is not enough evidence to justify the belief that their son is dead? Or is it motivated by something different, for example, by the fact that the parents want to consider more deeply the evidence before taking it at face value?

Indeed, there are (at least) two ways in which a suspension of judgement may be rational. First it may be rational because, while one is motivated to have a view on whether p , the evidence is insufficient to justify a belief that p rather than not- p . But, second, we might also think that it can be rational to suspend judgement – in the sense of not forming a view – because, while there is sufficient evidence to justify the belief, it is rational to pause and reflect on whether the evidence is sufficient.

We might think that the second way of suspending judgement has to do with the freedom of judgement given by the capacity to reflect on the reasons we have. Thus, for example, Korsgaard writes:

A lower animal's attention is fixed on the world. Its perceptions are its beliefs and its desires are its will. It is engaged in conscious activities but it is not conscious of them. That is, they are not the objects of its attention. But we human animals turn our attention onto our perceptions and desires themselves, onto our own mental activities, and we are conscious of them. That is why we can think about them. And this sets us a problem no other animal has. It is the problem of the normative. For our capacity to turn our attention onto our own mental activities is also a capacity to distance ourselves

²⁴See Owens (2000, ch. 3).

from them, and to call them into question. I perceive, and I find myself with a powerful impulse to believe. But I back up and bring that impulse into view and then I have a certain distance. Now the impulse doesn't dominate me and now I have a problem. Shall I believe? Is this perception really a reason to believe? I desire and I find myself with a powerful impulse to act. But I back up and bring that impulse into view and then I have a certain distance. Now the impulse doesn't dominate me and now I have a problem. Shall I act? Is this desire really a reason to act? The reflective mind cannot settle for perception and desire, not just as such. It needs a reason. Otherwise, at least as long as it reflects, it cannot commit itself or go forward." (Korsgaard, 1996, 92-93)

Assuming that we have such a freedom of judgement, it might be argued that, in (KILLED BOY), it is rational for the parents to suspend their judgement because, given the importance of the issue, it is rational for them to pause and to consider more deeply the available evidence. But note that, were they to believe that their boy is dead on the basis of the report, it is far from clear that we would count their belief as epistemically unjustified. It is clear that in such a case, they should not be blamed for their gullibility. If they believe on the basis of this report that their son is dead, it does not seem right to object to their belief by saying: "Look, of course you have been reported that your boy is dead. But this is not sufficient. You should not believe so easily what you are told. After all, it really matters whether your son is really dead". Thus, we might argue that what justifies the suspension of judgement is not the fact that a belief would be epistemically unjustified, or that the evidence is insufficient.

Also relevant here is the claim defended in the previous chapter, according to which there are no unconditional or absolute *epistemic* obligation to believe. If this claim is correct, we can explain why it is epistemically *permissible* for the parents to suspend their judgements, although it is also epistemically *permissible* for them to believe that their son is dead.

Of course, more can and should be said about these types of cases, and it is the purpose of the following chapters to discuss them in more details. For the time being, let us just note that these two cases are far from decisive to show that the evidential threshold for justified belief is (partially) determined by the subject's practical situation.

3.4.2 The connection between justified belief and rational action

Whatever we think of these particular cases, does adopting the idea that pragmatic considerations partially determine, in addition to inconclusive evidence, whether a belief is justified, solve the threshold problem? To show that it does, we need to show that there is a non-arbitrary rule that gives us, given the practical situation, the level of evidence required to justify a belief, or to know.

According to Fantl and McGrath, there is such a rule, which arises from the fact that there is the following pragmatic condition on epistemic justification (and knowledge):

(PCA) S is justified in believing p only if S is rational to act as if p. (Fantl and McGrath, 2002, 78).

What does ‘act as if p’ mean? Suppose that if it is raining, the best thing to do is to take your umbrella, but if it is not raining, the best thing to do is not to take your umbrella. If you are justified in believing that it is raining, they argue, you should act as if it is raining, that is, you should take your umbrella. If you should not take your umbrella, that means that you are not justified in believing that it is raining.

Fantl and McGrath argue in favour of (PCA) by arguing, first, that knowledge is sufficient for rational action (if S knows that p, then it should not be a problem for S to act as if p), and, second, that if knowledge is sufficient for rational action, epistemic justification, understood as knowledge-level justification, is also sufficient for rational action.

I will not dispute these claims here. But can (PCA) solve the threshold problem? (PCA) only gives a necessary condition on epistemic justification. According to (PCA), when it is not rational for S to act as if p, then p is not epistemically justified for S. But there are cases in which the degree of evidence for p is very weak, say, less than 0.5, although given the stakes, it is rational to act as if p. For example, suppose that there is little chance that it will rain, say, one chance over ten. However, if it rains you must absolutely have your umbrella. On the other hand, if it does not rain, the best thing would be not to have your umbrella, but this is not that important. In this situation, you should act as if it is raining, by taking your umbrella. But, of course, you are not justified in believing that it is raining. You might even say “I do not believe that it will rain, but I take my umbrella, just in case”. Thus, being rational to act as if p does not provide a sufficient condition for epistemic justification to believe.

However, if being rational to act as if p does not provide a sufficient condition for epistemic justification to believe, the threshold problem remains. What is the lowest possible level of evidence sufficient for epistemic justification to believe? This approach gives no answer to this question, because the evidence can be strong enough for S to be rational to act as if p but too weak for S to be rational in believing that p.

Fantl and McGrath (2009) propose a slightly different approach, in terms of having reasons to act. On this approach, a belief that p is epistemically justified for you if and only if p is warranted enough to justify your action. To say that p is warranted enough to justify your action is to say that if p does not justify your action, then the reason why p does not justify your action is not that the un-eliminated epistemic possibility that not-p is relevant and should be taken into account in your decision. In brief, p is warranted enough to be a reason you have if and only if the epistemic possibility that not-p is idle. So, they argue that a belief that p is epistemically justified for S if and only if the possibility that not-p is idle given S’s practical situation.

Alas, this approach does not solve the threshold problem either. Indeed, a threshold problem can be raised for when an epistemic chance should count as “idle” given the practical situation.

Is there a non-arbitrary rule telling us what the lowest possible degree of evidence or justification sufficient to exclude as idle an epistemic chance of error in the practical decision is? It does not seem so. Indeed, there are cases in which you should act as if p even if the chance that not- p is not idle. Again, consider the case in which you should take your umbrella even if it is very unlikely that it will rain. In such cases, you should act on the proposition that there is some chance that it will rain, rather than on the proposition that it will rain. That is, what can justify your action in such situations is the fact that p is probable enough for you, not the fact that p . But there is no non-arbitrary rule telling us when the fact that p , rather than the fact that p is probable enough for you, can justify your action. As a result, on this approach, there is no non-arbitrary rule telling us when the belief that p is epistemically justified for a subject.

Another problem, which is raised by Brown (2013), has to do with impractical propositions. There are plenty of propositions that we believe but that are unrelated to our practical situation. How are we going to fix the threshold for these propositions? If we say that the threshold is fixed by the standards applying to the proposition associated with the highest stakes faced by the subject in the situation, then a problematic kind of scepticism threatens. We often face high stakes decisions. If, in these situations, the threshold is very high for any proposition we consider, then we know very little in these situations. On the other hand, if we say that the threshold is fixed by the practical situation associated with each proposition, then this solution can have no application for impractical propositions.

We might think that there is a minimal evidential threshold which must be satisfied for any proposition, and that this threshold provides a sufficient condition for impractical propositions. However, we have seen that there is no non-arbitrary way of fixing this threshold, for the pragmatic condition is either too weak – it can be rational to act as if p even if it is not rational to believe that p – or a threshold problem can be raised regarding when we satisfy this pragmatic condition – how are we going to tell, in non-arbitrary way, when a proposition is warranted enough to justify action?

Finally, can't we imagine a being without desires and without preferences, unable to act, so that he can have no practical reason? Intuitively, the beliefs of such a being could still be assessed as rational or irrational. But this is surprising if the threshold for justification of beliefs is fixed by the subject's practical situation. By assumption, there is no practical situation to which this being is related.

In sum, it is unclear how appealing to practical factors can help to solve the threshold problem. This suggests that the third premise of ARGUMENT FROM INCONCLUSIVE EVIDENCE (RESULT) is incorrect. Even if we assume that inconclusive evidence alone cannot justify a belief, this does not suggest that practical factors are relevant to the justification of our beliefs.

3.5 Conclusion

In this chapter, I have considered the threshold problem for intellectualism (and epistemic evidentialism). This problem is that of determining in purely intellectualist terms what level of justification is required to epistemically justify a belief. I have considered two ways in which an intellectualist might want to address the threshold problem, and argued that they are unsatisfactory. Then, I turned to the idea that the threshold problem can be solved if we reject intellectualism and embrace pragmatic encroachment – the thesis that practical factors are required for epistemic justification. However, I have shown that endorsing pragmatic encroachment does not solve the problem either. Therefore, it can be concluded that the threshold problem does not place a specific threat on the orthodox account of knowledge, and hence, that it does not favour pragmatic encroachment over intellectualism.

Part II

'Knowledge' ascriptions

Chapter 4

Linguistic intuitions and the rule of relevance

4.1 Introduction

In the previous chapters, I have argued that the threshold problem should not lead us to endorse a shifty view of knowledge. In this chapter, I start considering another crucial argument in favour of a shifty account of knowledge. According to this argument, intuitively, a mere change in truth-irrelevant factors may affect the truth-value of a ‘knowledge’ attribution. Therefore, it is said, we should abandon the orthodox way of thinking about knowledge.

I discuss a popular alternative explanation of these intuitions appealing to pragmatics. According to such an approach, what is affected is not the truth-value of the ‘knowledge’ ascription, but its pragmatic appropriateness. More specifically, this approach appeals to Grice’s conversational rule of relevance and the distinction between semantic and pragmatic meaning. Briefly stated, the idea is that, in virtue of the maxim of relevance, saying a truth – e.g., that the subject knows the target proposition – can be misleading, whereas saying a falsehood – e.g., that she does not know the target proposition – can pragmatically communicate something true.

I consider five arguments against this particular approach. First, I consider DeRose’s argument according to which the (supposedly) false sentence under consideration does not lack relevance in the context in question. Hence, the rule of relevance cannot be used to explain its appropriateness and the account is *ad hoc*. I argue, however, that this argument is not decisive.

Second, I offer an argument showing that, even assuming that the rule of relevance can play a role, it is not sufficient to explain the supposed pragmatic meaning of the utterance under consideration. Given the background assumptions of this approach and Grice’s conversational maxim of

quantity (assert the stronger!), indeed, it is surprising that the hearer infer the supposed pragmatic meaning. So, either there is something wrong in the background assumptions, or the explanation is still insufficient and somewhat *ad hoc*.

Third, I explain that either this approach must assume the controversial claim that knowledge is not closed under known entailment, or we must recognize that there is more ad hocness than what we might have thought initially. Thus, this approach is not available to philosophers endorsing the principle of epistemic closure.

Fourth, I put forward that by focusing on salience, the proposal under-describes the conversational context, and hence, does not do well in identifying and explaining what is or is not conversationally relevant, and why. I consider connecting the rule of relevance to the practical context. I show that this refined approach is incompatible with the claim that knowledge is the norm of action, and that it should focus on a different speech act in the conversation.

Finally, I argue that this approach is at odd with two linguistic and conversational facts that are left without a satisfactory explanation. First, the utterance “I know that p” is *never* considered as conversationally inappropriate when the question is to make a practical decision. Second, the supposed pragmatic implication is not cancellable.

I conclude that this specific Gricean manoeuvre fails to provide a satisfactory explanation of our intuitions in an orthodox framework.

4.2 The bank cases and the rule of relevance

The orthodox account of knowledge is intellectualist and invariantist. According to this view, only truth-relevant factors can turn a belief into knowledge, and the relevant epistemic standards do not vary with the attributor’s conversational context.

Against this orthodox view, it has been argued that the truth-value of ‘knowledge’ attributions (utterances of sentences of the form ‘S knows that p’) are sensitive to practical factors. For epistemic contextualists, this is due to the fact that ‘know’ is a context-sensitive term.¹ For subject-sensitive invariantists, this is due to the fact that there is a pragmatic condition on knowledge that must be satisfied in the subject’s practical environment.²

This line of thought rests in great part upon pairs of cases showing that our ordinary ‘knowledge’ attributions vary according to the stakes and the possibilities of error that are mentioned. The most popular pair of cases is DeRose’s bank cases:³

¹Cohen (1988); Lewis (1996); DeRose (2009)

²Hawthorne (2004); Stanley (2005b); Fantl and McGrath (2009)

³See DeRose (2009). In order to decide between contextualism and subject-sensitive invariantism, it is useful to consider cases where the subject of the ‘knowledge’ ascription and the ascriber are different. If epistemic contextualism is true, this is the attributor’s conversational context that is relevant to assess the truth-value of the ‘knowledge’ attribution. If subject-sensitive invariantism is true, this is the subject’s practical situation that is relevant to assess this truth-value. Note, however, that the attributor’s conversational context may, in some way or other, be

Bank Case A (Low Stakes). My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.” I reply, “No, I know it’ll be open. I was just there two weeks ago on Saturday. It’s open until noon.”

Bank Case B (High Stakes). My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, “Banks do change their hours. Do you know that the bank will be open tomorrow?” Remaining as confident as I was before that the bank will be open then, still, I reply, “Well, no. I’d better go in and make sure.”

What Keith says in case A and B seems true. The only difference between the two cases is that the stakes are high in case B (“A large cheque has been written”) and a further possibility of error has been raised (“Banks do change their hours”).

This is a problem for the orthodox account. According to this view, since Keith is in the same epistemic position in both cases and the epistemic standards of knowledge do not shift, he knows in case A if and only if he knows in case B. Most epistemologists are also non-sceptics, and, on this basis, maintain that Keith knows in case A.⁴ Thus, according to the orthodox view, it is true that Keith knows in case A and in case B. However, what Keith says in case B, namely “I do not know”, seems true. If so, the orthodox account is in contradiction with our intuitions.

The orthodox philosopher is bound to deny that what Keith says in case B is true. Of course, he may say, what Keith says is appropriate. But does that mean that what he says is true? It can sometimes be appropriate, pragmatically, to say falsehoods. For instance, if asked by an automobilist whether there is a gas station in the area, you may well say “No, there isn’t”, if you know the only gas station is closed at that hour. Presumably, the automobilist is looking for an open gas station, and it would be inappropriate, although true, to reply: “Yes, there is one just on

determined by the practical situation of the subject under evaluation.

⁴Hereafter, I will put aside the sceptical approach to the problem.

your right after the next stop” if you know that it is closed.⁵

The truth-value of the semantic content of the uttered sentence (its literal meaning) is one thing. Quite another thing is the truth-value of the proposition that is communicated by the utterance of such a sentence. And, arguably, whether or not the utterance of a sentence is appropriate depends on these two things. For instance, suppose that John has three children. If I say “John has two children”, what I say is literally true. However, it does not seem inappropriate to say that, because that would be misleading. Due to a conversational rule of quantity (assert the stronger!), what I pragmatically communicate is that John has two children only, and therefore not three.

So, perhaps we can explain what happens in case B by saying that, although it’s true that Keith knows, it would be inappropriate for him to say it. Perhaps, given the conversational context, this would communicate something false. Further, perhaps it’s appropriate in this context to say something false, if this communicates something true.

To defend the legitimacy of this warranted assertability manoeuvre, it must be shown that there is a conversational rule operative in case B which explains that saying a truth (“I know”) would in this context prompt an implicature that is wrong in terms of conversational rationality, whereas saying something false (“I do not know”) prompts an implicature that is right in terms of conversational rationality. Otherwise, the manoeuvre is *ad hoc*.⁶

So far, the most influential account along these lines is Rysiew’s (2001; 2005). It has been improved by Brown (2006). That is why I will mainly focus on these views. The basic idea is that what Keith says is appropriate due to a conversational implicature based on the rule of relation (be relevant!). According to this view, on the assumption that Keith makes a relevant speech act, he addresses the worry raised by his wife. Therefore, it is natural to interpret what he says in light of this worry in the following way. In uttering “I do not know that p”, Keith implicates that he is not in a position to eliminate his wife’s salient alternative, which is true. By contrast, in saying the literal truth, that he knows, he would implicate that he is in a position to eliminate this alternative, which is false.

This approach rests upon the idea that in general speakers do not make the distinction between (or are not interested in) the semantic (or literal) meaning and the pragmatic meaning. If there is a confusion between the semantic and pragmatic meaning, one can explain our intuitions about the alleged variation in truth-value in the bank cases. In case A, what Keith says is literally true, and what is pragmatically communicated is identical to what is literally said. In case B, what Keith says is literally false, but what is pragmatically communicated is different, and true. Therefore, the intuition that it is true that Keith does not know in case B may rely on a confusion between the semantic and pragmatic meaning of Keith’s utterance.⁷

⁵See also Grice (1989, 32).

⁶DeRose (2009, 88-89)

⁷The idea that speakers are not always aware of, or interested in, the distinction between semantic content and pragmatic implicatures is disputable, but I will not consider this issue here.

4.3 The account

In order to see whether this Gricean manoeuvre is convincing, it is necessary to see how it works in more details.

First, what does a ‘knowledge’ attribution literally express? Rysiew adopts a relevant alternative theory about knowledge, even though his general view is compatible with other accounts. According to such a view, an utterance of ‘S knows that p’ literally expresses the proposition that S has a true belief that p, and that S’s epistemic position is good enough to eliminate all the alternatives to p that are relevant. An alternative is relevant in this epistemic sense if and only if it is knowledge-destroying.⁸ To illustrate, consider Henry in the fake barn county.⁹ He is driving and he tries to identify barns. Unbeknownst to him there are full of papier mâché facsimiles of barns in his vicinity. Given Henry’s environment, the fact that what Henry sees is a fake barn rather than a real barn is a knowledge-destroying alternative. Now, a subject is in a good enough epistemic position to eliminate a knowledge-destroying alternative if and only if she can discriminate its truth from its falsity. This means, roughly, that there is no *relevant* not-p world where the subject believes that p (on the same basis). Consider Henry. As a matter of fact, Henry is in front of the single real barn, and he forms the belief that it is a barn. His belief is true, but if what he is looking at were a fake barn, Henry would still believe that it is a barn (on the basis of the same perceptual experience). Therefore, Henry is not in a position to eliminate this alternative. As a result, since this alternative is knowledge-destroying, Henry does not know that the barn he is looking at is a barn.

According to Rysiew’s framework, in particular, a possibility of error is knowledge-destroying if its likelihood is above a certain threshold. This threshold does not vary depending on conversational and practical factors (such as what is salient in the conversation, what is at stake, etc.) but is set independently. That is why this account is intellectualist and invariantist. The degree of likelihood for a possibility of error is set by what a normal human would think of it in the circumstances in question. Thus, the set of knowledge-destroying alternatives does not vary across contexts.

Rysiew distinguishes these knowledge-destroying possibilities from possibilities of error that are merely salient. A possibility of error is salient if it has been raised in the conversational context. A possibility of error is salient and not knowledge-destroying if and only if it has been raised in the conversational context but does not prevent a subject from knowing, even if she is not in a position to eliminate it. The set of salient alternatives is a variable set.

As regards pragmatic meaning, it is standardly assumed that what an utterance pragmatically imparts depends on the conversational context. Let us suppose that there is a feature, X, that can roughly specify the conversational context, and that is determined by the goal or direction of the

⁸Rysiew uses the expression ‘relevant alternatives’, but to avoid ambiguities, I will use ‘knowledge-destroying alternatives’.

⁹See Goldman (1976).

conversation, what is salient, etc.¹⁰ Whether or not a speech act is appropriate in conversational terms depends on X. If the conversational participants are governed by a cooperative principle, they infer from what is literally said what best rationalizes others' speech acts given X. Therefore, what an utterance pragmatically imparts depends on X. In sum, given the variety of contexts, one can use an expression to communicate many different things, and what is communicated by a use of this expression in a given context depends on the semantic content of the expression in question, along with the general rules of conversation.

'Knowledge' ascriptions should not be different in this respect. Given the semantic content of the expression 'S knows that p' (namely, that S is in a quite good epistemic position with respect to p), it is natural to think that, in some conversational contexts, uttering it pragmatically imparts the proposition that S is in a good enough epistemic position to satisfy the epistemic standards required by X, at least if in these contexts this is what best rationalizes this utterance. In short, the idea is that X, the goal of the conversation, the interests of the speakers, etc., can determine some relevant epistemic standards in a context C, and an utterance of 'S knows' in C can pragmatically impart that S meets these epistemic standards.

On this basis, how can we explain that a 'knowledge' ascription would be inappropriate in the high stakes bank case, while a 'knowledge' denial is appropriate? Consider the conversational context first. In this context, a possibility of error has been raised. Next, consider the rules of conversation. An utterance is conversationally relevant in a context only if it takes the ingredients of the context into account and hence, in this context, only if it takes this possibility of error into account.

What would a 'knowledge' ascription communicate with respect to this possibility of error? As the semantic content of a 'knowledge' attribution has to do with the goodness of the subject's epistemic position, it is plausible to think that, with respect to the possibility of error in question, such an attribution would communicate that the subject is in a good enough epistemic position to eliminate it, which seems false in Keith's case. So, to avoid communicating this, he should refrain to say that he knows.

In contrast, what does a 'knowledge' denial communicate with respect to this possibility of error? As the semantic content of a 'knowledge' denial has to do with the badness of the subject's epistemic position, it is plausible to think that, with respect to the possibility of error in question, such a denial communicates that the subject is not in a good enough epistemic position to eliminate it, which seems true, in Keith's case. So, to communicate this, he should say that he does not know.

This explanation relies on the idea that when Keith utters "I do not know", the semantic content of the sentence uttered in this context lacks specific relevance, and that this would also have been the case had Keith uttered "I know that the bank is open". This lack of relevance is explained by the fact that, in this context, X has to do with the salient alternative raised by Keith's wife.

¹⁰One can think of X as the "conversational score". See Lewis (1979).

If the conversational context does not concern knowledge-destroying alternatives, but mere salient alternatives, one needs to find a meaning, different from the literal meaning, so as to make the utterance specifically relevant (so as to make the speech act in conformity with the maxim of relevance).¹¹

In sum, Rysiew's account relies on the following claims:

- (1) In general, an utterance of 'S knows that p' or 'S does not know that p' is specifically relevant in a conversational context only if it means something concerning the assessment of the strength of S's epistemic position with respect to the standards required by X, where X is set by the conversational context C.
- (2) In general, an utterance of 'S knows that p' in C, if it is specifically relevant, means that S's epistemic position is *good enough* with respect to the operative standards in C.
- (3) In general, an utterance of 'S does not know that p' in C, if it is specifically relevant, means that S's epistemic position is *not good enough* with respect to the operative standards in C.
- (4) When C is the context of case B, X is the salient alternative.

Suppose also that:

- (5) Keith is not in a position to eliminate the salient alternative in case B.

Then, by (1) - (5):

- (6) In case B, Keith's utterance of 'I knows that p', if it means something that is specifically relevant with respect to the assessment of Keith's epistemic position concerning the salient alternative, would mean something that is false (namely that Keith can eliminate the salient alternative)
- (7) In case B, Keith's utterance of 'I do not know that p', if it means something that is specifically relevant with respect to the assessment of Keith's epistemic position concerning the salient alternative, means something that is true (namely that Keith cannot eliminate the salient alternative)

Now, suppose that:

- (8) It is conversationally inappropriate to communicate something false and conversationally appropriate to communicate something true.

¹¹See, for example, Rysiew (2001, 491-492): "Whether or not what he [Keith] says is strictly speaking true, if the speaker *didn't* think that S's epistemic position were good enough *in the relevant sense of "good enough"*, whatever that is, he would not say "*S* knows that *p*"; and regardless of the truth-value of the sentence itself, he wouldn't say "*S* doesn't know that *p*" unless he thought there were salient no-*p* possibilities that the speaker could not rule out".

We can conclude:

(9) It is conversationally inappropriate for Keith to utter ‘I knows that p’ in case B.

(10) It is conversationally appropriate for Keith to utter ‘I do not know that p’ in case B.

It is important to insist that this account is supposed not only to explain why it would be improper for Keith to say “I know”, namely that this would prompt a false implicature, but also why it is appropriate to say “I do not know”. This utterance takes the context into account, and communicates something true.

4.4 Objection 1: The rule of relevance cannot be put to work

This Gricean manoeuvre makes crucial use of Grice’s rule of relevance. In using the maxim of relation, it seems, we explain that a hearer is led to seek for a different (and pragmatic) meaning because the semantic meaning of the uttered sentence lacks specific relevance. On this view, in bank case B, the rule of relation could do its explanatory work only if the ‘knowledge’ attribution or denial in this context lacks specific relevance. Thus, if one can show that the denial in case B does not lack specific relevance, what explains its appropriateness cannot be a pragmatic meaning inferred by the application of this rule. Indeed, if the semantic content of the ‘knowledge’ denial is relevant in case B, the hearer cannot be expected to infer a pragmatic meaning from its lack of relevance. Claim (1) above may well be true even though there is no pragmatic implicature at play in case B.

This objection is raised by DeRose (2009, 118-124), who argues that the semantic content of Keith’s utterance “I do not know that p” does not lack specific relevance in the context of case B. Indeed, it seems, if this utterance literally means that Keith is not in a position to eliminate the knowledge-destroying alternative, it (semantically) means that he is not in a position to eliminate the salient alternative. Let us consider DeRose’s analogy. Suppose that we are in a context where ‘tall’ means *at least 6 feet tall*. Suppose that I am interested in whether Sally is tall. If I ask you “Is Sally tall?” and you answer “She is less than 5.5 feet tall”, we can agree that this settles the question in a negative way. Therefore, your answer is not irrelevant, and the hearer is not led to seek for a different (pragmatic) meaning to explain your utterance.

Similarly, one might think, if the fact that you are not in a position to eliminate knowledge-destroying possibilities entails that you are not in a position to eliminate the not knowledge-destroying alternatives to p, then the semantic content of your utterance “I do not know that p” entails that you are not in a position to eliminate the not knowledge-destroying possibilities of error. If there is such a semantic implication, there is no need for a pragmatic inference. Thus, by saying “I don’t know”, Keith literally express something that semantically implies a truth concern-

ing his epistemic position with respect to his wife's possibility of error, and which is thereby fully relevant.

However, a first problem for this objection is that it is unclear why if one is not in a position to eliminate a knowledge-destroying alternative to p (if one does not know that p), then one is not in a position to eliminate a salient, but not knowledge-destroying, alternative to p . On the one hand, perhaps I am in a position to eliminate the salient possibility that there will be a terrorist attack on the bank tonight, which is a not knowledge-destroying alternative regarding the fact that the bank is open tomorrow. On the other hand, perhaps I'm not in a position to eliminate the possibility that the bank has changed its hours, and hence, if this is a knowledge-destroying alternative, to know that the bank is open tomorrow. Therefore, we cannot semantically infer from "S does not know that p " that S is not in a position to eliminate the not knowledge-destroying alternative q .

A second problem is that even assuming that the relevance of what is literally expressed by "I do not know that the bank is open" speaks against the idea that the hearer is led to infer the supposed pragmatic meaning by a lack of relevance, that still does not show that the rule of relevance cannot be used to explain why this (supposedly) false utterance is pragmatically appropriate. Indeed, there are different ways in which the rule of relevance can play its role in explaining why an utterance is appropriate. One is by explaining how a hearer is led to a pragmatic meaning due to a lack of relevance of what is literally expressed by the utterance in question. But there are others.

Consider the case of the automobilist above. He does not pragmatically infer from the irrelevance of what you literally say (namely, that there is no gas station in the area) what he needs to know (namely, that there is no open gas station in the area). Indeed, this is semantically implied by what you literally said. However, what explains why it was appropriate for you to say it, despite its falsity, is still the relevance of saying it. By literally saying this falsehood, the automobilist gets the information he needs.

The rule of relation can be used to explain why what Keith says is appropriate in the same way. It may be that Keith's wife is not led to the right information by an inference from the lack of relevance of his utterance, but by a semantic inference from the content literally expressed. This is the relevance of the information she thereby gets that explains why saying this falsehood is relevant.

According to DeRose, the problem here is that even if this utterance has a true (semantic) implication, this utterance is supposedly false. Hence, it is actually misleading. Since the hearer cannot appeal to a lack of relevance for working out the implicature (as this utterance does not lack specific relevance), and since the proposition expressed is false, the utterance violates the maxim of quality. One might think that it is incorrect to say something false to make someone reach the right conclusion. Therefore, the invariantist intellectualist cannot explain why Keith's utterance seems appropriate in this way.¹²

¹²A possible reply would be to say that it is common knowledge that S does not know, so that the hearer would not be misled. However, this is not a manoeuvre that uses the maxim of relation, but the maxim of quality. In

However, it is often acceptable to tell white lies when they are efficient, and hence, relevant. To illustrate, let's modify Sally's case. Suppose that she is more than 5.5 feet tall, but less than 6 feet tall. Is it inappropriate to falsely say "She is less than 5.5 feet tall" to make the hearer reach the right conclusion that she is not tall? I must say that this does not seem wrong to me. There are many cases in real life where one tells white lies, and this is acceptable insofar as this is efficient.¹³ As a result, it does not seem to me that DeRose's criticism is sufficient to undermine Rysiew's approach.

4.5 Objection 2: The rule of relevance is insufficient, and hence, the account *ad hoc*

A second objection that can be raised against Rysiew's account concerns whether the rule of relevance is sufficient to explain the appropriateness of "I do not know that p" and the inappropriateness of "I know that p" in the context in question. Doubts can be raised in particular by focusing on Grice's rule of quantity. If we accept Rysiew's background assumptions and we apply this rule of quantity, one should not expect the hearer to be led to the pragmatic meaning supposed by the theory under consideration.

Consider the claims (2) and (3) above. Why is it natural to think that, when relevant, a 'knowledge' attribution pragmatically communicates that the subject's epistemic position with respect to the operative standards in the context is "good enough", and a knowledge denial communicates that the subject's epistemic position is "not good enough"?¹⁴ Either there is a pragmatic explanation for this, or there is not. If there is not, the account is still *ad hoc*. Indeed, the rule of relevance can only explain why, when saying "I know that p" (or "I do not know that p"), this expression means something regarding the salient alternatives. But it does not explain the (alleged) pragmatic implication from "I know that p" to "I am in a *good enough* epistemic position with respect to X".

Indeed, at least three assumptions are required in this Gricean manoeuvre. Firstly, it must be assumed that "I know that p" lacks specific relevance (alternatively, semantically implies something relevant), secondly, that the pragmatic meaning of an utterance of 'S knows that p' is that S is in a good enough position with respect to X, and, thirdly, that X includes responding to the salient

addition, it is possible to figure out cases where this common knowledge is absent, as DeRose notes (DeRose, 2009, 122-123).

¹³See Fantl and McGrath (2009, 41-42) for further discussion. They insist that it is difficult to explain why a white lie is acceptable in case B. They also raise another worry concerning the first-person perspective. According to them, in case B Keith will be willing to say to himself that he does not know. This cannot easily be explained by the Gricean manoeuvre, even assuming that one easily confuses the semantic and pragmatic meanings. However, if Keith knows in case B, it is unclear why he would be willing to say to himself that he does not know.

¹⁴See Rysiew (2001, 487-488) and Rysiew (2005, 48): "[I]t's because speakers strive to conform, and are known to so strive, to the maxim of relation ("Be relevant") that, in uttering a sentence of the form, 'S knows that p', the speaker is naturally taken to intend/mean that S's epistemic position with respect to p is 'good enough' given the epistemic standards that are operative in the context in question." See also Rysiew 2007 638 639

alternative. The rule of relevance is supposed to explain on the basis of the third assumption the truth of the first one. It cannot be used to explain the second assumption.

Consider what Rysiew says on this issue:

For it is only if speakers are understood to be intending to communicate information about how the subject fares vis-à-vis the contextually operative standard(s) — hence, i.e., about whether the subject can or cannot rule out any contextually salient alternatives — that they can be seen as striving to be maximally relevantly informative (hence, as conforming to CP) (2005, 48).¹⁵

This speak in favour of claim (1) above. But how, from this, can we get to claims (2) and (3) above? There is a difference between communicating information about the value of a subject’s epistemic position with respect to an operative epistemic standard in a context (in Keith’s case, a standard which is a function of a salient alternative), and communicating that the subject’s epistemic position is “good enough” or “bad enough” with respect to the operative epistemic standards. I do not deny that it might well be that, in fact, the hearer is led to this pragmatic meaning. What I deny is that the rule of relevance explains this (putative) pragmatic inference.

The fact that this pragmatic inference is not based on the rule of relevance may lead to think that Rysiew just *assumes* that there is a pragmatic meaning different from the semantic meaning of ‘knowledge’ ascriptions, which of course would make his account *ad hoc*.¹⁶ Rysiew replies to this possible objection when he notes that, in his account, it is not assumed that there is pragmatic meaning besides the semantic meaning of a ‘knowledge’ ascription. Rather, he says, this is in virtue of the semantic meaning of the ‘knowledge’ ascription and the general rules of conversation, as well as the background situation, that the hearer is led to the pragmatic meaning in question.¹⁷ As one does not have a clear view of what the truth-conditions of a ‘knowledge’ attribution are (i.e. being in a good enough epistemic position to eliminate the knowledge-destroying alternatives), he says, it is plausible that one bases the inference upon a part of these truth-conditions (i.e. being in a good enough epistemic position with respect to p).

Thus, he writes:

All you need is to see that if I didn’t think that S did so measure up, it would be odd, indeed uncooperative, of me to say, “S knows that p”; for whatever exactly knowledge is, we know that S’s knowing that p entails that S is in a good epistemic position with

¹⁵See also Rysiew (2001, 491-492)

¹⁶See DeRose (2009, 118-120) for a criticism of Rysiew’s account thus understood.

¹⁷“The leading idea of that account, rather, is simply that we need to distinguish between the content of a given sentence and what one might convey in uttering it on a given occasion; and that, with some help from general conversational rules (as well as parties’ knowledge of the context, their general background knowledge, and so forth), uttering a given sentence (in context) can, and often does, serve to communicate information over and above, or even quite other than, what the sentence itself expresses” (Rysiew, 2005, 65). See also DeRose (2009, 121). DeRose does not challenge this assumption, and hence, he proposes the first objection above against Rysiew’s account

respect to p. “And why”, the hearer may well ask, “would he say something that means [semantically implies] that if he didn’t in fact think that S measured up to the epistemic standards that are in play, and so was able to put to rest any doubts or rule out any not-p possibilities that had just been raised?” (2005, 49).

However, it is still unclear how this works. On the one hand, if “S knows that p” semantically implies that S’s epistemic position is good enough with respect to p in an absolute sense, then it is true that if S knows that p, then S’s epistemic position is good enough no matter the context. Under this interpretation, uttering ‘S knows that p’ will semantically express that S’s epistemic position is good enough for any X, where X is determined by the context. However, this cannot be what Rysiew means, because then Keith would be in a position to eliminate his wife’s salient alternatives in case B.

On the other hand, if Rysiew means that “S knows that p” semantically implies that S’s epistemic position with respect to p is good enough *to a certain extent*, then it is still unclear in virtue of what should the hearer be led to the pragmatic meaning that S’s epistemic position is good enough for the epistemic standards of the context. Indeed, “being good enough” is not a semantic part of “being good enough to a certain extent”. It does not follow from “X is good enough to a certain extent” that “X is good enough”. Therefore, it is still unclear why we should assume that since “S knows that p” semantically imparts that “S is in a good enough epistemic position with respect to p (to know p)” then “S knows that p” naturally imparts, pragmatically, that “S is in a good enough epistemic position with respect to p”.

To see how problematic is this assumption, note that, contrary to what Rysiew claims, this pragmatic inference from “S knows that p” to “S is in a good enough epistemic position with respect to X” is quite odd given the general conversational rules. In particular, given the conversational rule of quantity, requiring to be maximally informative, this is not what should be expected. (The rule of quantity is the rule explaining, for example, why it is misleading to say “John has two children” if he has three children).

Suppose, for instance, that Keith says “I know”, and by this, literally means that his epistemic position is good enough to eliminate the knowledge-destroying alternatives. Suppose that he need to be in a stronger epistemic position to eliminate also the further not knowledge-destroying alternative in question. Given this background assumption and the rule of quantity, in Keith’s context, the hearer should be led to the right piece of information, namely that Keith is not in a stronger epistemic position than the one required for knowledge. So, “I know that p” uttered by Keith should pragmatically impart that Keith is not in a position to eliminate his wife’s salient alternative. After all, he did not assert the stronger “I know for sure”, or “I am certain”.

Thus, one cannot just assume that it is natural to be led, from the semantic meaning of an utterance such as “S knows that p”, to the pragmatic meaning that S’s epistemic position with

respect to p is *good enough* for X . This is a mysterious fact that should be explained.¹⁸ Again, I do not deny that this might well be what is communicated. What I claim is merely that this is an unexplained fact in this pragmatic account. So, this pragmatic approach fails to give a fully satisfactory account in which the inference in question is not semantic.

In sum, the idea that a ‘knowledge’ ascription would lack relevance in case B is not sufficient to explain why the hearer is driven to the idea that the speaker is (or is not) in a *good enough* epistemic position with respect to the salient alternative. If the background account of the semantic of ‘knowledge’ is true, it remains to explain why the conversational rule of quantity does not play its role and why there would be a natural pragmatic inference from “S knows that p ” to “S’s epistemic position is *good enough* with respect to X ”.

4.6 Objection 3: Rejection of the epistemic closure on knowledge or more *ad hocness*

Further, this account is entirely based on the idea that Keith cannot know that the salient alternative is false, or at least, that he cannot rule it out – in the sense of discriminating between its truth and falsity.¹⁹

However, it may be problematic to suppose that Keith cannot know that this alternative does not obtain. Indeed, this would amount either to rejecting the principle of epistemic closure on knowledge, or to embracing the claim that he does not know that the bank is open on Saturday. According to the principle of epistemic closure on knowledge, if a subject knows that p , and knows that p entails q , then she can know that q on the basis of a competent deduction from p . Keith is supposed to know that the bank is open. Keith presumably knows that *the bank is open on Saturday* entails *the bank has not changed its hours (in a way such that the bank is closed on Saturday)*. Thus, if the principle of epistemic closure holds, Keith should be able to know, by performing a competent deduction, that the bank has not changed its hours.²⁰ As a result, if the account is based on the idea

¹⁸It can also be argued in a similar way that (3) may be false. By uttering “I do not know that p ” in a context with epistemic standards higher than the ones required for knowledge, it should be possible to convey the idea that I do not *merely* know that p . This case would be similar to the one where one says that one does not believe p in order to convey that one knows p . The fact that we do not use “I do not know that p ” in this way is a fact in need of explanation.

¹⁹Thus, Rysiew (2001, 492-493) writes: “[W]ere the speaker in Case B to respond to his wife’s bringing up the possibility that the bank has recently changed its hours with ‘No: I know that the bank’ll be open tomorrow’, she might be expected – by us, as well as by the speaker himself – to reply in turn, ‘You mean, you know the bank hasn’t just recently changed its hours?’ But the speaker doesn’t know this, or doesn’t know that he knows it (given what we’re told), and he knows that that is what his insisting that he knows would communicate in Case B. And that, as the present account has it, is why he doesn’t say this”. See also Rysiew (2001, 493-494): “if there’s some salient not- p possibility which the speaker can’t rule out, we shouldn’t expect him to claim to know that p ; for in doing so, as he well knows, he’d be communicating that he can rule out the salient not- p possibilities.”

²⁰Likewise, if this principle is true, he may know that he knows that the bank has not changed its hours, if he knows that he knows that the bank is open on Saturday.

that Keith does not know that the bank has not changed its hours, it relies on the rejection of the principle of epistemic closure on knowledge. This weakens the account because many philosophers find some version or other of the epistemic closure principle plausible.

One might think that one way of preserving the principle of epistemic closure in this kind of account would be to say that the account only relies on the fact that Keith cannot rule out the salient alternative, but not in the sense of “rule out” that knowledge requires. To see how this could work, let’s distinguish two notions of discriminability. The first one is Dretske’s notion of discriminability.²¹ The second one is Rysiew’s notion of discriminability.

According to the first notion of discriminability, to say that S can discriminate between p and not-p means that if p were false S would not believe that p (on the same basis). According to this brand of relevant alternative view, if p is possible, there is always a not-p world that is relevant, namely, the nearest world in which not-p (not-p is always a relevant alternative to p).

Call this the “strong” notion of discrimination (s-discrimination), and the related notion of elimination based on it the “strong” notion of elimination (s-elimination). According to these notions, we may say that you cannot eliminate the possibility that you are a brain in a vat because you cannot discriminate between being a brain in a vat and not being a brain in a vat. However, you can eliminate the possibility that you do not have hands because you can discriminate between having hands and not having hands.

Now, suppose that knowledge requires s-discriminability. On this view, if Keith cannot s-eliminate the possibility of error according to which the bank has changed its hours, then, that means that in the nearest possible world where the bank has changed its hours, Keith still believes that it has not changed its hours. Thus, if knowledge requires s-discriminability, Keith does not know that the bank has not changed its hours. However, on this view, Keith may still know that the bank is open, if it is true that, were the bank closed, he would not believe it is open (if, for instance, in the nearest possible world in which the bank is closed on Saturday, it was not open the last Saturday when Keith came).

According to this interpretation of the bank cases, we should reject the closure of knowledge, and on the basis of this rejection, we may explain how the Gricean manoeuvre can work. Although Keith knows that the bank is open on Saturday, he does not know that his wife’s alternative does not obtain, hence, he should not communicate that he is in a position to know that it does not obtain.

However, if one thinks that the principle of epistemic closure of knowledge is a plausible principle, one should not hold that knowledge requires s-discriminability. Rysiew’s own account of knowledge does not rely on this. The notion of discriminability that Rysiew has in mind is weaker (call it “w-discriminability” and the related notion of elimination “w-elimination”). It only requires that in the relevant worlds, where these worlds are relevant if and only if they are probable enough, if

²¹See Dretske (1970).

not- p is the case, then the subject does not believe that p (on the same basis). On this view, note that the nearest not- p world may be an irrelevant world. If so, it may be trivially true that, given all and only the relevant worlds, if not- p were the case, the subject would not believe that p . On this view, indeed, there may be no relevant world in which not- p (not- p is not always a relevant alternative to p).

The principle of epistemic closure of knowledge seems to be compatible with the idea that knowledge entails w -discriminability. For example, if there is no relevant world in which one is not a brain in a vat (because this world is too unlikely), then S can w -discriminate between being a brain in a vat and not being a brain in a vat, and hence, he can know that he is not a brain in a vat.

Now, how should we understand Rysiew's claim according to which Keith cannot eliminate his wife's alternative? Is it supposed to mean that Keith cannot s -eliminate, or w -eliminate, his wife's alternative?

If it is supposed to refer to w -elimination, and if knowledge requires w -elimination, then the account requires to reject epistemic closure. Indeed, if Keith cannot w -discriminate between the falsity and truth of his wife salient alternative, that means, firstly, that the probability that the bank has changed its hours is somewhat too high, and hence, is relevant. Secondly, that means that it is false that if the bank had changed its hours then Keith would not believe that it has not (on the same basis). But, at the same time, if we want to claim that Keith knows that the bank is open, then we must deny that the possibility that it has changed its hours is relevant (likely enough) for Keith's knowledge that the bank is open. Indeed, the second claim, that if the bank had changed its hours, then Keith would not believe that it has not (on the same basis) is still true. However, this amounts to saying that this alternative is relevant to knowing that the bank has not changed its hours, but not relevant to knowing that it is open on Saturday, while the latter proposition obviously entails the former. This amounts to denying epistemic closure.

On the contrary, suppose that Rysiew's approach appeals to the distinction between the two notions of discriminability. One may say that, on the one hand, as closure holds and Keith knows that his wife's alternative does not obtain, he can w -eliminate it. On the other hand, one may argue, although Keith knows that his wife's alternative does not obtain, he cannot s -eliminate it. In Rysiew's framework, this would mean that, in the target situation, the probability that the bank has changed its hours is too low (according to standards of knowledge set independently of the practical situation) to be knowledge-destroying. However, it may still be true that in the nearest possible worlds in which the bank has changed its hours, Keith believes that the bank has not changed its hours.

Under this interpretation, the only reason why Keith's epistemic position with respect to the salient alternative is not good enough is that he cannot s -discriminate between its truth and falsity. Hence, on this view, what a 'knowledge' attribution is supposed to impart pragmatically is not

merely that the subject is in a good epistemic position to know that the salient alternative is false (or to w-discriminate it). It also supposed to impart that she can strongly rule it out, that is, that she is in such a strong epistemic position with respect to the falsity of this salient alternative that even knowledge does not require to be in such a strong epistemic position.

As a result, not only does the pragmatic meaning of “I know that p” would pragmatically impart that the subject is in a *good enough* epistemic position to know that the salient (not-knowledge destroying) alternative is false (which is a first unexplained fact). In addition, that would communicate that the subject is in a good enough epistemic position to *strongly discriminate* between this alternative and its falsity, which is much more demanding than merely knowing that it is false. The problem is that this further pragmatic meaning is also left without an explanation. Assuming without explanation that the hearer is led to this pragmatic meaning is, of course, *ad hoc*.

In sum, either the account must reject the principle of epistemic closure, or there is a further unexplained fact – namely, that the ‘knowledge’ attribution in bank case B pragmatically conveys powers of strong discriminability, something that knowledge does not semantically implies – and hence, there is more *ad hocness* than initially thought.

4.7 Objection 4: The conversational context is underdescribed

A further objection is that Rysiew underdescribes the conversational context. If we fully describe this conversational context, it becomes unclear why the pragmatic meaning conveyed by a ‘knowledge’ attribution should be that Keith is in a good epistemic position to *eliminate his wife’s alternative*, and why the pragmatic meaning conveyed by a ‘knowledge’ denial should be that Keith is not in a good epistemic position to *eliminate his wife’s alternative*.

Let’s consider claim (4) above. A possibility of error has been raised in the conversation. Presumably, this partially changes the conversational score. As Grice writes (quoted by Rysiew (2005, 50)): “I expect a partner’s contribution to be appropriate to the immediate needs at each stage of the transaction” (Grice, 1989, 28). However, in the high stakes bank case, this change should not be thought as so important as to change the general goal of the conversation. The aim of the conversation is still to decide what to do. Rysiew’s approach does not essentially appeal to the practical situation.²² This is problematic, because we have to explain why the possibility raised by Keith’s wife is conversationally relevant in the first place. If it is not, indeed, Keith’s utterance

²²According to Rysiew, the possibility of error is conversationally relevant because it has been raised in the conversation. One can explain why it has been raised by the fact that the stakes are high (when the stakes are high, it is natural to think about possibilities of error.) Nevertheless, there may be cases where the same manoeuvre is available, even though the possibility of error is raised in a different way. For instance, Hazlett (2009) uses as similar manoeuvre to account for knowledge denials in sceptical contexts where there are no stakes. Brown (2006) notes that the presence of high stakes in case B makes a difference as to the extent to which the salient possibility of error is conversationally relevant. I shall come back to this claim below.

“I do not know” cannot convey, due to the rule of relevance, that he is not in a position to eliminate it. Indeed, this would amount to imparting something irrelevant to the conversation, namely, that he is in a position to eliminate a conversationally irrelevant possibility of error that has been raised in the conversation.

To see the point, note that we may all agree that when a not knowledge-destroying possibility of error is raised in a conversation, it should be taken into account by the subsequent utterances. But there are many ways to do so. One is to take it seriously; another one is to reject it as irrelevant; still another one might be to claim that one knows the target proposition although one cannot eliminate this alternative.

Suppose one thinks that Keith’s wife alternative is somewhat irrelevant to the aim of the conversation. If so, one may also think that the question whether Keith is in a position to eliminate it is somewhat irrelevant. Hence, one won’t think that Keith saying “I know that p” would communicate, in virtue of the rule of relation, that he is in a position to eliminate it (and that Keith saying “I do not know that p” communicates that he is not in a position to eliminate it). A proposition that is supposed to be pragmatically communicated by applying the rule of relation cannot be somewhat irrelevant to the aim of the conversation. Instead, one would rather think, on this view, that the utterance “I know that p” would communicate that the salient alternative is irrelevant to the aim of the conversation (and the utterance “I do not know that p” communicates the (false) proposition that the salient alternative is relevant to the aim of the conversation).

To illustrate, suppose one thinks that knowledge is the norm of action, and that the aim of the conversation is to make a practical decision. On this view, one won’t agree that Keith’s wife alternative is relevant to the aim of the conversation. Hence, one won’t think that Keith saying “I do not know that p” can be appropriate, although Keith knows that p, due to the fact that this communicates that he is not in a position to eliminate his wife’s alternative. Rather, given the aim of the conversation, and the knowledge norm of action, such an utterance would be likely to communicate that Keith’s wife alternative is knowledge-destroying. Then, under these assumptions, the rule of relevance cannot be used to explain why it seems inappropriate for Keith to say “I do not know” in the high stakes case.

Thus, we can see that a manoeuvre such as Rysiew’s must assume that the salient alternative in question is conversationally relevant because it is relevant to make a practical decision, and not relevant merely to the extent that it has been raised. Further, and consequently, it must assume that knowledge is not the norm of action.

Brown (2006), in particular, has insisted upon the idea that Keith’s wife alternative is conversationally relevant due to the practical situation.²³ As she says, we must explain why and how the fact the Keith’s wife has mentioned a possibility of error can change the goal or the direction of the

²³Rysiew (2007, 639) grants this point: “a ‘knowledge’ ascription conveys either that the subject can rule the salient alternatives out (his/her evidence eliminates them), or that these alternatives aren’t worth taking seriously”.

conversation, so that “know” means “can eliminate the salient alternatives” and “not know” means “cannot eliminate the salient alternatives”. According to her, the practical importance can explain this. One can more easily resist to possibilities of error raised in contexts where it is not practically important not to be wrong, than to possibilities of error raised in contexts where it is practically important. And, as she notes, the fact that a large cheque has been written “makes it clear that what’s relevant to the conversation is a very strong epistemic position” (Brown, 2006). Therefore, clearly, an epistemic position stronger than knowledge.

According to this line of thought, the question is whether Keith is in a strong enough epistemic position to eliminate his wife’s alternative before acting on the target proposition. This is a natural way of understanding bank case B. After all, in the scenario, Keith acknowledges that he should check whether the bank has not changed its hours. This suggests that, after doing that, he will be in a good enough epistemic position to act on the target proposition. If not, indeed, there would be no point in checking. Therefore, it seems that the salient possibility of error is practically relevant in that it is necessary and sufficient to eliminate it to be in a good enough epistemic position to act on the target proposition.

However, even if one understand in this way the practical relevance of this possibility of error, the question whether Keith can eliminate it, although important, is still *less* important than whether he knows that the bank is open on Saturday. Indeed, being in a position to eliminate this possibility of error can be sufficient to act on the target proposition only if it is already clear that Keith knows that the bank is open. But recall that it is assumed that it is not common knowledge that Keith knows. Thus, it seems, the more important question is still whether Keith knows. And if this question is more important, it is difficult to see why a ‘knowledge’ ascription (or denial) would, in virtue of the rule of relevance, communicate something with respect to the salient alternative, which is less important.

Why is the question whether Keith knows that the bank is open more relevant than whether he is in a position to eliminate his wife’s alternative (granting that it is not common knowledge that Keith knows)? To see this, let me give an example. Suppose that p is relevant to your action, and you have a choice between two options: know whether p or eliminate a not knowledge-destroying possibility of error in which not- p . It seems clear that you will choose the first option. For instance, consider the following case:

POLICEMAN. You are a policeman in a country where there is only one gun, and you are about to enter into a house to arrest some dangerous criminal. It is very important that the criminal is not armed. Before entering, you just have enough time to read one of the two following messages on your mobile: (a) Is the criminal armed? or (b) Does the criminal have a gun?

It seems clear that it is rational to choose (a) if you have to decide what to do. Indeed, suppose

that you choose (b) and the answer is positive. Then it is not rational for you to go into the house. Suppose that you choose (b) and the answer is negative. Still, it is not rational for you to go into the house because the criminal might be armed in a different way. An answer to (b) does not make any difference to what it is rational for you to do. On the contrary, suppose that you choose (a). If the answer is positive, then it is not rational for you to go into the house. But if the answer is negative, it is more rational for you to enter the house. Even if one still doubts whether it is fully rational to go into the house given that there is no answer to (b), it is clear that (a) is more relevant than (b) to the extent that it can make a difference to the rationality of your action.

Similarly in bank case B. Suppose that Keith's wife does not know whether Keith knows. Does she raise this alternative to know whether Keith knows, or merely to know whether he can eliminate the alternative? It is implausible to think that she raises this alternative merely to know whether Keith can eliminate it. Indeed, whether or not the bank has changed its hours is not fully relevant. Maybe the bank has not changed its hours, but actually it is closed on Saturdays except when Keith came the last time. Maybe it has changed its hours but it is still open tomorrow. On the contrary, if she raises this possibility to challenge Keith's knowledge, that seems relevant insofar as Keith's knowledge that the bank is open might make a difference to whether it is rational to come back.

This is not a decisive objection against the kind of account favoured by Rysiew, although it shows that some refinement is necessary. First of all, one may grant that in case B, at some point or other, it is common knowledge that Keith knows. Indeed, at some point or other, it is common knowledge that, when Keith came to the bank last Saturday, it was open. If so, it should be clear to everyone that Keith knows that the bank is open. Then, the question is why his wife asks whether he knows. Presumably, this is at this moment that the rule of relevance should play its role. Given that she knows that Keith knows, there would be no point in her asking whether he knows unless, by this question, she meant something different than what is literally expressed. It is natural to think that the question must concern whether Keith can eliminate the alternative in question. If eliminating this alternative is necessary, given the practical situation, Keith's wife question is relevant. Subsequently, Keith's obviously false answer pragmatically means – given the pragmatic meaning that 'know' now has in the conversation – that he is not in a position to eliminate this alternative.

Let me recap. First, this warranted assertability manoeuvre should assume that the alternative is relevant to the conversational context due to the practical situation (and not due to mere salience). This implies, given that this alternative is not knowledge-destroying, that this manoeuvre is incompatible with the idea that knowledge is the norm of action. However, we should not grant this claim without discussion. Second, a Gricean manoeuvre should rely on the fact that it is common knowledge that Keith knows. As made clear, this latter difficulty only constitutes a problem for Rysiew's particular account. A different account in the same spirit need not be subject to this particular criticism.

4.8 Objection 5: Two unexplained facts

The last difficulty is that this account is at odd with two general linguistic or conversational facts. First, it seems that, when it comes to make a practical decision, ‘knowledge’ ascriptions are never conversationally irrelevant. Second, the pragmatic implicature of a ‘knowledge’ ascription that the account posits does not seem cancellable.

First, it is clear that, when deliberating on how to act, ‘knowledge’ ascriptions (or denials) are always relevant. For example, suppose that it is right to come back tomorrow only if the bank is open tomorrow. Suppose that I say “We should come back tomorrow. I know that the bank is open”. One may question whether I really know, but it is clearly inappropriate to question the relevance of my contribution to the discussion. The following reply, in particular, would be quite surprising: “Well, the question is not whether you know that the bank is open”.

This is not what we should expect if the conversational relevance of ‘knowledge’ ascriptions were sometimes merely due to a pragmatic implication. If so, indeed, there would also be cases where the ‘knowledge’ ascription is conversationally irrelevant and the pragmatic implication not triggered. In such cases, the irrelevance of the ‘knowledge’ ascription should appear. To my knowledge, there are no such cases. Even in high stakes situations, saying that one knows the target proposition is considered as sufficient to settle the issue, even when no alternative has been raised to salience. In high stakes cases, one might be less willing to cite knowledge, but that does not show that citing knowledge does not settle the question.²⁴

This feature of ‘knowledge’ ascriptions, namely, that they are always conversationally appropriate in this kind of contexts, also suggests that, in such contexts, whether or not a subject knows the target proposition always has a particular significance.²⁵ This is also at odd with the claim that there is no normative connection between knowledge and action.

Second, suppose that in some cases the subject knows that *p* but, although the choice is *p*-dependent, his knowledge is not (fully) relevant. Suppose that in such cases, the inappropriateness of a ‘knowledge’ ascription must be explained by a pragmatic implication having to do with the rule of relation. If so, this pragmatic implication should be cancellable, thereby making the ‘knowledge’ ascription in question appropriate. Indeed, the claim would then be true and without problematic pragmatic meaning. However, suppose that, in case B, Keith says “Well, I know that the bank is open, but you are right. It is not rational to come back without checking whether the bank has changed its hours. The stakes are high”. This seems odd. Still, it is true that Keith knows, and the

²⁴See Fantl and McGrath (2007, 562).

²⁵For instance, discussing a case similar to the bank cases, Reed (2010, 232) writes:

[Y]ou are trying to decide whether to check if the train stops in Foxboro because it is extremely important that you get there as quickly as possible. You have not yet decided whether it is rational to check if the train makes that stop. One of the relevant factors in your decision would presumably be an answer to the question, do you know the train will stop there?

supposedly problematic pragmatic implication – that Keith epistemic position is good enough for the epistemic standards of the context – has been cancelled. This Gricean manoeuvre, based on the rule of relation, provides no explanation why the cancellation seems odd. In contrast, this oddity may be easily explained semantically if the subject knows and knowledge is sufficient for action.²⁶

In summary, two conversational or linguistic facts are at odd with the approach required by this Gricean manoeuvre. First, in contexts where the goal of the conversation is to make a practical decision, ‘knowledge’ ascriptions are always conversationally relevant. Second, the supposed pragmatic implication does not seem cancellable.

4.9 Conclusion

The orthodox account of knowledge is challenged by pairs of cases with asymmetrical stakes and different salient alternatives. These cases show that, intuitively, the appropriateness of ‘knowledge’ ascriptions varies with the mere variation of these practical factors.

I have considered a popular explanation of this variation that is compatible with the orthodox framework. According to this approach, what varies with practical factors is merely the warranted assertability conditions of our ‘knowledge’ ascriptions, not their truth-value. On this view, the variation is due to a variation in the truth-value of propositions that would be pragmatically conveyed by such ‘knowledge’ ascriptions (or denials). The rule of relation plays a crucial role in explaining why these implicatures are generated.

I have contemplated five objections. While some are not decisive, and others merely speak against a particular way of fleshing out the general approach, yet others undermine the general strategy based on the rule of relevance. In particular, I have argued that the rule of relevance is not sufficient to avoid the *ad hocness* of the account, and I have shown that this approach is at odd with two general conversational and linguistic facts left unexplained. In chapter 6, I shall offer a different pragmatic explanation that does not rely on Gricean implicatures, but on the norm of assertion. I shall argue that this new approach avoids the *ad hocness* objection and can deal with these further unexplained facts. Before coming to this, however, it is useful to consider a different approach that has been proposed in order to explain our shifty intuitions in an orthodox framework.

²⁶See Fantl and McGrath (2007).

Chapter 5

Linguistic intuitions and belief

5.1 Introduction

In the previous chapter, I have shown that the Gricean approach based on the rule of relevance is not fully satisfying. In this chapter, I consider another way in which the orthodox philosopher might want to explain our intuitions regarding the appropriateness of ‘knowledge’ ascriptions. It appeals to the relation between practical factors and beliefs. Whereas the Gricean approach aims at explaining our intuitions concerning the appropriateness of ‘knowledge’ ascriptions (and denials) by referring to intuitions concerning the pragmatic effects of making these ascriptions (and denials), this approach aims at explaining these intuitions by referring to intuitions concerning whether the subject or the ascriber (justifiably) believes the target proposition. For example, according to this line of thought, what Keith says in the high stakes bank case is appropriate because true. This can be explained in an orthodox framework by the claim that, as a result of perceiving high stakes, his confidence is shaken. Hence, he does not believe – in the sense required by knowledge – that the bank is open.

I consider two ways of fleshing out this approach. The first one has it that practical factors may raise our threshold for confidently believing because we are prone to confuse the epistemic standards that we must satisfy to be in a good enough epistemic position to act on p with the epistemic standards that we must satisfy to be in a good enough epistemic position to believe p . Thus, for example, we might think that Keith does not satisfy the epistemic standards with respect to the proposition that the bank is open to act on this proposition. If he makes this confusion, he is also inclined to think that he does not satisfy the epistemic standards to rationally believe that the bank is open, and hence, he does not believe that the bank is open. Thereby, he does not know that the bank is open.

Against this approach, I argue that it is implausible to say that we confuse the epistemic

standards of practical rationality with the epistemic standards of theoretical rationality in this way. First, either this does not explain a further case, or we should also grant that we sometimes confuse the practical situation of the subject under evaluation and our own practical situation, which is absurd. Second, the underlying confusion would have to be general, which is at odd with the fact that, in general, we do not confuse practical and epistemic reasons.

The second way of fleshing out this approach says that the perception of practical factors causally (and rationally) affects the belief-formation process, and hence, our willingness to believe a proposition. Given this well-established feature of our psychologies, we can explain that practical factors are relevant to the truth-value of ‘knowledge’ ascriptions since they are relevant to whether a subject can believe a proposition, as well as relevant to how much evidence a subject needs, as a matter of psychological fact, to settle her mind on the issue if she correctly proportions her (practically affected) belief to the evidence. Thus, for example, we might think that when the stakes are high, Keith needs more evidence to be able to believe that the bank is open, and hence, he does not believe that the bank is open. Thereby, he does not know that the bank is open.

I raise four problems for this proposal. First, I show that it cannot explain cases of rationally confident high stakes subjects. Second, this approach appeals to (hidden) compromising-accuracy factors to explain intuitions of theoretical irrationality, but I argue that this explanation is not satisfactory. Third, I argue that this approach faces a dilemma. Either it relies on the problematic assumption that we naturally favour assessments in terms of theoretical irrationality over assessments in terms of practical irrationality, or it does not, but, as a result, it cannot explain all the intuitions. Fourth, this proposal relies on the assumption that knowledge is incompatible with epistemic anxiety, and I argue that this is problematic.

5.2 The relevance of practical factors to our beliefs

It is a standard assumption in epistemology that knowledge requires, in some sense or other, a belief. As mental states, beliefs enter into various causal relationships. They can figure as causes in explanation of other mental states or actions. For example, if asked why John took his umbrella, we may say that the reason is that he believed it was raining and desired not to be wet. Beliefs can also figure as effects. Thus, we can causally explain the fact that Smith believes in God by the fact that he was influenced by an hypnotist. Relationships involving mental attitudes (and actions) are often not merely causal. Mental attitudes (and actions) are reason-sensitive.¹ Often, the explanation of an attitude (or action) also rationalises the attitude (or action) in question. Thus, the fact that John has a belief that it is raining (and a desire not to be wet) rationalises his intention of taking an umbrella. His belief (presumably along with his desire) provides him with a reason (or, perhaps, merely an apparent reason) to intend to take an umbrella. In contrast, suppose

¹See Scanlon (1998).

that John has taken a pill the effect of which is to prompt an intention of taking an umbrella when he believes that he is tired. If John believes that he is tired, and as a result acquires the intuition of taking an umbrella, his intention is not rationalised by his belief (or the content of his belief). Normally, that one is tired is not a reason for one to take an umbrella, and, any way, John did not acquire this intention as a result of being motivated by the reason *I am tired*. He did not reason from the content that he is tired to the intention of taking an umbrella.

When the explanation of an attitude is rationalising, the attitude is justified by the considerations that produces it, or, at least, the attitude is excused if the consideration in question is merely an apparent reason, or something in light of which a rational being would have acquired the attitude in question even if it is wrong. If the attitude is not produced by a rationalising consideration, it can still be explained by this consideration, but not justified.²

According to the orthodox account of knowledge, the attitude that knowledge requires is never rationalised, or justified, by truth-irrelevant factors. However, it can sometimes be explained by truth-irrelevant (practical) factors. We might think that practical factors, such as stakes, can influence this attitude in at least two ways. The first one is when the subject makes a mistake, and confuse the question whether it is epistemically right to believe *p* and the question whether it is epistemically right to act on *p*. Being granted that practical factors can affect whether it is epistemically right to act on *p*, if the subject can make such a mistake, practical factors can affect her willingness to believe that *p*. The second one is when the belief-formation process is causally (and rationally) affected by such factors. It is a well-established fact, in particular, that when we perceive high stakes, we become anxious, and hence, that it is more difficult to form a belief.

Now, suppose that our beliefs can be affected by practical factors, in such a way that when stakes rise on whether *p*, we are less prone to believe that *p*. If so, when the stakes rise for *S*, we should be less prone to ascribe ‘belief’ to *S*. Given that knowing that *p* requires believing that *p*, in ascribing ‘knowledge’, we implicitly ascribe ‘belief’. If so, it is not surprising that when the stakes rise for *S*, we are less prone to ascribe ‘knowledge’ to *S*. Hence, a change in stakes for *S* can affect our intuitions about the appropriateness of ‘knowledge’ ascriptions to *S*.

It is important to note that, depending on the case, practical factors can affect the attributor’s belief, or the subject’s belief, or both. Thus, the explanation why it is (intuitively) inappropriate for an ascriber *A* to ascribe ‘knowledge’ to a subject *S* can take different forms. For example, if an ascriber does not believe that *p*, she will not be willing to ascribe ‘knowledge that *p*’ to others. If she ascribes ‘knowledge that *p*’, indeed, that means she thinks that believing *p* is justified, and hence, disbelieving *p* is unjustified. Hence, she should see herself as having a wrong attitude toward *p*. Thus, it may be inappropriate for *A* to ascribe ‘knowledge that *p*’ to *S* even if *S* knows that *p*.

In the remaining of this chapter, I shall examine these two ways in which one might think that practical factors affect the belief-formation process, and the ways in which we might think that they

²See chapter 7.

are thereby relevant to explain the (intuitive) shifty appropriateness of ‘knowledge’ ascriptions.

5.3 The confusion between the epistemic standards of action and those of belief

5.3.1 Bach’s account

Suppose that a subject knows that p only if she confidently believes that p . Next, suppose that when a subject perceives high stakes on being right about p , her threshold for confidently believing p is raised. Thus, following Bach (2005; 2008; 2010) an orthodox epistemologist might explain our intuitions in the bank cases as follows.

In the low and high stakes cases, Keith has the same evidence, and this evidence is sufficient to justify the belief that the bank is open (the belief that the bank is open is propositionally justified for Keith). Thus, it seems, in both cases, if Keith believes that the bank is open on the basis of the evidence, he knows that the bank is open. In the low stakes case, the threshold for confidently believing is low. Therefore, the evidence is sufficient to produce an outright belief. Keith knows that the bank is open. In the high stakes case, however, the threshold for confidently believing is higher. Although the evidence is sufficient to justify a belief, Keith does not believe that the bank is open. Therefore, he does not know that the bank is open. This obviously explains why, when he says that he does not know, he says something appropriate. Indeed, he says something true.

Against this approach, firstly, one might insist that it is stipulated in the case that Keith remains as confident as before, even when the stakes rise. Therefore, it is stipulated that the threshold for confidently believing is not raised, or that even if it is raised, Keith still believes that the bank is open on Saturdays. However, this stipulation should strike us as incoherent. As Bach (2005, 76) notes, when you take yourself as not knowing that p , in general, you are not as confident that p as when you take yourself as knowing that p . Thus, it seems, either Keith believes in the same way that the bank is open in both cases, and hence, it is unclear how we can have the intuition that he may appropriately say that he does not know in the high stakes case. Or his confidence is somewhat shaken when the stakes rise, and hence, we can see why it is appropriate for him to say that he does not know. Thus, the bank cases seem to rely on a problematic assumption. Once this is made clear, the problem disappears.³

However, secondly, this approach cannot explain third-person ‘knowledge’ denials as easily. Consider Cohen’s airport case:

AIRPORT. Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago.

³See also Nagel (2010, 421).

They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds “Yes I know it does stop in Chicago.” It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, “How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute.” Mary and John agree that Smith doesn’t really know that the plane will stop in Chicago. They decide to check with the airline agent. (Cohen, 1999, 58)

In this case, Mary and John seem to appropriately deny ‘knowledge’ to Smith. They are in a high stakes situation, and they are assessing a low stakes subject who clearly believes the proposition in question, on the basis of some evidence. They know that the subject (Smith) believes the target proposition on this basis. If, further, the evidence is sufficient to justify this belief, then the subject (Smith) knows. So, contrary to the high stakes bank case, the attributors’ intuitively appropriate judgement that the subject does not know cannot be explained on the basis of the fact (or on the basis of the attributors’ intuition) that the subject does not believe the target proposition. We must explain in a different way why Mary and John’s judgement seems appropriate.

We should note that the attributors are in a high stakes situation. If Bach’s approach is correct, we must grant that their threshold for believing is high, so that they do not believe the target proposition. But, of course, when we do not believe that p , we cannot coherently think that someone else, who possesses the same evidence for p , knows that p . Hence, it is incoherent to ascribe ‘knowledge that p ’ to this person. Indeed, this would amount to granting that a belief in p , given this evidence, is justified. But this is at odd with the fact that we do not believe that p . Therefore, it is clear that the intuitive appropriateness of Mary and John’s claim that Smith does not know cannot be explained by Smith’s lack of belief (or by Mary and John’s intuition that Smith does not believe the target proposition). But we can explain this by the fact that Mary and John do not believe the target proposition. As a result, Bach’s view has it that Mary and John are making a false judgement, which is nonetheless intuitively appropriate in light of their lack of belief.⁴ Their false judgement manifests what we should expect from them in term of coherence.

Now, we must explain why Mary and John do not believe the target proposition, if such a belief is in fact justified by the evidence that they have. More generally, why does the belief threshold rise when stakes rise? According to Bach:

One’s threshold for (confidently) believing a proposition is a matter of what one implicitly takes to be sufficient reason to believe it (...). Even if in fact one is in a position to know something, thinking one is not in a position to know it is enough to keep one from believing it (at least not without reservations) and to lead one, if it matters enough, to

⁴See, e.g., Bach (2005, 76-77).

look into it further (Bach, 2005, 77).

On this view, one's threshold for believing rises to the extent that one makes a normative mistake concerning one's reasons to believe, or as far as one displays "excessive epistemic caution" (Bach, 2010, 117).⁵ Why such a mistake?

In several places, Bach suggests that such a mistake comes from a fundamental confusion between the epistemic standards of practical rationality and the epistemic standards of theoretical rationality. In some cases, the practical situation requires to gather further evidence before acting. But this can lead us to think erroneously that more evidence is required to have a justified belief, and hence, to know. This explains why the threshold for believing is raised.⁶

Thus, this approach assumes that knowledge is not the norm of action. It assumes that a subject can know that *p* even if, for prudential reasons, she should check whether *p*. For example, if Keith really believes that the bank is open on Saturday (let's assume that his belief threshold is not raised), he may know that the bank is open. Still, he should check before acting.

A problem for this approach might come from cases where there can be no incoherence between a 'knowledge' ascription and the attributor's (and the subject's) lack of belief. Let us assume that the high stakes attributor believes (and knows) the target proposition because she has more evidence than the low stakes subject under evaluation. For instance, modify AIRPORT. Suppose now that Mary has more information than Smith (she has double-checked), and thereby self-ascribe 'knowledge' (call this case "GOOD EVIDENCE AIRPORT"). Mary may still deny, in a way that seems appropriate, that Smith knows. However, we cannot explain the intuitive appropriateness of this denial by Mary's lack of confidence.

There is a reply, though. The same supposed confusion can explain this 'knowledge' denial. Mary confuses the epistemic standards of practical rationality – the amount of evidence required to rationally act on the target proposition – with the epistemic standards of theoretical rationality – the amount of evidence required to rationally believe the target proposition. This explains why

⁵See also Bach (2005, 52, 77, 81).

⁶See Bach (2005, 30), (2010, 117-118). See also Bach (2008): "It is hard to argue with the claim that as the stakes go up, then, everything else equal (for example, there is plenty of time and the cost of checking further and guarding against counterpossibilities is small compared to the stakes), one should take greater care in the ways indicated. But is this a matter of epistemic or merely practical rationality? Being practically rational requires risk management, and guarding against error is part of this. Since the size of the stakes increases the cost of being wrong, it also increases the range of possibilities one should guard against. It requires not reflexively acting on assumptions one would normally make and act on without question, not if there is a chance they are false and acting on them if they are false would be costly. This much makes perfectly good sense. However, it does not follow that the standards of knowledge go up (or that the threshold of justification is higher or that the range of relevant alternatives to be ruled out is wider). Rather, the higher stakes raise the threshold of confident, doubt-free belief. One's practical reasons give one reason not to treat certain possibilities as closed, hence not to act as if they do not obtain. One's practical interest explains the rise in the threshold of confident, settled belief, and thoughts of counterpossibilities make it more difficult for this threshold to be crossed. Lack of doubt-free belief keeps one from meeting the doxastic condition on knowledge. Proponents of pragmatic encroachment assume that the justification condition on knowledge has gone up in cases of higher stakes. In my view, one doesn't know in these cases because the stakes have raised the threshold of confident, doubt-free belief."

her threshold for believing is raised. She then believes that Smith's evidence is not good enough for knowledge. And since Smith does not base his belief on the amount of evidence required for Mary's action, Mary mistakenly denies that Smith knows. In sum, beyond the attributor's belief, or lack of belief, what really matters is her threshold for confidently believing, and the fact that it can be raised by practical factors.

5.3.2 Two problems: low stakes assessors and the kind of underlying error

According to Bach, we confuse the epistemic standards of practical rationality and those of theoretical rationality. Therefore, practical factors affect our threshold for confidently believing. This explains why 'knowledge' ascriptions shift with practical factors. In some cases the attributor's belief threshold is excessively high, which is why it seems appropriate for her to deny 'knowledge' to subjects who actually have knowledge.

Now, although this approach can explain HIGH STAKES BANK CASE, AIRPORT, GOOD EVIDENCE AIRPORT, it fails to account for a different type of case. Suppose that an assessor believes the target proposition. She is in a low stakes situation. Her belief threshold is low, and she bases her belief on modest, but sufficient, evidence. Hence, she knows that *p*. Suppose she is assessing a high stakes subject possessing exactly the same evidence. This assessor may judge, in a way intuitively appropriate, that the high stakes subject should not say that she knows the target proposition. To illustrate, consider AIRPORT. Suppose that Smith considers Mary and John's situation. Intuitively, he should judge that they should not say that they know, no matter what they believe (call this case "SMITH AIRPORT"). To take another example, consider our own intuitive judgements concerning Keith in the high stakes case. Whatever Keith believes, one thing is certain: it would be inappropriate for him to say that he knows.⁷

How can Bach's account explain the intuition we may have, to the effect that a high stakes subject with modest (but knowledge-level) evidence should not say that she knows, being taken for granted that we ourselves are not in a high stakes situation? On Bach's view, in the high stakes case, if Keith believes that the bank is open on the basis of the evidence that he has, then he knows. If he says that he knows, this is true. So what kind of error can be involved in the intuitive assessment that it is inappropriate for Keith to say that he knows?

As we have seen, there are two options available in Bach's framework. First, we might invoke the

⁷According to Fantl and McGrath (2009, 44), if we see Keith as outright believing the target proposition (and hence, according to the view under consideration, as knowing the target proposition), we should see him as saying something like: "Aw come one, I know [the bank] will be open. We'll just come back tomorrow". They write: "Intuitively, Keith's knowledge-claim, 'I know' is false", and, of course, it is irrational to come back without checking. However, it is unclear why postulating that Keith outright believes (and knows) the target proposition would amount to see Keith as willing to say "I know" (even if this is true) and "We'll just come back tomorrow". This depends on how we understand the notion outright belief, what we take the norm of action to be, and whether we think that such a 'knowledge' ascription would be appropriate just in case it is true.

lack of belief of the assessor (e.g. Smith). If the assessor does not believe the target proposition, then it is clear why she can think that the subject (e.g. Mary) should not say that she knows. However, in the kind of cases under consideration, it is assumed that the assessor believes the target proposition. For example, in SMITH AIRPORT, it is assumed that Smith believes that the flight has a layover in Chicago.

Second we might invoke the fact that the assessor's (Smith's) threshold for believing is high, because she confuses the epistemic standards of practical rationality and the epistemic standards of epistemic rationality. However, in the case under consideration the assessor is in a low stakes situation. Therefore, even if she confuses the epistemic standards of her own practical situation with the epistemic standards of theoretical rationality, this should make no difference. We may assume that her threshold for believing is low. She should still think that the high stakes subject satisfies these epistemic standards. Hence, she should be willing to ascribe 'knowledge' to this subject. If the low stakes assessor still thinks that the high stakes subject should not say that she knows, that cannot be due to the fact that the assessor thinks that this high stakes subject does not satisfy the standards of knowledge. But then, Bach's approach offers no explanation here.

Perhaps we might think that, after all, the assessor's threshold for believing is high. But then, two problems arise. First, how can she still believe that p and self-ascribe 'knowledge', if she has no more evidence than the subject under evaluation? Second, according to this line of thought, if the assessor's threshold for believing is high, that must be because it is raised by a practical consideration. But only the practical situation of the subject under evaluation might explain that the assessor's threshold rises. This means that the assessor does not merely confuse the epistemic standards of her own practical situation with the epistemic standards of theoretical rationality. She also confuses the epistemic standards of her own practical situation with the epistemic standards of the subject's practical situation. This is implausible.

There is a general reason why Bach's approach cannot explain this kind of cases. In such cases, the subject and the assessor are seen as knowing that p , although the subject is seen as inappropriately saying 'I know that p '. As Bach's approach aims at explaining our intuitions regarding the appropriateness of 'knowledge' ascriptions in terms of the (supposed) lack of knowledge of the assessor, or the (supposed) lack of knowledge the subject, it cannot explain cases in which a subject is seen as inappropriately saying 'I know that p ', although it is clear to the assessor and the subject that they know that p .

In other words, the problem with Bach's general strategy here is that it aims at explaining the intuition that it is inappropriate for a high stakes subject to say that she knows by appealing to a (sometimes mistaken) judgement that she does not know. However, it is clear that, without incoherence, we may judge that a high stakes subject knows that p although she should not say that she knows. For example, it is not incoherent for Smith to judge that Mary knows the target proposition (if she believes it), while at the same time judging that it would be inappropriate for

Mary to say that she knows (or to say that Smith knows). As there is no (mistaken) judgement that Mary (or Smith) does not know, the intuitive judgement that it would be inappropriate for Mary to say ‘I know’ cannot be explained by invoking a lack of knowledge.

A second problem for Bach’s approach has to do with the kind of error which is postulated. According to Bach, the error theory upon which his account relies is not troubling – in contrast to the semantic error postulated by epistemic contextualism. Indeed, “it has people sometimes denying knowledge of people who have it, it does not have them generally confused about knowledge or “knowledge”” (Bach, 2005, 78) Of course, a theory relying on the idea that our intuitions are systematically mistaken is *prima facie* less plausible than a theory relying on the idea that they are *sometimes* mistaken. There is no denying that we sometimes make mistakes. In contrast, a theory assuming systematic mistakes run the risk of being *ad hoc*.

However, contrary to what Bach claims, there still is a systematic error involved in his theory. It concerns what we take to be the epistemic standards of knowledge.⁸ Indeed, this approach is not based on the idea that we *sometimes* confuse the epistemic standards of practical rationality with the epistemic standards of theoretical rationality, but that, *in general*, we are not sensitive to the difference between the two. This assumption is required to explain why, in general, the belief threshold is raised by high stakes. Therefore, there is something troubling in Bach’s error theory.

Further, this supposed error is all the more surprising that we do not make such errors when it comes to considering the difference between epistemic reasons and practical reasons. For example, very few people, if any, take the fact that someone desires to believe that it is raining as an epistemic reason for this person to believe that it is raining, nor reason from the belief that believing p would satisfy a desire that they have to the belief that p. Thus, it is quite surprising that we confuse the epistemic standards of practical rationality with the epistemic standards of theoretical rationality.

In sum, Bach’s account is not satisfactory because it does not do well in explaining our intuitions regarding the inappropriateness of claims to know from high stakes subjects that are additionally (properly) assessed as knowing. Further, it relies on the controversial thesis that we often confuse the epistemic standards of practical rationality with the epistemic standards of theoretical rationality.

5.4 High stakes and desire for truth

5.4.1 Need-for-closure and epistemic anxiety

In light of the previous considerations, we may want to abandon the claim that the belief threshold is determined by what the subject takes to be a sufficient reason to believe, while preserving the general idea that when the stakes are high, the threshold for confidently believing rises. Instead, we can say that, as a matter of psychological fact, perceiving high stakes provokes anxiety, and

⁸See also Fantl and McGrath (2009, 53).

hence, affects the belief-formation process.⁹ It is important to stress that, on this approach, the explanation why the belief threshold rises with the stakes has nothing to do with a putative error regarding the epistemic standards of theoretical rationality.

Following Nagel (2008), in particular, we should note that the belief-formation process is affected by motivational practical factors.¹⁰ We can say that these factors, integrated and weighed together, determine the need-for-closure of the subject. The concept of need-for-closure is borrowed from Kruglanski. The need-for-closure of a subject determines the point where “a belief crystallizes and turns from hesitant conjecture to a subjectively firm ‘fact’ ” (Kruglanski and Webster (1996, 266) quoted by Nagel (2008, 287)). If we achieve closure, we do not “continue juggling alternative possibilities” (Nagel, 2008, 287).¹¹ In other words, practical factors determine when the belief-formation process is liable to end in producing an outright belief.

For example, suppose there is some pressure of time and the cost of error is low. It is to be expected that the subject will easily form a belief. Given a strong motivation to have a belief on the issue and a weak motivation to have an accurate belief, the subject’s need-for-closure is high. In contrast, suppose the subject can take her time to settle the issue and the reward in accuracy is high. We must expect her not to form her belief as easily. Given a weak motivation to settle the issue and a strong motivation to have an accurate belief, the subject’s need-for-closure is low.

In general, the practical factors determining the need-for-closure of a certain subject are of two kinds. Some are source of theoretical irrationality, whereas others are source of theoretical rationality. The former have to do with the utility of having a belief, that is, the utility of having a belief even if it is an epistemically unjustified belief. The latter have to do with the utility of having an accurate belief, that is, the utility of avoiding a belief which is likely to be false, and hence, the utility of avoiding an epistemically unjustified belief.

For example, a desire to believe a particular proposition because some gain is to be expected is a force (that may be practically rational) liable to influence the belief-formation process, like in wishful thinking. However, the presence of such a factor affecting the belief-formation process also makes it likely that the resulting belief will not be epistemically rational. Given how the belief-formation process is affected (perhaps for a practical reason) it is likely that its result will not be an epistemically justified belief. Haste, distraction, etc., are other factors affecting the belief-formation process in ways that are not source of epistemic or theoretical rationality. The presence of all these factors can conflict with the motivation for having an accurate belief, and hence, with the production of an epistemically rational belief.

⁹See Nagel (2008, 2010).

¹⁰See also chapter 2 section 4 on the idea that there are practical reasons to form beliefs or to refrain from forming beliefs, but that this does not impugn intellectualism about justified beliefs.

¹¹We must distinguish specific and non-specific need-for-closures. The former has to do with a motivation for having a specific belief (believing p rather than not- p). This motivation, for example, biases us in wishful thinking. The latter concerns a motivation to have a settled belief (either p or not- p). The notion of closure relevant here is the non-specific one.

In contrast, other practical factors may affect the belief-formation process in a sense that is not epistemically or theoretically bad. Such factors are factors prompting a desire for more evidence, or for a more thorough process of assessment. Although these are practical factors affecting the belief-formation process that cannot justify a belief, they are source of epistemic rationality in that it is more likely that a belief-formation process thus affected will produce a justified belief.

In brief, whereas the practical factors that are source of epistemic irrationality are factors that affect the subject's belief-formation process in such a way that it becomes less sensitive to the belief's accuracy, and hence, less sensitive to epistemic reasons to believe, the practical factors that are source of epistemic rationality are factors that affect the subject's belief-formation process in such a way that it becomes over sensitive to the belief's accuracy.

While these two kinds of practical factors can explain why a subject has (or lack) a certain belief, neither of them can rationalise, or justify, the belief in question.¹² They are not (apparent) reasons for this belief. The fact that you will be rewarded if you believe *p* may be a reason to affect your belief-formation process so as to cause in yourself the belief that *p*. But this fact is not a reason to believe that *p* that can justify this belief. The fact that you will be punished if you believe *p* when *p* is false may be a reason to affect your belief-formation process so as to prevent in yourself the belief that *p* unless *p* is absolutely certain. But that is not a reason that can justify the suspension of your judgement. These considerations do not affect the (epistemic) rationality of your beliefs, but they affect your belief-forming process and your willingness to believe.

The set of practical factors that are source of epistemic rationality, that is, factors which prompt a desire for accuracy, determines what Nagel calls the "level of epistemic anxiety", where epistemic anxiety is defined as the force or desire for increased cognitive activity (either in terms of greater amount of evidence or in terms of more thorough assessment of the evidence).¹³ The level of epistemic anxiety is in particular a function of the perceived stakes.¹⁴ The higher the stakes are, the more motivated the subject will be to have an accurate belief. Hence, the stronger the epistemic anxiety she will experience.

In short, we adapt our cognitive effort to the practical situation.¹⁵ As epistemic anxiety is led by a desire for accuracy, responding to epistemic anxiety by satisfying this desire is source of epistemic or theoretical rationality. On the other hand, factors that conflict with epistemic anxiety are the kind of factors that compromise accuracy. As a result, there are two ways of forming a belief given some level of epistemic anxiety. Either one satisfies this desire for accuracy and cognition (by gathering more evidence or by assessing more thoroughly the evidence), or this desire is overshadowed by another practical factor that is not source of epistemic rationality.¹⁶ In the

¹²See chapter 2 section 4.

¹³See Nagel (2010, 214).

¹⁴The level of epistemic anxiety can also be influenced by other factors such as whether the question is posed in a familiar way (Nagel, 2010, 214).

¹⁵See Nagel (2010, 409).

¹⁶See Nagel (2010, 414). It is worth stressing that the way in which we respond to epistemic anxiety is not supposed

second case, the resulting belief is supposed to be epistemically unjustified since (partially) formed on the basis of truth-irrelevant factors affecting the belief-formation process.

5.4.2 Epistemic anxiety and ‘knowledge’ ascriptions

How does this approach explain our intuitions about ‘knowledge’ ascriptions and denials? DeRose’s high stakes bank case is treated as in Bach’s proposal. In the high stakes case, Keith does not outright believe that the bank is open because he has a low need-for-closure.¹⁷ This low need-for-closure is explained by his epistemic anxiety (a high desire for accuracy) resulting from the fact that he perceives that the cost of error is high. As knowledge requires an outright belief (i.e., to have closed the issue), he does not know that *p*. Therefore, it is appropriate for him to say that he does not know that the bank is open because this is true.

We might think that, similarly, Cohen’s airport case can be roughly treated as in Bach’s account. It may be granted that Mary says that Smith does not know because she does not outright believe the target proposition (she feels high epistemic anxiety), and that, as a result, she is not in a position to ascribe ‘knowledge’.

However, important differences between Bach’s account and Nagel’s account are worth noting. Bach’s explanation relies on the idea that we confuse the epistemic standards of practical rationality with those of theoretical rationality. According to Bach, in AIRPORT, Mary has a high belief threshold because she confuses these epistemic standards. This confusion does not disappear when she considers the beliefs of other people. Given her practical situation, Mary takes as epistemic standards of theoretical rationality excessively high epistemic standards. But Smith does not satisfy these excessively high standards. Therefore, Mary denies ‘knowledge’ to Smith. This also explains why she denies ‘knowledge’ to Smith even after she acquires sufficient evidence to believe the target proposition (see GOOD EVIDENCE AIRPORT).

Nagel’s account of AIRPORT is very different in this respect.¹⁸ It appeals to the so-called “hind-sight bias”. This bias affects ascribers of propositional attitudes, so that they tend to mistakenly project their own doxastic situation onto the subjects of these ascriptions.¹⁹ Given this bias, Mary

to be voluntary. Epistemic anxiety is like an epistemic feeling to which our cognitive effort automatically respond in a subconscious way (by automatically selecting more accurate but more costly strategies (this is the “toolbox model”) or by setting the evidence threshold (this is the “evidence accrual model”). Thus, the cognitive response to epistemic anxiety is not achieved at the level of personal conscious reflection.

¹⁷Nagel (2008, 289).

¹⁸Also, Nagel considers a slightly different case, from DeRose (2009). The case is the following. Lena, Thelma and Louise share the same evidence that John was at the office today (and they know that they share the same evidence). They saw his hat and they heard someone who called him. Now, Thelma at the tavern appropriately says that Lena knows that John was at the office. However, before the police asking for a testimony that John was at the office, Louise appropriately says that she and Lena do not know that John was at the office.

¹⁹See Nagel (2010, 425): “We have a well-documented tendency to misread the mental states of those who are more naive than we are, to evaluate them as though they were privy to our concerns, without being aware that we are doing so”. Nagel (2010) also proposes an alternative explanation of Thelma, Lena and Louise Case. When the police is asking Louise whether Lena knows, presumably, they are asking about Lena’s capacity to testify. However,

will erroneously ascribe the anxiety she feels to Smith. Now, given Smith supposed anxiety, Nagel argues, Mary cannot see him as having a justified belief.

This last claim need clarification. After all, it is clear to Mary that Smith believes the proposition on the basis of modest, but normally sufficient, evidence. So why does Mary think that Smith's belief is unjustified? The reason why Mary takes Smith's belief to be unjustified, given that she erroneously ascribes to him high epistemic anxiety, is not due to a confusion between the epistemic standards of practical rationality and those of theoretical rationality. Rather, Nagel suggests that this is due to the fact that, by ascribing a high epistemic anxiety to Smith, Mary expects Smith either not to have formed a belief on such modest evidence, or, if he has formed a belief, to have formed it on the basis of epistemically bad considerations. If Smith has formed his belief on the basis of epistemically bad considerations, his belief is unjustified.

But why should Mary expect that, given his (supposed) epistemic anxiety, if Smith believes the target proposition, his belief is formed on the basis of epistemically bad considerations? Recall that either a desire for truth (epistemic anxiety) can be satisfied (calm down) by gathering more evidence (or by assessing the evidence more thoroughly), or else this desire can be overshadowed by accuracy-compromising factors (haste, distraction, ...). Now, Smith is seen (erroneously) by Mary as having a very strong desire for truth. Smith cannot calm down his supposed epistemic anxiety with modest evidence. So, if he outright believes that p – if his epistemic anxiety is calmed down – while he has made not further cognitive effort, nor acquired further evidence, the only explanation why he outright believe that p is that his desire for truth has been overshadowed by compromising-accuracy factors, such as haste, distraction, or whatever. Of course, a belief resulting from a process thus affected is not epistemically rational. Therefore Smith does not know. In sum, according to Nagel, the intuitive appropriateness of Mary's (mistaken) denial that Smith knows is based, first, on the (mistaken) projection of her own epistemic anxiety onto Smith, and, second, on the psychological fact that a subject with such an epistemic anxiety and modest evidence can outright believe the target proposition only if her epistemic anxiety (her desire for truth) is overshadowed by the presence of a compromising-accuracy factor.

The two other cases to consider are those that involve no lack of confident belief from the attributor and from the subject. They are also treated differently in this account. Consider first GOOD EVIDENCE AIRPORT. This case involves a high stakes attributor knowing the target proposition (Mary, after she acquires more evidence) and assessing a confident low stakes subject who has less evidence (Smith). The attributor cannot project her lack of settled belief. Still, she will say – and it may seem, appropriately – that the subject does not know.

Here also, we can appeal to the hindsight bias, though. If the attributor projects her own doxastic situation, presumably, she ascribes to the low stake subject the same need-for-closure, and hence,

if asked, given the epistemic anxiety she will feel, Lena will probably say that she does not know. That explains why it is appropriate for Louise to say that Lena does not know, thereby imparting that Lena would not testify that John was at the office.

she takes the subject to need the same amount of evidence to satisfy it. In the attributor's point of view, the subject has a low need-for-closure, and hence, she believes on the basis of (psychologically) insufficient evidence. As in AIRPORT, if the subject has an outright belief, the attributor thinks, this belief must be based on non-evidential or epistemically bad considerations. It is unjustified.²⁰

Second consider SMITH AIRPORT. This case involves a confident low stakes subject (Smith) assessing a high stakes confident subject (Mary, confidently believing the target proposition on the basis of modest evidence). It seems that this confident high stakes subject should be taken by the assessor as knowing the target proposition, since this confident high stakes subject believes the proposition on the basis of sufficient, although modest, evidence to justify the proposition. Intuitively, though, the low stakes assessor will appropriately judge that it is inappropriate for the high stakes subject to self-ascribe 'knowledge'. This was the problematic case for Bach's account.

The same kind of problem seems to arise here. If the assessor has a normal epistemic anxiety, we cannot explain, on the basis of a projection of her own doxastic situation onto the subject, why the assessor can erroneously (although appropriately from her point of view) judge that the confident high stakes subject (who believes on the basis of the same evidence) does not know. Hence, we cannot explain in this way why the assessor can (appropriately) judge that the subject should not self-ascribe 'knowledge'. The hindsight bias is useless here.

Nagel's reply could be that, actually, the confident low stakes assessor (Smith) does not erroneously judge that the confident high stakes subject (confident Mary) does not know. On the contrary, the confident low stakes assessor rightly judge that the confident high stakes subject does not know. And if this is correct, we can explain the appropriateness of the assessor's judgement that the high stakes subject should not self-ascribe 'knowledge' by the fact that this 'knowledge' self-ascription is false.

Why think that the high stakes confident subject does not know? Again, the assessor see that the confident high stakes subject perceives high stakes. She then correctly ascribes to her a high level of epistemic anxiety. Given this high level of epistemic anxiety, this subject should have a low need-for-closure, unless compromising-accuracy factors (e.g., haste, distraction) overshadow this epistemic anxiety. The evidence at the disposal of the confident high stakes subject (as well as her cognitive effort) is modest. Given that she has a high epistemic anxiety, if there are no compromising-accuracy factors at work, we should not expect her to be in a position to settle her mind on the issue with such modest evidence and cognitive effort. Thus, if the subject has an outright belief, that must be on the basis of non-evidential or epistemically bad factor. Therefore, the subject does not know.

Again, this explanation is based on the claim that we can expect of someone who has a high epistemic anxiety (a strong desire for truth) on whether p either not to have formed a belief on

²⁰Nagel (2010, 426) uses a similar strategy to explain our intuitions concerning 'knowledge' denials in cases in which the assessor is aware of the high stakes situation of a subject ignoring that the stakes are high for her (see Stanley (2005b) for this kind of case).

whether p if the evidence is modest, or else to have formed this belief on the basis of epistemically bad factors. In this sense, as Nagel notes, the representation of a subject's interests (along with the evidence at her disposal and the cognitive effort deployed) can help us to assess whether she knows. The subject's interests manifest the level of epistemic anxiety she (psychologically) should have. If she has an outright belief while the evidence is insufficient to satisfy this level of epistemic anxiety, then this subject is affected by a compromising-accuracy factor (such as haste or distraction). She does not know the target proposition.²¹

To sum up, Nagel's approach relies on the idea that perceiving high stakes raises epistemic anxiety. As in Bach's approach, this proposal aims at explaining our intuitions concerning the appropriateness of 'knowledge' ascriptions (and denials) by referring to intuitions concerning whether the subject knows. In this respect, it differs from Gricean approaches, which aim at explaining our intuitions concerning the appropriateness of 'knowledge' ascriptions (and denials) by referring to intuitions concerning the pragmatic effects of making 'knowledge' ascriptions (and denials).

More precisely, Nagel's approach explains why we have the intuition that high stakes subjects with modest evidence do not know the target propositions by the fact that we correctly judge that they do not outright believe the target proposition, or else, that their belief are unjustified (they are partially formed on the basis of compromising-accuracy factors). On the basis of the hindsight bias, this approach explains why, when the attributor feels high epistemic anxiety, she mistakenly denies 'knowledge' to those who in fact can know.

Nagel's proposal improves on Bach's account in that it can explain why confident low stakes assessors can (appropriately) judge that confident high stakes subjects with modest evidence should not assert that they know, without appealing to a problematic confusion between epistemic standards of practical rationality and epistemic standards of theoretical rationality. Indeed, on Nagel's view, this judgement is appropriate because it is true.

Still, although interesting, Nagel's account faces at least four problems. I now turn to consider them.

5.4.3 Four problems for Nagel's account

Problem 1: Stipulation that the confident high stakes subject's belief is epistemically well-formed

A first possible objection to Nagel's account appeals to a stipulation to the effect that the confident high stakes subject has formed her belief in an epistemically impeccable way. Still, according to this objection, it does not seem appropriate for this subject to say that she knows. Nagel's account cannot explain this.

²¹Nagel (2010, 409).

There are two ways in which one may suppose that the high stakes subject has formed her belief in an epistemically impeccable way. First, we may suppose the subject does not experience epistemic anxiety. Either this subject is psychologically abnormal, or she is merely practically irrational. Recall, indeed, that epistemic anxiety is determined by a strong desire for truth, which, in the case under consideration, is based on a practical consideration (the practical cost of error). If the subject does not care about the cost of error, she is merely practically irrational. Second, we may suppose that the confident high stakes subject has experienced high epistemic anxiety but has produced intense cognitive effort in assessing the evidence at her disposal.

Consider the first option. Suppose that the confident high stakes subject has formed her belief in exactly the same way as the confident low stakes subject.²² What does this stipulation amount to? According to Nagel, when the perceived stakes are high, the subject's epistemic anxiety rises. To calm down her epistemic anxiety, the subject must improve her belief-forming process, and as a result, she normally changes the method she uses. Thus, if the high stakes subject forms her belief exactly in the same way as the low stakes subject, that means that she does not feel epistemic anxiety. Hence, that means that she is psychologically abnormal (or perhaps practically irrational).²³

Indeed, it is a well-established fact in psychology that people do not appeal to the same belief-formation process (or method) when the stakes are low and when they are high.²⁴ If we suppose that the subject is psychologically normal (or practically rational), either she will not use the same method in the two cases, or if she uses the same method and her belief-forming process is not impaired by compromising-accuracy factors, given her epistemic anxiety, the belief won't be formed. In sum, according to Nagel, it is not psychologically coherent, or plausible, to assume that the subject uses the same method and believes in both cases, without assuming that the high stakes subject's belief has been formed on the basis of compromising-accuracy factors.

However, we may assume that the subject is psychologically abnormal (or practically irrational). In such cases, according to Nagel's view, we make an error in denying 'knowledge' to the subject (as we mistakenly ascribe epistemic anxiety to the subject). Still, do we make an error when we judge that it is inappropriate for this subject to assert that she knows? It does not seem so. Even if the subject does not feel epistemic anxiety, she should still see that the cost of error is high. And, we might think, this is sufficient to explain why she should not say that she knows (even it is true that she knows).

Thus, whether it is appropriate for this confident high stakes subject to say that she knows the target proposition does not (merely) depend on whether she believes the target proposition in the same way (epistemically speaking) as the confident low stakes subject. Therefore, if we want to

²²See also Fantl and McGrath (2009, 44-46).

²³See, for example, Nagel (2008, 281) on the idea that it is not psychologically plausible that a normal subject lacks epistemic anxiety when she perceives high stakes.

²⁴See Nagel (2010, 409-412)

defend the orthodox account of knowledge, the problem remains of explaining why such a claim to know from this high stakes subject is inappropriate. Indeed, that it still seems inappropriate for this confident high stakes subject to say that she knows might well suggest that the method or the evidence she uses is epistemically good enough to know in a low stakes case but not in a high stakes case.

Second, consider a confident high stakes subject who has calmed down her epistemic anxiety by a factor source of epistemic rationality. We can show that depending on the type of epistemically good factor we invoke, this confident high stakes subject may still not be in a better position to assert that she knows. This is problematic, if the intuition that it is inappropriate for such a subject to assert ‘I know’ is to be explained in terms of an intuition that she lacks knowledge, which, in turn, is to be explained by the presence of a compromising-accuracy factor affecting her belief-forming process.

To see this, let us grant, following Nagel, that there are two ways of improving our cognition: either by acquiring further evidence, or by thinking more thoroughly about it.²⁵ Let’s figure out a case where the high stakes subject has calmed down his epistemic anxiety by a cognitive effort consisting in assessing more thoroughly the modest (but knowledge-level) evidence she has. Suppose the subject’s epistemic anxiety is calmed down as a result of such a cognitive effort. If Nagel’s account is right, we should not think that in such a case the subject’s belief-formation process is impaired by a compromising-accuracy factor.

However, imagine that, in the high stakes case, Keith and his wife consider the evidence at their disposal more thoroughly. For instance, imagine they can take all the time they want to decide. They give a weight to each clue, and they sum the result in favour or against the fact that the bank is open. In this case, a huge cognitive effort is made. If the evidence they have is sufficient to know, this evidence should be sufficient for them to conclude that the bank is open. At least, if they outright believe the proposition on such a basis, they cannot be seen as being subject to a compromising-accuracy factor. Still, even after such a cognitive effort, they do not seem to be in a position to assert that they know the target proposition. If they have enough evidence to know, and hence, can have a justified outright belief, it is still unclear why they should not assert that they know.

To make the problem more vivid, compare with a pair of cases from Nagel:²⁶

HIGH STAKES, HIGH NEED-FOR-CLOSURE MACK. Mack is an oncologist. It’s late in the day, he is behind in his work, and needs to clear one more case before he leaves. He reviews some highly complex evidence, pronounces a diagnosis, and rushes out the door to go buy his husband a present before the shops close. The evidence was of a sort that doctors sometimes misread. Someone’s life hangs in the balance.

²⁵Nagel (2010, 414).

²⁶Nagel (2008, 291).

LOW STAKES, LOW NEED-FOR-CLOSURE JACK. Jack is Mack's twin brother, and has had identical medical training. He is calmly taking a practice test with no consequences for himself or others. He reviews some highly complex evidence of a sort that the doctors sometimes misread (say, the actual patient file from the Mack case). He is enjoying the experience of writing the practice test (a nice break from seeing patients), the room he is writing the test in is quiet, and he is not under any perceived time pressure. Jack eventually arrives at a diagnosis, and when he does so he feels the same level of confidence as Mack had.

In Mack's case, it is clear that there are compromising-accuracy factors, which can explain why we have the intuition that Mack does not know. This, in turn, explains our intuition that it would be inappropriate for him to say that he knows. However, there is a fundamental disanalogy between Mack's case and high stakes Keith's case. Modify Mack's case, so that he now takes his time, takes pleasure in doing his job, etc. Then, clearly, if Mack assesses the evidence in this way, he will be in a position to have a justified outright belief, and he will be in a position to say that he knows, even if we grant that he is in a high stakes situation. In Mack's initial case, there is no missing evidence. The only problem is how Mack assesses the evidence. This explains why we may have the intuition that Jack can justifiably outright believe the diagnosis he arrives at, and can appropriately claim to know. Indeed, Mack and Jack have the same evidence.

However, intuitions are clearly different in Keith's high stakes situation. Even if Keith and his wife thoroughly assess the evidence, they still seem not to be in a position to assert that they know that the bank is open. Their problem is that they lack evidence to assert that they know, not that they are cognitively biased or subject to a compromising-accuracy factor.

This strongly suggests that the reason why Keith should not appropriately say that he knows in the high stakes case does not essentially concern his psychological state and the cognitive effort required given his psychological state. Rather, it concerns whether the evidence is sufficient for Keith to appropriately say "I know". Nagel's account does not explain this.

Problem 2: The role of compromising-accuracy factors in our intuitions

A second problem is that Nagel's proposal relies on the idea that our (supposed) intuitions that the confident high stakes subject with modest evidence does not know is driven by the (hidden) presence of a compromising-accuracy factor that we postulate to make sense of the subject's outright belief. According to Nagel, the postulation of such a compromising-accuracy factor is enough to lead us to appropriately deny 'knowledge' to the subject.

However, there are cases where explicitly mentioning a compromising-accuracy factor in a high stakes case does not incline us toward the judgement that the subject does not know, but, rather toward the judgement that she knows (or, at least, toward the judgement that she can appropriately

say that she knows). Consider the following case from Schaffer (2006, 90):

HIGH-AND-FAST. On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, and indeed he has pressing financial obligations that require a deposit before Monday. His entire financial future is at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. As Sam is about to stop to double-check the bank hours, he remembers that he promised to buy a present for his wife. She will be furious if he forgets – his whole relationship is at stake. The stores are about to close. Sam must choose. So Sam makes a split-second decision to drive past the bank and pick up a present for his wife instead, thinking that after all, the bank will be open this Saturday. So, does Sam know that the bank will be open this Saturday?

This is a case where there is time pressure, an accuracy-compromising factor.²⁷ Still, we may have the intuition that the subject can appropriately assert that she knows that the bank is open (at least, we have the intuition that it is more appropriate for her to say that she knows in this situation than if she was not short of time).

This is problematic for Nagel's approach. Suppose that the intuition that it is inappropriate for a confident high stakes subject with modest evidence to say that she knows is to be explained in terms of the subject's absence of knowledge, which, in turn, is to be explained by the presence of a hidden compromising-accuracy factor. On this view, when a compromising-accuracy factor is made explicit, we should not be inclined to judge that this subject knows (and hence, we should not be inclined to judge that she can appropriately say that she knows). At least, our intuition should be somewhat driven toward the judgement that she does not know (and, hence, toward the judgement that she should not say that she knows).

Therefore, it seems that the intuition that a confident high stakes subject does not know (and inappropriately says "I know") is not merely driven by the (hidden) presence of accuracy-compromising factors. If this is correct, the fact that a high stakes subject outright believes a proposition on the basis of modest evidence because she is affected by the presence of some compromising-accuracy factors is insufficient to explain the intuition that this subject does not know (or should not say "I know"), and Nagel's proposal is unsatisfactory in this respect.

Problem 3: Why should we deem the subject epistemically irrational rather than practically irrational?

A third problem for Nagel's approach is that it relies on the idea that we should (rightly) see the confident high stakes subject as having an unjustified belief, on the basis of the fact that we have to

²⁷See Nagel (2010, 291) for the claim that pressure of time is a compromising-accuracy factor.

postulate the presence a compromising-accuracy factor to make sense of her outright belief. Indeed, the subject is supposed to experience epistemic anxiety, and hence, unless she is affected by the presence of a compromising-accuracy factor, she supposedly cannot outright believe the proposition on the basis of modest evidence.

We should note, however, that in this framework we might as well make sense of the subject's outright belief by seeing this subject as practically irrational. Suppose that the subject is practically irrational, and does not care about the cost of error. Hence, she is not epistemically anxious.²⁸ If so, we can easily explain why she outright believe that p , without postulating the presence of a compromising-accuracy factor. But then, there is no reason to think that her belief is epistemically unjustified, and hence, no reason to think that she does not know the target proposition.

Of course, if the subject knows the target proposition, we cannot explain why it is inappropriate for her to say that she knows this proposition by appealing to the idea that this claim is false. Hence, if the subject knows the target proposition, Nagel's approach fails to explain our linguistic intuition that the subject should not say that she knows.

Is there a reason to think that the subject is epistemically irrational (in having an outright belief on the basis of a compromising-accuracy factor), rather than pragmatically irrational (in neglecting the cost of error)? We may understand why we should ascribe irrationality rather than psychological abnormality. It is *prima facie* less plausible that the subject is psychologically abnormal rather than not fully rational. However, there is no explanation why we should ascribe epistemic irrationality rather than practical irrationality. There is no reason to think that one is more often epistemically irrational than practically irrational. Further, ascribing epistemic irrationality is more costly. Whereas we need to postulate the presence of compromising-accuracy factors to see the confident high stakes subject as epistemically irrational, we need not postulate the presence of this kind of factors to see the subject as practically irrational.

To illustrate this point, consider what Nagel writes:

If we see the wife's request for additional evidence as rational, [Keith]'s refusal to take it seriously points to some epistemically problematic disposition in him. (Nagel, 2012, 682)

The objection here is precisely the following: why would this refusal point toward an *epistemically* problematic disposition (concerning the rationality of beliefs) rather than a *practically* problematic disposition (concerning the rationality of action)? If we see the request as *epistemically* rational, of course, we might think that Keith's refusal to take it seriously points to some epistemically problematic disposition in him. But if, as non-sceptical orthodox philosophers hold, this request

²⁸Nagel grants that insensitivity to epistemic anxiety can sometimes be explained in terms of practical irrationality. For example, she writes: "the underlying reason why we should feel increased epistemic anxiety in high-stakes circumstances has to do with what is pragmatically desirable: the pragmatic benefits of increased accuracy offset the pragmatic costs of increased effort and time spent searching for evidence. If the (...) subject does not feel high anxiety in her current predicament, he is at risk of failing to maximize utility." (Nagel, 2010, 427)

is merely practically rational (the alternative raised by Keith's wife – that the bank may have changed its hours – is not knowledge-destroying), we should not see Keith's refusal as epistemically problematic. If we do not see his refusal as epistemically problematic, we need not see the fact that he knows or outright believes as problematic. But that does not mean, of course, that we should not see his refusal as *practically* problematic. And that does not mean that, if he asserts that he knows, we should see his assertion as appropriate (because true). However, again, Nagel's approach does not have the resources to explain the intuition that Keith should not assert that he knows that the bank is open even if he is (intuitively) taken to know that the bank is open.

We can see the problem that Nagel faces as a dilemma. On the one hand, suppose that Nagel's proposal relies on the idea that we should judge the confident high stakes subject as epistemically irrational (which, in turn, is supposed to explain our intuition that it would be inappropriate for her to claim to know). If so, it remains to explain why we should deem the subject as epistemically irrational rather than practically irrational. Indeed, this evaluation is surprising if there is a sharp distinction between epistemic and practical rationality, and if ascribing epistemic irrationality is more costly (we need to postulate the presence of a compromising-accuracy factor). An advocate of pragmatic encroachment (or epistemic contextualism) may put forth that, precisely, this evaluation is explained if the truth-value of 'knowledge' ascriptions is affected by practical factors, but unexplained otherwise.

On the other hand, suppose that Nagel's proposal does not necessarily rely on the idea that we should judge the confident high stakes subject as epistemically irrational. If so, this proposal cannot explain the remaining intuition to the effect that the epistemically rational confident high stakes subject with modest evidence should not assert that she knows. Indeed, on Nagel's approach, we are supposed to explain our intuitions concerning the inappropriateness of 'knowledge' ascriptions (concerning confident high stakes subjects with modest evidence) in terms of intuitions about epistemic irrationality.

In sum, either Nagel's proposal is insufficient because it does not explain why we should deem the subject as epistemically irrational rather than practically irrational, or it is insufficient because it does not explain why we can assess the subject as epistemically rational while inappropriately asserting that she knows.

Problem 4: Does knowledge require the absence of epistemic anxiety?

A fourth problem is that Nagel's account supposes that the psychological attitude required by knowledge is incompatible with epistemic anxiety. However, one might think that although the high stakes subject feels epistemic anxiety, that does not prevent her from knowing the target proposition. For instance, it might be that Keith knows the target proposition in the high stakes case, even if he naturally feels increased epistemic anxiety. If so, Nagel's proposal cannot explain why it seems inappropriate for him to say that he does not know.

As she writes:

[A] critic might agree that the anxious person will keep collecting evidence longer than his casual counterpart, while still thinking that this anxious person would nonetheless be seen as already having an outright belief as soon as he has the moderate level of evidence possessed by his counterpart. This worry is hard to answer without taking on the task of defending some particular view of what it is to have an outright belief. (Nagel, 2010, 417)

Nagel defends the idea that outright belief – the attitude required by knowledge – is incompatible with epistemic anxiety by considering in particular the idea that, on influential functionalist accounts of outright belief, there is a practical condition on outright belief, such that S outright believes that p only if, in some sense, S takes p as true in her practical reasoning, or act (/ prefer) as if p.²⁹ This suggests that if S outright believes that p, then S is not epistemically anxious.

Indeed, we may argue, if S acts as if p, S does what it is rational to do conditional on p (at least on the assumption that S is practically rational). But, typically, the rational thing to do given p (for example, the action that has the highest expected utility given p) is not to check whether p. Thus, typically, if the subject acts as if p, she is not inclined to check whether p. Hence, if S is epistemically anxious about p, she does not outright believe that p, and does not know that p.

However, we should note that this pragmatic account of outright belief is incompatible with a non-sceptical orthodox account of knowledge. First, suppose that knowledge requires an outright belief, and suppose that if S outright believes that p, then S is willing to act as if p in all practical contexts in which p is relevant to action. On this context-insensitive pragmatic view of belief, we face the problem of scepticism. There are very few (if any) propositions you believe in such a way that, in all practical situations (where they are relevant to action), you would be willing to act as if there were true. In some cases, the cost of error is extremely high, and you would not be willing to act as if there were true.

Second, suppose that, to avoid scepticism, we say that if S outright believes that p in certain circumstances, then S is willing to act as if p *at least in these circumstances*. This means that, in different circumstances, where the cost of error is higher and you are not willing to act as if

²⁹According to Weatherston (2005, 421): “To believe p is to have preferences that make sense, by your own lights, in a world where p is true.” More precisely, according to Weatherston, S outright believe that p if S has a sufficiently high credence in p, such that her preferences (over a relevant class of action) conditional on a relevant class of propositions do not change if she conditionalises on p. According to Ganson (2008, 453), to outright believe that p, an agent must believe that p to a high enough degree such that she is willing to act as if p in the circumstances, where acting as if p means doing what she would be willing to do given p (and not necessarily what it would be rational for her to do given p, although in cases in which the subject is practically rational, the two come to the same thing). According to Fantl and McGrath (2009), S outright believes that p if S’s credence in p is high enough for p to be her motivating reason. On all these views, outright belief is reducible to credences and a pragmatic condition. Some other philosophers distinguish outright belief and credences. Thus, Ross and Schroeder (2014) argues that to outright believe that p is to have a defeasible disposition to treat p as true in practical reasoning (i.e., to evaluate the options as if p were certain).

p, you do not outright believe that p. But this context-sensitive pragmatic view of belief leads to pragmatic encroachment on knowledge: whether you are justified in having an outright belief varies with your practical situation

To see this, suppose we think of outright belief in terms of degrees of belief (credences) sufficiently high to be willing to act on the proposition in the circumstances. Given that we may also assume that a certain amount of evidence justifies the same degree of belief in all practical circumstances, it follows that whether a certain amount of evidence justifies an outright belief that p (i.e., justifies a degree of belief sufficiently high to be willing to act on p) depends on the practical context.³⁰

Alternatively, suppose that we think that outright beliefs are irreducible to credences (and a pragmatic condition on these credences). We may rather suggest that outright beliefs are dispositions to treat propositions as true in practical reasoning, i.e., dispositions to evaluate the options without considering the possibility that these propositions are false.³¹ It is important to stress that treating p as true in reasoning is different from having a sufficiently high credence in p to motivate acting as if p. When you treat p as true in your reasoning, you evaluate your options as if p were certain. Thus, if we represent your practical problem by a decision table, the propositions that you treat as true in your reasoning are those that are assumed. For example, the proposition that if you perform a certain action A, and the state of the world is p, then the outcome is O, is treated as true in your reasoning. In contrast, if you represent the possible state of the world not-p, and assign to this possibility a certain degree of probability, you do not treat p as true. That is, on this view, you do not outright believe that p, even if the probability that p is so high that you should do what is best given p (act as if p).

Now suppose that knowledge requires being justified in occurrently believing outright that p, that is, being justified in treating p as true in your practical reasoning. Whether you are justified in treating p as true in your practical reasoning depends on your practical situation. Again, if the cost of error about p is high, intuitively, you should not treat p as true in your reasoning. You should consider the possibility that not-p. If so, whether you know that p depends on your practical situation.³²

In sum, if we accept the idea that outright belief is incompatible with epistemic anxiety, either because we understand outright belief in terms of credence sufficiently high to act on the target proposition, or in terms of a disposition to treat the proposition as true in practical reasoning, we are led either to scepticism or to pragmatic encroachment. We are led to scepticism if we grant that outright belief requires a sufficiently high credence to act on p, or to treat p as true, in all practical circumstances. We are led to pragmatic encroachment if we grant that outright belief requires a sufficiently high credence to act on p, or to treat p as true, in the present circumstances only.

Still, Nagel is right that it is plausible to think that there is a pragmatic condition on outright

³⁰See Weatherston (2005); Ganson (2008); Fantl and McGrath (2009).

³¹See Ross and Schroeder (2014).

³²See Ross and Schroeder (2014).

belief. But some other philosophers have suggested a weaker pragmatic condition, consistent with epistemic anxiety and a non-sceptical orthodox view of knowledge.

To begin with, it pays to note that there are practical circumstances in which we would not be willing to act as if p (treat p as true) although, intuitively, we outright believe that p in these circumstances. Suppose you are offered a bet on whether Dublin is the capital of Ireland. If you take the bet and Dublin is the capital of Ireland, you win one dollar. If you take the bet and Dublin is not the capital of Ireland, you are tortured to death. The rational thing to do for you, it seems, is to decline the bet. Still, given that Dublin is the capital of Ireland, the best option is to take the bet. Therefore, if outright belief in some circumstances requires being willing to act as if p (treat p as true) in these circumstances, if you decline the bet, we should say that you do not outright believe that Dublin is the capital of Ireland in these circumstances. But this verdict is counterintuitive. As Wedgwood notes:

In this case, it would be natural for you to say, “Look: I’m not going to take this bet. Even though I definitely believe that Dublin is the capital of Ireland, I’m not so fantastically confident of that to make it rational for me to take this insane bet!”
(Wedgwood, 2012, 321)

Wedgwood offers a pragmatic account of outright belief compatible with the idea that you can outright believe that p in circumstances C even if you are not willing to act as if p in C . On his view, to outright believe that p is to have a disposition to assign a *practical* credence of one to p *for all normal practical purposes* (i.e., to have a disposition to treat p as certain in the practical reasoning in normal contexts). You may assign maximal practical credence to p without assigning maximal *theoretical* credence to p . In this way, scepticism is avoided. More precisely, theoretical credences register your degree of justification for p , whereas practical credences determine your intentions and actions (how you reason to make a decision). Your practical credences can vary across contexts. For the purpose of deciding whether to take the bet above, your practical credence in the proposition that Dublin is the capital of Ireland is less than one. For normal purposes, however, it is one. In contrast, your theoretical credences are invariant. They determine the practical credences it is rational to assign to a proposition given the practical context. The fact that you do not have maximal justification for the proposition that Dublin is the capital of Ireland explains why you do not have a practical credence one in this proposition when you face such a bet. But even if you do not assign a practical credence one to this proposition in these circumstances, you still are disposed to assign a practical credence one to this proposition in all normal contexts. Therefore, you still outright believe this proposition in this context.

Another pragmatic account of outright belief on which you can outright believe that p even if you do not act as if p is proposed by Williamson (2000). On this view, you outright believe that p if

you are willing to rely on p in practical reasoning.³³ Still, there are degrees of outright belief. When the stakes are high, if your degree of outright belief is not strong enough, you can outright believe that p even if you do not rely on p in this case. Thus, we may propose the following pragmatic condition on outright belief: S outright believe that p only if S is disposed to treat p as true in her practical reasoning when the stakes are low.

Is there a reason to prefer this weaker pragmatic account of outright belief, other than the fact that this weaker account is compatible with a non-sceptical orthodox account of knowledge? I think so. In the following chapter (section 2.3), I will argue on the basis of independent grounds that knowledge does not entail certainty (psychological or epistemic). Now, it seems that epistemic anxiety just is psychological uncertainty. So, if it is correct that knowledge does not entail psychological certainty, and that knowledge entails outright belief, it follows that outright belief, and knowledge, are compatible with a certain degree of epistemic anxiety.

In summary, we may grant that there is a pragmatic condition on outright belief. But this does not imply that outright belief is incompatible with epistemic anxiety. It is plausible to think that in abnormal (or high stakes) cases, you can outright believe that p even if you are not willing to rely on p in these circumstances. On the other hand, if we grant that an outright belief that p in C implies being willing to act as if p in C (/ treat p as certain in C), or experiencing no epistemic anxiety as to whether p in C , either we should be sceptic, or we should endorse pragmatic encroachment. Since Nagel's account relies on the claim that outright belief in C is incompatible with epistemic anxiety in C , it does not provide an explanation compatible with a non-sceptical orthodox account of knowledge.

5.5 Conclusion

There may well be cases where an approach in doxastic terms correctly explains our intuitions that a high stakes subject with modest evidence lacks knowledge, and hence, should not say that she knows. When we ourselves are in a normal doxastic situation with modest evidence, it is possible that we often correctly judge that a high stakes subject with the same evidence, and on the basis of a normal cognitive effort, would not outright believe the target proposition, or that if she does believe, her belief is formed in a way affected by accuracy-compromising factors (at least if she is a psychologically normal subject). When we ourselves are in an abnormal doxastic situation, it is possible that we often mistakenly judge that a low stakes subject with modest evidence, and on the basis of a normal cognitive effort, cannot know the target proposition, and hence, should not say that she knows.

However, this approach cannot explain all the intuitions. In particular, we may well think that a confident high stakes subject possessing modest (but knowledge-level) evidence has not formed her

³³See chapter 3, section 3.2 for more details.

belief on the basis of a process affected by compromising-accuracy factors, or that these factors are not sufficient to prevent her from knowing. Even so, it still seems that this subject should not assert that she knows. If she knows, this approach does not explain why it would be inappropriate for her to assert that she knows. The main problem of the doxastic approach is that it aims at explaining a pattern of intuitions concerning the appropriateness of ‘knowledge’ ascriptions and denials on the basis of a pattern of intuitions concerning whether the subject under evaluation knows. Thus, in contrast to a Gricean approach, the doxastic approach cannot explain a judgement that a high stakes subject with modest evidence can know that p but should not say “I know”. In the next chapter, I shall defend a new approach relying on the norm of assertion. I will argue that this approach can accommodate this judgement while avoiding the problems of the Gricean approach discussed in chapter 4.

Chapter 6

Linguistic intuitions and the certainty norm of assertion

6.1 Introduction

I have argued that the Gricean account and the doxastic account do not satisfactorily explain our intuitions as far as ‘knowledge’ ascriptions and denials are concerned. The Gricean account aims at explaining these intuitions in terms of the distinction between the semantic and the pragmatic meaning of such attributions and a specific conversational context. The doxastic account aims at explaining these intuitions in terms of the assessor’s (sometimes mistaken) intuitions about the subject’s knowledge. In this chapter, I shall defend a mixed account appealing to a doxastic approach and to the claim that certainty is the norm of assertion.

First, I present the certainty account of assertion. I argue that it does not lead to scepticism about warranted assertion, because, first, it is unclear that ‘certain’ is a maximum-standard absolute gradable adjective, and second, it is a context-sensitive word.

Second, I present the main arguments in favour of the certainty account of assertion. This account can easily explain in a unified way a wide range of linguistic and conversational data having to do with Moore-paradoxical sentences, appropriate challenges and parenthetical uses. To these arguments, I add six considerations showing that the certainty account of assertion is the right account.

Third, I defend this account against two objections. The first one is put forth by Pritchard (2008) and states that a certainty norm is too strong. The second one is proposed by Turri (2010) and states that, in comparison with the knowledge account, the certainty norm does not do well in explaining how we prompt and challenge assertions.

Finally, I show that if we adopt the certainty account of assertion, we can propose a mixed account of our shifting intuitions concerning the appropriateness of ‘knowledge’ ascriptions, appealing partially to a doxastic approach and partially to the certainty norm of assertion. I show that this mixed approach outperforms the simple doxastic account criticised in chapter 5 and the Gricean account considered in chapter 4.

As this mixed approach is compatible with the non-sceptical orthodox account of knowledge, I conclude that shifting ‘knowledge’ ascriptions can be explained in an orthodox framework.

6.2 The certainty norm of assertion

6.2.1 The epistemic norm of assertion

Assertions are speech acts. As any other action, speech acts are governed by norms. An assertion, for example, can be imprudent, rude, irrelevant, etc. It has been argued that, besides these practical norms, speech acts are governed by specific epistemic norms. It is said, in particular, that we can individuate the different speech acts in terms of the specific epistemic norms applying to them.¹ Thus, for example, one may think that S’s question whether p is good or right *qua* question only if S does not know whether p, that S’s assertion that p is good or right *qua* assertion only if S knows that p, etc.

It is important to emphasize that the claim that there is an epistemic norm governing a certain speech act does not imply that if this speech act is performed and the norm in question is violated, then the subject is, overall, criticisable for making such a speech act. For example, suppose that knowledge is the epistemic norm of assertion. That means that you ought to assert that p only if you know that p. But if you assert that p without knowing that p, thereby violating the norm, even if your assertion *qua* assertion is incorrect, it does not follow that you are criticisable. It can be overall required to make an incorrect assertion. Perhaps making this assertion has good effects – concerning for example politeness – outweighing the fact that the assertion is epistemically wrong.

It is also sometimes suggested that the epistemic norm is constitutive of the speech act that it governs.² Consider, for example, the rules of chess. They constitute what it is to play chess in that to play chess is to make certain moves that can be evaluated as right or wrong in view of these rules. Necessarily, your move is a move in a play chess only if it is assessable as right or wrong in light of the rules of chess. In the same way, it is sometimes said, a particular epistemic norm constitutes what it is to assert something rather than, for example, recommending: a speech act is an assertion only if, necessarily, this speech act is correct only if its content is known (or believed, etc.)

¹See Williamson (2000, ch. 11).

²See Williamson (2000, ch. 11).

The claim that a certain epistemic norm is constitutive of a certain speech act does not mean that if we try to make this speech act but the epistemic norm in question is violated, then the speech act is not performed. This only means that if this speech act is performed, it is criticisable as a bad exemplar of this kind of speech act. Likewise, for example, a bad cooker is still a cooker, but who is bad as a cooker insofar as she lacks some property that a perfect cooker should have.

The epistemic norm supposed to govern the associated speech act is somehow related to the success conditions of the speech act in question, that is, to the conditions that must obtain for this speech act to be achieved. For example, we may think that S asserts that p only if S says “p” with the intention of conveying, or implying, that she believes, or knows, that p (and with the intention that the audience sees that intention). Making a statement “p” is insufficient for assertion. When an actor plays a role in a movie, he does not assert the propositions expressed by the sentences that he utters, for he does not intend to convey that he believes, or knows, them. Arguably, if the speaker fails to convey that she says “p” with the intention of conveying that she believes or (takes herself to) know(s) the proposition that p, the speaker fails to assert that p. But the speaker can succeed even if she violates the supposed norm. Thus, the claim that there is an epistemic norm of assertion, say, knowledge or belief, must be distinguished from the claim according to which, by asserting that p, the speaker hereby represents herself as believing, or knowing, that p.

In what follows, I will leave aside the questions surrounding the nature of assertions, and whether assertions are governed by an epistemic norm which is individuating (i.e., which applies only to assertions) and constitutive (i.e., in virtue of which a given speech act counts as an assertion).³ I will just assume that assertions are subject to an epistemic norm and that you can make assertions even if you violate the epistemic norm, like when you lie (in such cases your assertion is epistemically bad but it still is an assertion). This assumption is widely shared among philosophers nowadays. My aim is merely to argue, following Stanley (2008), that the epistemic norm of assertion is certainty.

6.2.2 Subjective and epistemic certainty

According to the certainty account of assertion, certainty is the norm of assertion. This proposal will strike many as implausible, and it does not have many advocates nowadays. Philosophers have focused on other potential norms, such as a rational credibility norm, a knowledge norm, a warrant norm, or a truth norm.⁴ Indeed, it is generally thought that a certainty norm is a non-starter because it is obviously too strong. However, I will argue, correctly understood, this norm is not too strong.

First of all, we should note that the idea that there is a connection between assertion and certainty is not new. It was already present in Moore’s “Certainty”. Thus, after making a certain number of assertions, like “I am not absolutely naked”, Moore writes the following:

³See Engel (2013) and Pagin (2015) for further discussion.

⁴See, respectively, Douven (2006), Williamson (2000), Gerken (forthcoming) and Weiner (2005).

Now I have made a number of different assertions; and I have made these assertions quite positively, as if there were no doubt whatever that they were true. That is to say, though I did not expressly say, with regard to any of these different things which I asserted, that it was not only true but also *certain*, yet by asserting them in the way I did, I *implied*, though I did not say, that they were in fact certain – implied, that is, that I myself knew for certain, in each case, that what I asserted to be the case was, at the time when I asserted it, in fact the case. (Moore, 1959)

More recently, Stanley (2008) has explicitly proposed the idea that certainty is the norm of assertion. Stanley’s aim in this paper is to defend the idea that knowledge does not imply certainty. By endorsing the certainty norm of assertion, Stanley seeks to explain pragmatically why utterances of ‘I know that p but I am not certain that p’ are odd.

Now, it seems that a certainty norm is quite strong, and perhaps too strong. Before coming to this objection, however, it is necessary to say a few words about the notion of certainty.

How should we understand the notion of certainty? In general, the notion of certainty has two aspects, a doxastic (or subjective) aspect and an epistemic aspect. Subjective certainty can be characterised as follows: You are subjectively certain that p if you have the highest degree of confidence in p. For example, we refer to subjective certainty when we say “Keith is certain that the bank is open”. What we say is that Keith has the highest degree of confidence toward the proposition that the bank is open.

Epistemic certainty can be characterised as follows: It is epistemically certain for you that p only if you are in a position to know that p on the basis of the highest degree of justification (or if p is maximally justified for you). For example, when we say “It is certain for Keith that the bank is open”, we express the proposition that Keith is in a position to know that the bank is open on the basis of a justification such that, in some sense, there is no better justification.

According to the certainty account of assertion proposed by Stanley (2008), certainty is the epistemic norm of assertion in the subjective and the epistemic sense. That is, on this view, you ought to assert that p only if you are *subjectively* certain that p, and only if p is *epistemically* certain for you. It is plausible that either the epistemic certainty norm, or the subjective certainty norm, is more fundamental. As Stanley notes, we might think that the subjective certainty norm is derived from the epistemic certainty norm, or vice versa, given a plausible assumption regarding the normative relation between epistemic and subjective certainty.⁵

Thus, we may suppose that subjective certainty is the norm of assertion, but that epistemic certainty is the norm of subjective certainty, in the sense that you ought to be subjectively certain that p only if p is epistemically certain for you. It follows that, as a matter of fact, there is something wrong with someone who asserts something which is not epistemically certain for her, namely, that

⁵See Stanley (2008, 51-52).

she should not be subjectively certain that this proposition is true. In this case, the relevance of epistemic certainty for assertion is derived from the subjective certainty norm for assertion and the epistemic certainty norm for subjective certainty.

Alternatively, we may suppose that epistemic certainty is the norm of assertion, but that it is also the norm of subjective certainty, in the sense that you ought to be subjectively certain of the propositions you consider and that are epistemically certain for you. If so we may explain why there is something wrong with someone who asserts something without being subjectively certain of it. Although the assertion is appropriate, the assertor violates a norm governing subjective certainty.

Be that as it may, on this view, when you assert that p , you represent yourself as knowing that p . Indeed, epistemic certainty requires being in a position to know, and subjective certainty requires believing outright. Further, on the assumption that the assertor is rational, she correctly bases her subjective certainty on the factors that make the proposition certain for her. However, if knowledge does not entail (epistemic and/or subjective) certainty, knowing that p is not sufficient to satisfy the certainty norm of assertion with respect to p .

6.2.3 Knowledge and certainty

The claim that knowledge does not entail certainty, in either of these senses, is shared by most contemporary epistemologists. However, the main reason why most philosophers think that knowledge does not entail (subjective and/or epistemic) certainty is that this would lead to scepticism about knowledge. If this is correct, we should also think that endorsing the certainty norm for assertion leads to the idea that most of our assertions are unwarranted, that is, to what I will call ‘scepticism about assertion’. But why think that if knowledge requires certainty, then scepticism about knowledge is true?

Consider subjective certainty. As explained in chapter 3 (section 3.2), one might think that if you are subjectively certain that p , you are willing to bet on p whatever the odds. Further, suppose that certainty is the highest degree of confidence. Then, it seems, if I am more certain that p than that q , then I am not certain that q . I am more certain that I exist than I am of most of the other things (e.g. that Paris is in France). Consequently, I am not certain of these latter propositions. If knowledge requires certainty, then, it seems, we should be sceptic about most of our knowledge.⁶ If we in fact know many things, knowledge does not require subjective certainty.

Consider epistemic certainty. It seems that if you need to have maximal justification for p in order to know that p , then you do not in fact know many things. For example, it seems that you merely have non-entailing grounds for the proposition that you have hands. Your justification would be better if you had entailing evidence. So, you do not have maximal justification for the proposition that you have hands. If knowledge requires maximal justification, you do not know

⁶See Unger (1975, ch. 2).

that you have hands. Also, if knowledge requires maximal justification, it seems impossible for you to know anything unless you have an infinite mind. Indeed, your epistemic position with respect to p can always be improved by adding further evidence for p , or by improving the reliability of the belief-forming process, or by acquiring higher-order justified belief about the fact that your first-order belief that p is justified, or about the fact that your belief that p is justified is justified, etc.⁷ Therefore, if scepticism about knowledge is false, knowledge does not require epistemic certainty.

Now, if subjective and epistemic certainty are not (rationally) achievable, one might worry that the certainty account of assertion renders all our assertions unwarranted. That is, the certainty norm for assertion would be too strong.

However, the notions of certainty that the certainty account of assertion uses need not be these notions leading to sceptical worries. There are many different ways of thinking about certainty. For example, against scepticism about subjective certainty, it has been claimed that being subjectively certain that p does not imply being willing to bet on p whatever the odds, and that we are as certain that we exist as we are that, say, there are cars.⁸ Against scepticism about epistemic certainty, it is sometimes argued that we should distinguish a weak and a strong notion of epistemic certainty, and that only the second one leads to sceptical worries.⁹

Also, taking inspiration from Williamson, we may want to distinguish between degrees of uncertainty (partial beliefs and partial justification) and degrees of certainty (degrees of outright belief and degrees to which it is rational to rely on p in practical reasoning). On this approach, certainty

⁷See Brown (2011, 161-162) and Fantl (2003).

⁸See for example Miller (1978, 53). According to him, absolute conviction that p requires to unhesitatingly base his actions on p in normal circumstances (when the stakes are not abnormally high) and to be unwilling to accept correction, understood as a shift from the belief that p to the belief that not- p . You are unwilling to accept correction when, were you to doubt that p , you would doubt that there is any means to come to decide whether p or not p . It seems that we have this attitude toward many propositions (e.g., that there are cars, that we have hands, etc.). Klein (1981, 91-92) argues that psychological (or subjective) certainty does not lead to dogmatism.

⁹For example, Malcolm defends the idea that many propositions concerning the external world are certain for us at time t in that our epistemic position with respect to these propositions at t is such that we can imagine no future negative test that could make believing not- p at t warranted, although we could imagine future negative tests that could make believing p less warranted at t (see Firth (1967)). Miller (1978) argues that absolute epistemic certainty does not require absolutely conclusive justification, that is, a super-justification to believe p such that, whatever the future evidence, one always have reason to believe p and no reason to believe not- p . What absolute epistemic certainty merely requires is, rather, immunity against future disconfirmation. Thus, p can be epistemically certain for S even if not- p may be a reasonable alternative for S in the future. Klein (1981, 123) argues that “ p can have the highest degree of epistemic warrant on the basis of some evidence, e , without e entailing p ”. Thus, he distinguishes two senses of ‘certainty’, one having to do with nondefective justification that p , roughly, justification confirming p that is “attack proof” in the actual world (such that it cannot be weakened by a proposition that is true in the actual world), and the other having to do with justification logically impossible to defeat (roughly a justification that is attack proof in all the possible worlds) (See in particular Klein (1981, 178-189). While, according to Klein, knowledge entails certainty in the first sense, it does not entail certainty in the second sense, and there are many propositions that are certain for S in the first sense. See also Firth (1967) for a distinction between the truth-evaluating, the warrant-evaluating and the testability-evaluating uses of the notion of epistemic certainty. In particular, Firth (1967, 10-12) shows that there are many ways of construing the notion of “maximum warrant” or “highest degree of credibility”. Not all of them entail that a maximally warranted proposition is true, or entailed by the evidence the subject has, or better justified than any other proposition, or such that one cannot imagine to be better warranted in believing this proposition, etc.

should be distinguished from absolute certainty. Absolute subjective certainty is the highest possible degree of outright belief, and absolute epistemic certainty is the highest possible degree to which you can rationally rely on p in practical reasoning. But subjective certainty merely requires outright belief (being willing to use p as a premise in practical reasoning when the stakes are low), and epistemic certainty merely requires sufficient justification for outright belief (evidential probability one).¹⁰ We can then say that certainty thus understood, and not absolute certainty, is the norm of assertion.

Yet, on these kinds of non-sceptical approaches to the notion of certainty, which appeal to a distinction between certainty and absolute certainty, it is unclear whether there is a real difference between knowledge and certainty, and hence, it is unclear why we should think that knowledge does not entail certainty. And if so, it is unclear that there is a real difference between the knowledge norm for assertion and the certainty norm for assertion. In sum, it seems that we face a dilemma: either knowledge and certainty differ, but then certainty seems too strong to be the norm of assertion. Or we have a non-sceptical notion of certainty, but then knowledge and certainty do not differ, and the certainty norm of assertion is identical to the knowledge norm.

Still, it pays to note that there are degrees of certainty, whereas, arguably, there are no degrees of knowledge. If this is correct, we can distinguish the certainty norm for assertion from the knowledge norm by invoking this gradability. Even if we grant that knowledge entails certainty (i.e., on this approach, justified outright belief), and certainty entails knowledge, knowledge does not entail the satisfaction of all the degrees of certainty (i.e., all the degrees of outright belief and epistemic justification). Thus, we can propose to reformulate the certainty norm of assertion as follows: you ought to assert that p only if p is sufficiently certain (epistemically and subjectively), or certain to a sufficient degree. What ‘sufficient’ means depends on the context of the assertion. In this way, the certainty norm and the knowledge norm differ – knowledge is not always sufficient for assertion – and scepticism about assertion is avoided – the certainty norm can often be satisfied.

On this model, knowledge entails a minimal degree of certainty (i.e., justified outright belief), and as we distinguish certainty and absolute certainty, scepticism about knowledge is avoided. Yet, there are reasons independent from the threat of scepticism to think that knowledge does not entail certainty, even a minimal degree of certainty.

First, it seems that there are cases of knowledge without certainty. It is clear that knowledge entails outright belief, but as I have argued previously (see ch. 5 section 4.3 problem 4), it is unclear that outright belief entails the absence of uncertainty. Again, suppose you are offered a bet on whether Dublin is the capital of Ireland. If you take the bet and Dublin is the capital of Ireland, you win one dollar. If you take the bet and Dublin is not the capital of Ireland, you are tortured to death. Presumably, you will decline taking such a bet.

I concur with Wedgwood (2012) in saying that, in this case, you presumably outright believe the

¹⁰See chapter 3 section 3.2 for more on Williamson’s approach.

proposition, although you are not willing to take the bet because you are not subjectively certain that Dublin is the capital of Ireland (in Wedgwood's terms, your practical credence is less than one in the circumstances, but it would be one in normal circumstances; see chap. 5 section 4.3 problem 4). And, presumably, you are not subjectively certain, in these circumstances, that Dublin is the capital of Ireland because this proposition is not epistemically certain for you (in Wedgwood's terms, your theoretical credence is less than one). Yet, it seems to me that you can know, in these circumstances, that Dublin is the capital of Ireland. If this is correct, knowledge does not require epistemic certainty either.

Another example in which we ascribe 'knowledge' to an unconfident subject is the following:¹¹

STUDENT. An examiner asks a student when a certain battle was fought. The student fumbles about and, eventually, unconfidently says what is true: The Battle of Hastings was fought in 1066. It is supposed, quite properly, that this correct answer is a result of the student's reading. The examiner, being an ordinary mortal, and so unconfident of many things himself, allows that the student knows the answer; he judges that the student knows that the Battle of Hastings was fought in 1066

It is clear that the student is uncertain, and yet the examiner seems to correctly judge that the student knows the answer.

Second, there are linguistic considerations speaking against the claim that knowledge entails certainty. Thus, Stanley (2008) notes that uttering sentences the truth of which entails that knowledge is compatible with the absence of truth, belief, or evidence – the traditionally accepted necessary conditions on knowledge – sound like contradictions, whereas uttering sentences the truth of which entails that knowledge is compatible with uncertainty seems acceptable. Consider the following claims:

- (1) There is something someone knows that isn't true.
- (2) There is something someone knows that she doesn't believe.
- (3) There is something someone knows that she doesn't have any reason to believe.
- (4) There is something someone knows, of which she is uncertain.

(1), (2) and (3) seem obviously false, whereas (4) is not obviously false.

Further, there is an asymmetry between these two sentences:

- (5) I know that *p* but I am not certain that *p*
- (6) John knows that Bush is a Republican, though, being a cautious fellow, he is somewhat uncertain of it.

¹¹See Unger (1975, 83-84) and Stanley (2008, 41).

(5) seems odd, whereas (6) seems more acceptable. But if knowledge semantically entails (subjective) certainty, there should be no difference between first-person and third-person cases.

Moreover, consider the following sentences:

(7) Everything anyone knows is true.

(8) Everything someone knows she believes.

(9) If someone knows something, she has a reason to believe it.

(7), (8) and (9) look like trivial truths.

In contrast,

(10) I'm certain of everything I know.

(11) Everyone is certain of everything she knows.

do not look like trivial truths. This suggests, again, that knowledge does not entail (subjective) certainty.

Similar intuitive verdicts concern epistemic certainty:

(12) Everything I know is certain to be true.

(13) There are some things I know, which are somewhat uncertain.

(12) does not look like a banal truth and (13) does not look like an obvious contradiction.

In sum, it is far from clear that knowledge is linked to subjective and epistemic certainty as closely as it is to truth, belief, and reason (or evidence). Stanley's considerations give us a second reason to think that knowledge does not entail certainty.

If we accept that knowledge does not entail certainty, we should offer a different model of the relation between knowledge and certainty than the one above, inspired from Williamson. Recall that this model assumes that knowledge entails a minimal degree of certainty, but that knowledge does not entail absolute certainty. Now, if we maintain that knowledge does not entail certainty, we must make sense of the idea that an outright belief can be uncertain, and we must accept that S can be in a position to know p even if p is not epistemically certain for S.

To accommodate the idea that knowledge does not entail certainty, we might be tempted to say that although the degree of justification and the degree of confidence required by knowledge do not amount to certainty, any degree of confidence and any degree of justification stronger than that amount to certainty. However, consider again the bet on whether Dublin is the capital of Ireland. We may suppose that, with respect to the proposition that Dublin is the capital of Ireland, you have a justification stronger than the one required for knowledge, and a degree of confidence stronger than the one required for outright belief (we may suppose that you are disposed to treat the proposition that Dublin is the capital of Ireland as true in all normal (or low stakes) contexts, as believing

outright requires, but also in some abnormal (with higher stakes) contexts). Still, intuitively, you do not take the bet because you are not certain that Dublin is the capital of Ireland, and that seems to be the right thing to do because it is not epistemically certain to you that this proposition is true.

A more plausible view, proposed by Stanley (2008), is that the degree of confidence and the degree of justification which count as sufficient for (subjective and epistemic) certainty depend on the context. When you say “I am certain that the restaurant is open”, you do not say that you would be ready to take a bet on this proposition whatever the odds. Rather, “I am certain that p” is true in your mouth, it seems, if the degree to which you are confident (or the degree to which you outright believe the proposition) exceeds (or is equal to) some contextually fixed degree, arguably, a degree sufficient to be willing to take p for granted in the deliberative context you are in. That is, it is plausible to say that when S says “I am certain that p”, she expresses the proposition that she is willing to treat p as true *in the circumstances of the utterance*. This is what distinguishes certainty and outright belief, for, on this approach, outright belief merely requires a disposition to treat p as true *in normal circumstances*. Or, when you say “It is certain that the restaurant is open”, you are not committed to the claim that you satisfy the Cartesian infallibilist’s standards, or that no evidence could be added in favour of this proposition. Rather, it seems, you mean that no further evidence is required in favour of this proposition given your aims.

On this approach, we must relativise the certainty norm of assertion to the context of the assertion, as follows:¹²

CNA-R. S ought to assert that p in C only if S satisfies the standards of subjective and epistemic certainty in C.

On this approach, even if there is no distinction to be made between certainty and absolute certainty, we avoid scepticism about certainty, and hence, about warranted assertion, for we can often meet the standards of certainty that are operative in the context.

Still, we should note that invoking the idea of context-sensitivity, together with the idea that we should distinguish certainty and absolute certainty, suggests other possible options. Perhaps ‘certain’ is context-sensitive, but ‘absolutely certain’ is not context-sensitive, in which case we should say that certainty, but not absolute certainty, is the norm of assertion. On this view, ‘certain’ might be considered as a relative adjective, similar to ‘tall’. Alternatively, perhaps ‘certain’ is not context-sensitive but ‘absolutely certain’ is context-sensitive, in which case we can say that absolute certainty is the norm of assertion. On this view, ‘absolutely certain’ should be considered as an absolute adjective, similar to ‘full’. This option is particularly appealing if we adopt the certainty norm of assertion and the model of certainty inspired by Williamson, in which knowledge implies certainty but not absolute certainty. Indeed, we can thus maintain that there is a distinction

¹²See DeRose (2009, 99) for a similar relativised version of the knowledge norm.

between the knowledge account of assertion and the certainty account of assertion, but that the absolute certainty account of assertion does not lead to sceptical worries about warranted assertions. Finally, perhaps we should say that both ‘certain’ and ‘absolutely certain’ are context-sensitive, in which case it is open whether certainty or absolute certainty is the norm of assertion.

In general, the idea that we are (rationally) certain of many things, or that many things are certain for us, for example, that we have hands, that there are cars, is rather intuitive. However, given that most philosophers are sceptic about certainty, the claim that certainty can be (rationally) achieved deserves to be further articulated and defended. As said earlier, we can defend the idea that scepticism about (absolute) certainty is false by distinguishing certainty and absolute certainty and/or by invoking the context-sensitivity of the epistemic standards of (absolute) certainty. The idea according to which the degree of confidence and justification which is sufficient for (absolute) certainty is context-sensitive is made plausible if ‘certain’ is a context-sensitive word, as Stanley assumes. I will defend Stanley’s model.¹³ I will argue that we must distinguish certainty and absolute certainty and that both ‘certain’ and ‘absolutely certain’ can be context-sensitive. I will assume that certainty, and not absolute certainty, is the norm of assertion.¹⁴ However, it is important to keep in mind that there are other options available to the advocate of the certainty norm for assertion.

6.2.4 Against scepticism about certainty: the meaning of ‘certain’

In this subsection, I argue against scepticism about (absolute) certainty by distinguishing certainty and absolute certainty, and by arguing that the epistemic standards of certainty and absolute certainty are context-sensitive. I start by presenting the notions of context-sensitivity and gradability in general. Then, I consider what kind of adjective ‘certain’ is. I show that ‘certain’ does not pass all the tests for maximum-standard absolute adjectives. This suggests that it is not a maximum-standard absolute adjective, and hence, that we can distinguish certainty and absolute certainty. This is sufficient to undermine an influent argument for scepticism about certainty. I suggest that the pair ‘certain’ / ‘uncertain’ is more like the pair ‘happy’ / ‘unhappy’, the meaning of which is relative to standards partitioning the associated scale into segments. Finally, I argue that the standard for ‘absolutely certain’ can be relativised to a context, as other maximum-standards adjectives. I conclude that the main reason to think that (subjective and/or epistemic) certainty is

¹³See Stanley (2008, 54): “A person’s belief satisfies the property expressed by a subjective use of “certain” relative to a context if and only if that person holds that belief at or above the contextually salient degree of confidence; mutatis mutandis for epistemic certainty and degrees of justification. Just as many beliefs may satisfy “certain”, many beliefs may, in context, satisfy “absolutely certain”. The semantic function of “absolutely” is to raise the degree on the scale above that for “certain”. So in any context, it will be harder for a belief or a proposition to satisfy “absolutely certain” than it will be for it to satisfy “certain”. But it still is possible for a belief to satisfy “absolutely certain” in one context, while not satisfying “absolutely certain” (or even “certain”) in another.”

¹⁴The reason why it is more plausible to think that certainty, rather than absolute certainty, is the norm of assertion, will emerge below. In a nutshell, challenging assertions in term of certainty seems appropriate whereas challenging assertions in terms of absolute certainty seem too harsh.

not, or cannot be, (rationally) realised – namely that ‘certain’ is a maximum-standard absolute adjective and that it is not context-sensitive – is wrong.

Context-sensitivity and gradability

It is beyond dispute that some linguistic expressions are context-sensitive, in the sense that their reference, or denotation, depends on the context of use. For instance, suppose that Keith says “I am a philosopher” and Bill says “I am a philosopher”. It is clear that their utterances do not semantically mean the same thing, even if the sentence they use is the same. What Keith says is that Keith is a philosopher, and what Bill says is that Bill is a philosopher. The first utterance may be true while the second one is false.

Following Kaplan (1989), we must distinguish the linguistic meaning of an expression, its “character”, from its semantic content. The linguistic meaning of an expression depends on the syntactic and semantic rules of the language. It is a function from the context of utterance to truth-conditions (a content). For expressions that are not context-sensitive, since the content is fixed across the contexts of use, their character is their content. For context-sensitive expressions, their content varies across the contexts of utterance. Therefore, their linguistic meaning is different from their content. The content is the proposition that is expressed by the utterance of a sentence.¹⁵ The content of an uttered sentence can be represented by a function from possible worlds to truth-values.

To illustrate, consider the expression ‘It is raining’. It may be argued that it is context-sensitive in the sense that it does not express a complete proposition (a truth-evaluable content) unless one adds a propositional constituent (which, in this case, does not correspond to a constituent articulated at the level of the sentence) regarding the place where it is raining, and which depends on the context of utterance.¹⁶ An expression may also be context-sensitive due to the context-sensitivity of one of its articulated constituents. Consider the sentence ‘I am a philosopher’. Out of context, it is neither true nor false, because it expresses an incomplete proposition. This is due to the fact that the character of its articulated constituent ‘I’ is insufficient to provide a definite semantic content out of context. The linguistic meaning of ‘I’ is that ‘I’ refers to the speaker. Out of context, there is no determinate speaker. In the same way, sentences containing demonstratives (‘this’) are neither true nor false out of context. The content of ‘this’ is determined by the linguistic meaning of this expression which is that ‘this’ refers to the object that the demonstration demonstrates in the context of utterance.

In sum, in cases of context-sensitive expressions, the context in which the expression is used provides what is necessary to complete the linguistic meaning of the sentence in order to get the

¹⁵According to the paradigm of direct reference that Kaplan puts forwards, a propositional component can be an object. For instance, the proposition that is expressed by an utterance of ‘I am a philosopher’ is constituted by the speaker and the property of being a philosopher. By contrast, on the descriptivist paradigm that goes back to Frege, propositional components are part of the sense of the expression, which in turn determines the reference.

¹⁶See Perry (1986).

complete semantic value of the utterance. Thus, two different tokens of the same type of utterance, in different contexts, can express different propositions.

There are various kinds of context-sensitive terms: indexicals ('I'), demonstratives ('this'), relational terms ('enemy'), perspectival terms ('left'), gradable adjectives ('tall'), etc. Is 'certain' a context-sensitive term?

'Certain' is a gradable adjective. Gradable adjectives are adjectives associated with a scale. They can take the comparative form (-er/more), they can be modified by degree modifiers such as adverbs (very, extremely, totally, absolutely...), and they can be combined with 'how' to form questions. 'Certain' satisfies these conditions. We can say "John is more certain than Bill that the bank is open", "John is absolutely certain that the bank is open", and it makes sense to ask "How certain is John"? (Similar sentences can be formed with epistemic certainty).

In addition, many gradable adjectives, like 'tall', 'big', etc., are context-sensitive. Arguably, whether "John is tall" is true, for example, depends on a standard of comparison, which is determined by a comparison class provided by the conversational context. If we are talking about children, the standard may be not too high on the scale of height, so that John may meet this standard. In contrast, if we are talking about basketball players, the standard is presumably quite high on the scale of height, so that John might not meet this standard.¹⁷ Likewise, one might think, as 'certain' is associated with a scale, it is also context-sensitive.

However, being associated with a scale does not imply context-sensitivity.¹⁸ As Stanley (2005b, 77) notes:

an expression can be semantically linked with a scale, even though there is a designated point on that scale that, for any context *c*, is the point of demarcation between F and not-F. This is what we see with expressions such as 'taller than 6 feet'. Something can be much taller than 6 feet, even though it is a context-independent matter whether or not something is taller than 6 feet.

Similarly, it might be that 'certain' refers to a certain point (or a set of points) on a scale that does not vary in a context-sensitive way. Some have argued, in particular, that when the scale associated with a gradable adjective is closed at one endpoint, in general, this adjective denotes this endpoint, and hence, this adjective is not context-sensitive. Indeed, this endpoint provides a potential fixed value for the standard of application of the adjective. If, in general, we prefer interpretations that minimize context-dependence, then, in general, adjectives associated with close scales will take an endpoint of the scale as their standard of application.¹⁹

Adjectives that denote an endpoint of the associated scale are called *absolute* gradable adjectives,

¹⁷See Kennedy and McNally (2005).

¹⁸See Kennedy and McNally (2005, 355-356).

¹⁹See Kennedy and McNally (2005, 361).

and those that do not are called *relative* gradable adjectives.²⁰ In the case of relative adjectives, the standard is relative to the context in that a contextually fixed comparison class determines the degree on the scale that counts as the standard. In the case of absolute adjectives, there is a clear sense in which the standard denoted on the scale is not relative to the context: it is always the endpoint on the scale. Thus, there are two kinds of absolute gradable adjectives, depending on whether they refer to the minimum limit or the maximum limit on the scale. The first ones may be called, following Kennedy and McNally (2005), ‘minimum-standard absolute gradable adjectives’, and the second ones ‘maximum-standard absolute gradable adjectives’.

However, we must take note of the fact that what counts as the endpoint of a scale may be context-sensitive. If so, absolute adjective can be context-sensitive too.

Thus, to see what kind of gradable adjective ‘certain’ is we must determine three things. First, we must determine the associated scale. Second we must determine the point (or the set of points) that it refers to, that is, whether it is an absolute or a relative gradable adjective, or something else, and if it is an absolute gradable adjective, whether it is a minimum-standard absolute adjective or a maximum-standard absolute adjective. Third, we must determine whether it is a context-sensitive word. Concerning the two first questions, I will mainly rely on Kennedy and McNally (2005)’s general discussion of gradable adjectives.

What kind of gradable adjective is ‘certain’?

What is the kind of scale associated with ‘certain’? To begin with, it is necessary to distinguish at least four kinds of scales. Scales that have no lower nor upper limit may be called ‘open scales’. Scales that have upper and lower limits may be called ‘totally closed scales’. Scales that have a lower limit but no upper limit may be called ‘lower closed scales’. Scales that have an upper limit but no lower limit may be call ‘upper closed scales’.

To see whether a scale associated with an adjective is totally open, closed at one point, or totally closed, we may use maximality modifiers (e.g. ‘totally’, ‘completely’, etc.). If an adjective does not accept maximality modifiers, then the associated scale is open at one point. If the antonym of this adjective does not accept maximality modifiers either, then the associated scale is totally open.

Let me give examples. Consider ‘tall’. “He is absolutely (fully, completely) tall” is infelicitous. Therefore, the associated scale is open at one side. Consider ‘small’, the antonym of ‘tall’. “He is absolutely (fully, completely) small” is infelicitous. Therefore, the associated scale is open at the other side. We can conclude that the associated scale for ‘tall’ is totally open.

Consider ‘flat’. “He is absolutely flat” is felicitous. Therefore, the associated scale is closed on one side. Consider ‘bumpy’, the antonym of ‘flat’. “He is absolutely bumpy” is infelicitous. Then, the associated scale is open on the other side. The scale is upper closed. (Whether or not a scale

²⁰See Kennedy and McNally (2005)

is lower closed or upper closed depends on whether the adjective referring to the endpoint refers to a minimum-standard or a maximum-standard. I explain below how minimum-standard adjectives and maximum-standard adjectives can be distinguished).

Consider ‘famous’. “He is absolutely famous” is infelicitous. Therefore, the associated scale is open on one side. Consider ‘unknown’, the antonym of ‘famous’. “He is completely unknown” is felicitous. Then the associated scale is closed at the other side. The scale is lower closed.

Consider ‘open’. “The room is totally open” is felicitous. Therefore, the scale is closed at one side. Consider ‘closed’, the antonym of ‘open’. “The room is completely closed” is felicitous. Therefore, the scale is closed at the other side. The associated scale is totally closed. (Further, note that “The room is half open” is felicitous. If an adjective accepts a proportional modifier, then the related scale is totally closed. There must be two limit points on a scale to have a proportion).

What about ‘certain’ and ‘uncertain’? “He is absolutely certain” is felicitous. “He is absolutely uncertain” is felicitous. Further, “He is half certain / uncertain” is felicitous. Therefore, the scale associated with ‘certain’ and ‘uncertain’ is totally closed.²¹ (We get the same result for epistemic certainty.)

Next, consider what kind of gradable adjective ‘certain’ is. First of all, it is crucial to realise that the distinction between absolute (maximum / minimum-standard) and relative adjective does not exhaust the set of possibilities for gradable adjectives. Consider ‘taller than 6 feet’. It is not an absolute gradable adjective because it is associated with a totally open scale (“John is absolutely taller / smaller than 6 feet” is infelicitous). Adjectives associated with a totally open scale cannot be absolute adjectives, because there is no upper or lower limit of the associated scale to which they can refer. Still, it is not a relative adjective, because it is not context-sensitive.

As we have seen, ‘certain’ is associated with a totally closed scale. However, in the same way that a gradable adjective associated with an open scale may be not relative (e.g. ‘taller than 6 feet’), it might be that an adjective associated with a totally closed scale is not absolute, but relative, or of another kind.

Is ‘certain’ a maximum-standard absolute adjective, a minimum-standard absolute adjective, a relative adjective, or an adjective of another kind? Often, philosophers consider ‘certain’ as a maximum-standard absolute adjective. For example, Unger (1975) argues that it is similar to ‘empty’ or ‘flat’, and on this basis, argues that (almost) nothing is really certain, as (almost) nothing is completely flat or empty. After all, in most cases, there are still small bumps or grains of dust that prevent the thing to be *completely* flat or empty. We often use these words loosely, which does not matter to the extent that for practical purposes, having a false belief that is close enough to

²¹Surprisingly, Kennedy and McNally (2005, 355) take the scale associated with ‘certain’ and ‘uncertain’ to be closed only at the upper limit. According to them, “I am fully uncertain” is infelicitous. However, it is clear that ‘certain’ and ‘uncertain’ may be modified by proportional modifiers such as ‘half’ or ‘mostly’, which indicates that the associated scale is totally closed. Further, to my hear, “I am completely / fully uncertain” is not infelicitous, although perhaps a more common way of saying it is “I am completely unsure”.

the truth may be better than having no belief at all. According to Unger, as ‘certain’ is like ‘flat’ or ‘empty’, a maximum-standard adjective, (almost) nothing is certain, because almost nothing is completely certain. In addition, as Unger thinks that knowing entails being certain, according to him, one knows (almost) nothing. However, is ‘certain’ like ‘empty’ or ‘flat’, a maximum-standard adjective?

In order to distinguish maximum-standard absolute adjectives, minimum-standard absolute adjectives, and relative adjectives, different tests have been proposed. At least nine tests can be found in the literature. I consider them in the following table, with the results for ‘certain’, ‘happy’ and ‘taller than 6 feet’. ‘F’ refers to the adjective (e.g. ‘tall’), ‘A’ to its antonym (e.g. ‘small’), ‘X’ means that the adjective fails the test, ‘V’ that it passes it, ‘?’ that this is unclear, ‘*’ that, although this is the intuitive result, it has been disputed in the literature that it is the correct result, and ‘-’ that the test does not seem to apply. It emerges that ‘certain’, ‘happy’ and ‘taller than 6 feet’ seem to form separate categories, irreducible to the categories formed by absolute and relative adjectives.

Tests	Max	Min	Re	certain	happy	taller than 6 feet
a. If x is F, nothing is more F than x	V	X	X	X*	X	-
b. If x is not F, x has no degree of F	X	V	X	V	V	-
c. If x is F-er then y, x is F	X	V	X	?	?	-
d. x is less F than y, then x is not F	V	X	X	?	?	-
e. If x is partially F, x is not-F	V	X	X	?	?	-
f. If x is partially F, x is F	X	V	X	?	?	-
g. x is rather F, x is F	X	V	X	V	V	V
h. x is slightly F, x is F	X	V	X	V	V	V
i. If x is almost F, x is slightly A	V	X	V	V	V	V

While, presumably, intuitions concerning ‘certain’ may diverge, at least one may agree that in light of all these tests, it is far from clear that ‘certain’ is a maximum-standard adjective, or a minimum-standard adjective, or a relative adjective. Rather, ‘certain’ behaves more like ‘happy’, and not like ‘close’, ‘dirty’ or ‘tall’. To illustrate some of the differences, here are some examples.

Consider ‘certain’ *versus* ‘flat’:

- (1) John is certain that he has hands, but he is even more certain that he exists
- (2) #The road is flat, but the table is even more flat

Consider ‘certain’ *versus* ‘dirty’:

- (3) ?John is more certain than Bill that the restaurant is open, therefore John is certain that the restaurant is open
- (4) Room A is more dirty than room B, therefore room A is dirty

Consider ‘certain’ *versus* ‘tall’:

- (5) #John is rather certain that the restaurant is open, but he is not certain that the restaurant is open
- (6) John is rather tall, but he is not tall

In light of these intuitive differences, we should at least acknowledge that it is far from clear what kind of gradable adjective ‘certain’ is. The argument for scepticism about certainty starting from the idea that ‘certain’ is a maximum-standard absolute gradable adjective is rather fragile.

Many results with ‘certain’ or ‘happy’ are unclear. The intuitions related to tests c, d, e, and f are hazy (mine at least). How can we explain this?

A natural explanation has to do with the scale structure associated with these adjectives. We saw that the scale associated with ‘certain’ is closed. However, there are different kinds of closed scales. In particular, a (closed) scale can be partitioned in different segments. Suppose that the pair ‘certain’ / ‘uncertain’, like the pairs ‘happy’ / ‘unhappy’ and ‘taller than 6 feet’ / ‘smaller than 6 feet’, refers to a scale that is partitioned into segments. The first segment is constituted of negative degrees of certainty and the second segment is constituted of positive degrees of certainty. My suggestion is that depending on which segment of the scale we associate with the expressions, tests c, d, e, and f give different results, which explains why intuitions are hazy.

For example, consider ‘happy’ / ‘unhappy’. If we associate to the expression ‘John is happier than Jill’ the segment of the scale which is constituted of the degrees of happiness, then it follows that if John is happier than Jill, then John is happy. But if we associate the segment constituted of the degrees of unhappiness, then it does not follow that John is happy. In this case, John is just less unhappy, which does not imply that John is happy. Arguably, which segment of the scale is associated with the meaning of the comparative expression depends on the context.

In favour of the claim that the scale associated with ‘happy’ / ‘unhappy’ is a segmented scale (more precisely, as ‘happy and ‘unhappy’ are contraries but not contradictories, the associated scale is partitioned into three segments, for you can be neither happy nor unhappy), we should note that ‘happy’ and its antonym ‘unhappy’ do not express inverse ordering relations, in contrast to antonyms like ‘open’ / ‘close’ or ‘empty’ / ‘full’. Indeed, if two antonyms F and A express inverse ordering relations, then x is F-er than y if and only if y is A-er than x.²² However, it is false that John is happier than Mary if and only if Mary is more unhappy than John. John may be happier than Mary even if Mary and John are not unhappy. If Mary and John are not unhappy, then Mary is not more unhappy than John. Therefore, degrees of happiness are different from degrees of unhappiness. These are not degrees of the same property with an inverse ordering.

We can express the same idea as follows. Scales can differ with respect to their degree, their

²²See Kennedy and McNally (2005, 351).

dimension, and their ordering.²³ Arguably, the scales associated with ‘happy’ and ‘unhappy’ do not differ in the way in which their domain of the objects is ordered. These antonyms do not use the same scale while just imposing inverse ordering. But these scales do not differ by their dimension either. Indeed, the following comparison is not anomalous: “She is as happy that John comes as she is unhappy that Jill comes”. Presumably, they differ with respect to their degrees. The scale associated with ‘happy’ / ‘unhappy’ is partitioned into segments distinguishing degrees of happiness and degrees of unhappiness.

The idea that ‘happy’ and ‘unhappy’ refer to standards partitioning a scale into segments with different degrees, those of happiness and those of unhappiness, is reinforced by the fact that these adjectives behave in some respect like minimum-standard adjectives. Thus, if S is somewhat unhappy that p, then S is unhappy that p. If S is somewhat happy that p, then S is happy that p. This is explained if they respectively refer to the minimal degree of happiness and the minimal degree of unhappiness on a partitioned scale with two maximum limits, absolute happiness and absolute unhappiness, and two minimum limits, minimum happiness and minimum unhappiness. The maximum of happiness (or unhappiness) is not the minimum of unhappiness (or happiness). Rather, it seems, the minimum of unhappiness is closer to the minimum of happiness.

Moreover, the pair ‘happy’ / ‘unhappy’ is similar to the pair ‘more than 6 feet’ / ‘less than 6 feet’, which is clearly associated with a scale of height partitioned into two segments by the 6 feet threshold. In particular, ‘happy’ and ‘unhappy’, like ‘more than 6 feet’ and ‘less than 6 feet’, do not seem to refer to a single point on the associated scale, but to a set of points or degrees.

The same seems to hold for ‘certain’ and ‘uncertain’. It seems false that John is more certain that he exists than that he has hands if and only if John is more uncertain that he has hands than that he exists. John may be certain of both. So, it seems, these antonyms do not refer the same degrees on the scale. In the same way that degrees of unhappiness are not degrees of happiness, degrees of uncertainty are not degrees of certainty. The maximum of uncertainty (or certainty) is not the minimum of certainty (or uncertainty). Rather, it seems that the minimum of uncertainty (or certainty) is closer to the minimum of certainty or (uncertainty).²⁴ In addition, ‘certain’ (/ ‘uncertain’) does not seem to refer to a single point on the associated scale, but to a set of points (or degrees).

How can we explain on this basis the fact that intuitions are hazy with c, d, e and f? Suppose that ‘certain’ and ‘uncertain’, or ‘happy’ and ‘unhappy’, respectively refer to positive and negative

²³See Kennedy and McNally (2005, 351).

²⁴The idea that there are degrees of certainty in this sense can be found, for example, in Ayer (1973, 232): “When philosophers talk about certainty, they often treat it as a matter of degree. Thus, a philosopher who does not deny the truth of such a proposition as that there are trees in his garden (...) may still wish to say that it is less certain than some proposition which records his current visual or tactual sensations.” (Quoted by Klein (1981, 133)). It seems that a philosopher who does not deny that there are trees in his garden, but still wish to say that it is less certain than the proposition that it seems to him that there are trees in his garden, takes as certain that there are trees in his garden, although not as certain as the fact that it seems to him that there are trees in his garden.

degrees on a segmented scale. The fact that intuitions are hazy with ‘certain’ and ‘happy’ may be explained by the fact that we can associate the scale with these adjectives in two different ways. We can consider the positive part of the scale or the negative part of the scale. To see this, consider:

- (9) If John is more happy than Mary, then John is happy.
- (10) If John is partially happy that it is raining, then John is not happy that it is raining
- (11) If John is partially happy that it is raining, then John is happy that it is raining.

(9) seems true if we are considering degrees of happiness but false if we are considering degrees of unhappiness. (10) seems true if we are considering degrees of unhappiness, but false if we are considering degrees of happiness. (11) seems true if we are considering degrees of happiness but false if we are considering degrees of unhappiness.

Likewise, the intuitive results for these tests c, d, e, and f, regarding ‘certain’, depend on whether we are considering the negative or the positive part of the associated scale:

- (12) If John is more certain than Bill that he has hands, then John is certain that he has hands
- (13) If John is partially certain that he has hands, then he is unsure that he has hands
- (14) If John is partially certain that he has hands, then he is certain that he has hands

(12) seems false if we are considering degrees of uncertainty, because in such a case, ‘more certain’ means ‘closer to certainty’ or ‘less uncertain’. However, it seems true when we are considering degrees of certainty. In such a case, ‘more certain’ means ‘closer to absolute certainty’ and not ‘less uncertain’. (13) seems false if we are considering degrees of certainty but true if we are considering degrees of uncertainty. (14) seems true if we are considering degrees of certainty and false if we are considering degrees of uncertainty.

This general picture of the scale associated with ‘certain’ and ‘uncertain’, and the related intuitions that I take for granted, is likely to be controversial. However, the most important point here is that it is far from clear that ‘certain’ (in its subjective and epistemic sense) is a maximum-standard gradable adjective that can be compared with ‘empty’ or ‘flat’. Unger’s argument for scepticism about certainty based on this comparison is fragile. There is no strong reason to endorse scepticism about certainty on this basis.

In addition, a second defence of non-scepticism about certainty can be proposed. For the sake of argument, let us grant that ‘certain’ is a maximum-standard adjective. We may question that it follows from this that ‘certain’ is not context-sensitive. If ‘certain’ is a context-sensitive word, then the fact that “y is more certain than x” is true when uttered in context A may show, if ‘certain’ is a maximum-standard adjective, that “x is certain” is false in this context. But this cannot show that “x is certain” is false when uttered in context B. Therefore, that “y is more certain than x” is true

in some contexts does nothing to show that we should be sceptic about certainty concerning x , in the sense that “ S is certain that x ” would be (almost) always false.

In favour of the idea that (maximum-standard) absolute terms can be context-sensitive, it may be argued that what counts as the limit of a given scale can be context-sensitive.²⁵ The standards for ‘clean’ seem to vary depending on the thing to which this adjective is predicated, for example, whether it is a kitchen knife or a surgical instrument. Similarly, whether a glass is full seems to depend on the liquid. A glass is full of beer when entirely filled up, whereas it is full of wine when filled up to about half of its capacity.

We may think that ‘certain’ and ‘uncertain’ are context-sensitive in roughly the same way. What it takes to be certain (or uncertain) (or for p to be certain / uncertain) can vary with the context. For example, the standards of epistemic certainty do not seem to be the same in a trial when the issue is to sentence someone to death, and when it comes to decide whether to eat out and the stakes are low. To show that it is plausible to think that ‘certain’ is context-sensitive, let me consider two tests for context-sensitivity.

First, context-sensitive words can be used with two different meanings in the same conversation, leading to a misunderstanding. This misunderstanding can be dispelled by distinguishing the two different meanings at stake. To illustrate, consider the following dialogue:²⁶

John - I know that he will come at 10 a.m.

Paul - You can tell that Robert won't be late due to traffic jam?

John - No.

Paul - So you do not know that he will come at 10 a.m, right?

John - I did not mean that. I was not talking about Robert, but about Jack.

John's last answer may be perfectly correct, and this is due to the fact that ‘he’ is context-sensitive. Now, consider a similar dialogue:

John - I know that Robert will come at 10 a.m.

Paul - You can eliminate the possibility that he will be late due to traffic jam?

John - No.

Paul - So you do not know that Robert will come at 10 a.m, right?

John. - I did not mean that. I was not considering the traffic jam possibility.

Clearly, John's last utterance is puzzling. Intuitively, he should rather say something like “You're right, I should not have said this”. This fact may be taken to speak against the idea that ‘know’ is

²⁵See Cruse (1980); McNally (2011); Burnett (2014).

²⁶See also Stanley (2005b, 52-53).

context-sensitive. At least, if we maintain that ‘know’ is context-sensitive, we should also maintain that we are blind to this semantic fact, which is potentially problematic.²⁷

Consider a similar dialogue with ‘certain’:

John - I am certain that Robert will come at 10 a.m.

Paul - Would you bet on this at any odds?

John - No

Paul - So you are not certain that Robert will come at 10 a.m., right?

John - I did not mean that. I was not considering such a level of certainty (/ I didn’t mean I was *that* certain).

John’s last answer is not puzzling. This speaks in favour of the idea that ‘certain’ is context-sensitive.

Second, if a term is context-sensitive, it can block inter-contextual disquotational indirect reports.²⁸ Suppose for instance that John says “You should wear that”, in a context where Justine is the contextually salient audience and a blue hat is being demonstrated. Now, suppose that Bob indirectly report what John said by saying “John said that you should wear that”, in a context where Ludlow is the contextually salient audience and a pair of sandals are being demonstrated. Bob did not report correctly what John said. This is explained by the fact that ‘you’ and ‘that’ are context-sensitive and do not mean the same thing in different contexts.

Do expressions with the word ‘certain’ block inter-contextual disquotational indirect report? Consider the following scenario, an adaptation of DeRose’s scenario involving Lena, Thelma and Louise:²⁹

The Office. Today is Thelma’s day off, but she decides to walk up to the office in order to finding out whether a certain colleague, John, is at work. As she passes the door to John’s personal office, she sees his hat hanging on the hook in hallway, which, in her long experience, has been a sure-fire sign that John is in fact at work. She also hears one working colleague shout to another, “Why don’t you clear that letter with John quick before you send it off?” Satisfied that John is at work she goes out to a local tavern to meet friends.

At the Tavern. At the tavern, Thelma meets a friend who bet that John wouldn’t be at work, and so owes Thelma \$2. Thelma says, “Hey, John was at work today. Pay up, sucker!” When her friend asks, “Are you sure?” Thelma replies, “I am certain that John was at work. His hat was hanging in the hall outside his office door, and I heard Frank

²⁷See Schiffer (1996).

²⁸See Cappelen and Lepore (2005, 88-90).

²⁹See DeRose (2009, 4-5.)

telling someone to quickly check something with John before sending it off.” Satisfied with Thelma’s evidence, the friend pays up. Louise overhears the conversation.

Louise with the Police. On her way home, Louise is stopped by the police. They are conducting an extremely important investigation of some horrible crime, and, in connection with that, are seeking to determine whether John was at work that day. It emerges that they have some reason to think that John was at work and no reason for doubting that, but as the matter has become so important to the case, they are seeking to verify that he was there. When the police ask her whether she could testify that John was at work, Louise replies, “Well, no, I am not certain that John was at work.” When the police ask Louise whether someone else could testify, Louise answers, “Yes, Thelma. She said she was certain that John was at work.”

It does not seem that Louise correctly reports what Thelma said. Thelma did not say that she was certain enough to testify to the police, but only that she was certain enough to make her friend pay. Note also that we would not find it strange if the police asked in what context Louise said she was certain, which is what we can expect if ‘certain’ is context-sensitive.

The fact that ‘certain’ passes these two tests for context-sensitivity is good evidence that it is a context-sensitive word.³⁰

To summarize, the certainty account of assertion does not imply that most of our assertions are unwarranted. First, it is unclear that ‘certain’ is a maximum-standard absolute gradable adjective. Second, even if ‘certain’ were a maximum-standard absolute adjective, it is unclear that this would lead to scepticism about certainty, because maximum-standard adjectives can be context-sensitive, and, arguably, ‘certain’ is context-sensitive.

The non-sceptical picture of ‘certain’ that I propose is thus the following. The pair of gradable adjectives ‘certain’ / ‘uncertain’ is, like the pairs ‘happy’ / ‘unhappy’ and ‘more than 6 feet’ / ‘less than 6 feet’, associated with a partitioned scale. These antonyms respectively refer to positive and negative sets of points (or degrees) on the associated scale. Like the pair ‘happy’ / ‘unhappy’, but

³⁰Other tests for context-sensitivity have been proposed. In addition to inter-contextual disquotational indirect report tests, Cappelen and Lepore (2005, 88-108) propose collective description tests, according to which if ‘F’ is context-sensitive, then if “x is F” is true in context A and “y is F” is true in context B, one cannot thereby conclude that “x and y are F” is true, and an indirect disquotational test, according to which if ‘F’ is context-sensitive, then “There are (or can be) false utterances of ‘x is F’ even though x is F” may be true. In addition to speech act report tests, Stanley (2005b) considers propositional anaphora tests, according to which if ‘F’ is context sensitive, it can make sense to say things like “x is F, but now that I am in a different context, x is not F, although what I said earlier is true”. Cappelen and Hawthorne (2009) propose agreement tests, according to which if ‘F’ is context-sensitive and S says “x is F”, and S* says “x is F”, we cannot thereby conclude “S and S* (both) agree that x is F”, and if S says “x is F” and S* says “x is not F”, we cannot thereby say that they disagree. Some of these tests have been criticised. Stanley (2005b, 49-50) criticises the collective description test, and Cappelen and Hawthorne (2009) criticise the inter-contextual disquotational test. However, the reliability of these tests is questioned in the sense that some uncontroversial context-sensitive expressions do not pass them, not in the sense that uncontroversial insensitive expressions pass them. Therefore, even if ‘certain’ do not pass some of these tests, this is insufficient to show that it is not a context-sensitive word, and if it passes some of them, this seems sufficient to show that it is context-sensitive.

unlike the pair ‘more than 6 feet’ / ‘less than 6 feet’, the pair ‘certain’ / ‘uncertain’ is associated with a totally closed scale. Unlike the pair ‘more than 6 feet’ / ‘less than 6 feet’, and presumably also unlike the pair ‘happy’ / ‘unhappy’, the threshold dividing the scale associated with the pair ‘certain’ / ‘uncertain’ is context-sensitive. If so, what counts as ‘certain’ or ‘uncertain’ depends on the context.

6.3 Arguments in favour of the certainty account of assertion

In the previous section, I have presented the certainty account of assertion and argued that we should not be sceptic about (epistemic and subjective) certainty. In this section, I shall consider the arguments in favour of the certainty account of assertion. I shall show that this account does better than other rival accounts – in particular the knowledge account – in explaining the relevant data.

6.3.1 Moore-paradoxical sentences

The first way of arguing in favour or against one or another epistemic norm of assertion is to consider Moore-paradoxical sentences. Moore noted that it sounds incoherent to say “p but I do not believe that p”.³¹ It is clear that this sentence does not semantically express a contradiction, for it may well be true that p and that the subject uttering this sentence does not believe the proposition. So, how can we explain the fact that asserting this sentence seems incoherent?

A plausible explanation appeals to a norm of assertion. Suppose, for example, that you should not assert what you do not believe. Suppose that you are seen as following the norm of assertion. Then, if you assert p, you are seen as believing that p. In other words, in virtue of the supposed norm of assertion, by asserting that p, you represent yourself as believing that p. In the second half of your assertion, though, in virtue of this norm, you represent yourself as believing that you do not believe that p. Thus, an assertion of this sentence sounds incoherent because the subject represents herself as believing and not believing that p.

There are at least four Moore-paradoxical assertions to consider:

- (1) p but I do not believe that p
- (2) p but I do not know that p
- (3) p but I am not certain that p
- (4) p but it is not certain that p³²

³¹See Moore (1942, 543).

³²Some philosophers might find that there are cases in which assertions of (3) and (4) do not sound incoherent or infelicitous. They might in particular point to the case that Radford (1966) proposes to argue that knowledge without belief is possible. Suppose that, in a quiz, Albert is asked the question “When did Queen Elisabeth died?”.

The certainty account easily explains why assertions of the form (1)-(4) sound incoherent in a unified way.³³ According to this account, supposing that you follow the norm of assertion, when you assert that *p*, you take yourself to be certain that *p*, and you take *p* to be certain. Certainty (subjective and epistemic) entails belief and knowledge. So, in each case, if you assert the sentence, you represent yourself as being certain that *p* and taking *p* to be certain, while at the time taking yourself not to satisfy a necessary condition for certainty. So, you represent yourself as having incoherent attitudes. In contrast, an advocate of a weaker norm will have to propose a different explanation for the oddity of (3) and (4).

In *Knowledge and its Limits* (ch.11), Williamson defends the idea that knowledge is the norm of assertion against weaker norms by focusing in particular on the fact that this account can explain the oddity of (2). However, when considering utterances of sentences of the form (3)-(4) he writes:

What seems to be at work here is a reluctance to allow the contextually set standards for knowledge and certainty to diverge. Many people are not very happy to say things like ‘She knew that *A*, but she could not be certain that *A*’

The suggestion seems to be that (3)-(4) do not speak in favour of the certainty account and against the knowledge account because we have a general reluctance to dissociate the standards for knowledge and those of certainty in our assertions. According to this line of thought, we should explain the oddity of (3)-(4) in the following way. By saying that *p*, you represent yourself as knowing that *p*, and in virtue of a general reluctance to dissociate the standards of knowledge from those of certainty (which can be shown by the reluctance of saying things like “*S* knows that *p* but *S* is not certain that *p*”), we expect you to be reluctant to say “I am not certain that *p*”. For this would amount to represent yourself as endorsing “I know that *p* but I am not certain that *p*”, the assertion of which we are reluctant to accept as felicitous or coherent.

The problem, however, is that this general reluctance is not explained in the knowledge account framework, but merely taken for granted. If this reluctance has to do with the fact that knowledge entails certainty, the knowledge account is not different from the certainty account, and the oddity of (3)-(4) do not challenge the knowledge norm. But if we assume that knowledge does not entail

Albert thinks he does not know the answer. Yet, he reliably answers “Elisabeth died in 1603”. Albert does not trust his answer, and hence, according to Radford, Albert does not believe that Elisabeth died in 1603. Still, he knows the answer. Whether or not Radford is right, we might think that this case illustrates a situation where a subject can utter without incoherence or infelicity “*p* but I am not certain that *p*”. If so, assertions of (3) and (4) are not always infelicitous. Two remarks are in order here. First, suppose we accept that Albert’s case is a case in which an assertion of (3) and (4) from Albert would be felicitous. Then, it must also be granted that such a case is also a case in which (1) and (2) could be asserted without infelicity by Albert, for Albert does not believe that Elisabeth died in 1603, and yet he says “Elisabeth died in 1603”; also, Albert believes that he does not know that Elisabeth died in 1603, and yet he says “Elisabeth died in 1603”. This would undermine the general strategy of appealing to supposed Moorean paradoxical sentences. But, second, note that it is far from clear that if Albert utters “Elisabeth died in 1603 but I am not sure”, he is really asserting the first half of this sentence. Albert is in a quiz. As Radford stresses, Albert takes his answer to be a mere guess. So, it is plausible to think that Albert’s utterance of “Elisabeth died in 1603” is not a case of assertion, but a case of guessing.

³³See Stanley (2008).

certainty, this reluctance should be explained. I will discuss below an attempt to explain this oddity, which has been proposed by an advocate of the knowledge norm, but I will show that it fails. For the moment, it suffices to note that the certainty account can easily explain the reluctance of saying “I know that p but I am not certain that p ” pragmatically (the third-person version “ S knows that p but S is not certain that p ” does not seem as bad as the first-person version).³⁴ So, appealing to this reluctance does not show that the knowledge account is in a better position than the certainty account.

Further, to defend the knowledge account against the certainty account, Williamson claims that when we distinguish knowledge and certainty “assertability goes with knowledge, not with the highest possible standards of certainty” (Williamson, 2000, 224). Suppose that this is true. Still, the fact that you can say without apparent incoherence “ p but by Descartes’s standards, I cannot be absolutely certain that p ” does not show that certainty is not the norm of assertion, but only that Cartesian certainty is not the norm of assertion.³⁵ Therefore, that does not show that knowledge, rather than certainty, is the norm of assertion.

Finally, we should note that it is far from clear that assertability always goes with knowledge and not with the highest possible standards of certainty. Suppose we are in a philosophy class. We are considering the brain in a vat hypothesis. Suddenly, a student assert “But I know that I am not a brain in a vat. I have hands”. The assertion “I have hands” seems inappropriate given the context. Intuitively, it is not warranted enough to refute the sceptic. The context requires that to be warranted in asserting “I have hands” it should be epistemically certain for the student that she has hands in the sense that she should be in a position to eliminate the possibility that she is a brain in a vat. By raising the sceptical possibility of error in a context where the aim is to decide whether or not scepticism can be refuted, the standards for certainty are raised up to absolute Cartesian certainty. It seems that, in such contexts, assertability goes with the standards of Cartesian certainty rather than with knowledge.

6.3.2 Challenges

A second way of arguing for or against one or another epistemic norm of assertion is to put forth data concerning the conversational patterns we use to challenge assertions. Thus, in favour of the knowledge account, Williamson (2000) put forth that it is natural to challenge an assertion that p by asking “How do you know that p ?” or “Do you know that p ?”.³⁶ If knowledge is not the norm of assertion, it is unclear why such challenges are always appropriate.

Although the fact that we use ‘know’ to challenge assertions may speak in favour of the knowledge account, because it suggests that knowledge is necessary, it does not speak against the certainty

³⁴See Stanley (2008).

³⁵See also Turri (2010, 459).

³⁶See also Unger (1975, 263-264).

account. For according to this account, knowledge is always necessary.

However, one might think that the following argument can be made against the certainty account. If knowledge is not sufficient, we should not be satisfied with the mere fact that the assertor knows what she asserts. Hence, how shall we explain that a positive answer to a challenge in terms of knowledge is satisfactory, and that a challenge in terms of knowledge seems always sufficient?

An advocate of the certainty account may reply that by saying that she knows that *p*, the assertor represents herself as being certain that she knows that *p*, and hence, by factivity of knowledge, as certain that *p*. Thus, although knowledge is not the norm of assertion, we can explain why it may be sufficient to challenge an assertion by asking whether the assertor knows that *p*, and why a positive answer to this question from the assertor is always sufficient to settle the issue whether her epistemic position is good enough. Therefore, the fact that assertions are most of the time challenged by using ‘know’, and that a positive answer to this challenge is always sufficient to settle the issue whether the assertor is in a good enough epistemic position to assert, does not count against the certainty account.

Moreover, in favour of the certainty account, and against the knowledge account, we may emphasize that “we often demand certainty, and not merely knowledge, of our interlocutors” (Stanley, 2008, 51). If knowledge does not imply certainty, and knowledge is the norm of assertion, this is a surprising fact.³⁷

On this issue, Williamson (2000) writes:

[I]t would often be inappropriate to respond to the assertion ‘A’ by asking ‘How can you be so certain that A?’. The word ‘so’ flags the invocation of unusually high standards of certainty. By ordinary standards you may have had warrant to assert that A even if you could not be so certain that A.

However, the certainty account of assertion is not committed to the claim that assertion requires unusually high standards of certainty (although this may sometimes be the case). So, this does not speak to the problem.³⁸

Turri (2010) proposes an explanation of these challenges in term of certainty that is supposed to be compatible with the knowledge account. The first point made by Turri is that the (alleged) fact that knowledge is the epistemic norm of assertion is not inconsistent with the fact that we

³⁷Against the knowledge account, Kvanvig (2009) writes that the line of argument using appropriate challenges “may prove too much. It is also appropriate to ask, ‘Are you certain?’ [and] ‘Are you absolutely sure?’. If the conversational propriety of various questions is an argument in favour of the knowledge account, the propriety of these questions is an argument in favour of a stronger account: that one must be absolutely certain in order for a claim to be assertible.” However, note that it is unclear that challenges in terms of absolute certainty are as appropriate as challenges in terms of certainty. So, I take it that this line of argument may only prove that certainty is required. Gerken (forthcoming) also takes the fact that challenges in terms of certainty are appropriate to undermine the argumentative strategy focusing on appropriate challenges. Indeed, he takes it as implausible that certainty is the norm of assertion. However, once it is granted that we should not be sceptic about certainty, I see no reason why certainty could not be the norm of assertion.

³⁸See Turri (2010, 459).

appropriately challenge assertions in terms of certainty. According to Turri, the challenge may concern not whether the assertor satisfies the epistemic norm for assertion, but whether she is certain that she satisfies this epistemic norm. Turri proposes the following analogy. Suppose that an official has the authority to wed people. This is independent from whether or not she is certain that she has this authority. Having a warrant for doing something is one thing. Another thing is having certainty that one has this warrant to do that thing. Thus, it is not impossible that challenges using ‘certain’ do not concern the warrant for assertion, and if so, the fact that we challenge assertions by using ‘certain’ do not speak against the knowledge norm of assertion.

However, there is a disanalogy here. The knowledge account of assertion is supposed to capture the epistemic condition required to warrant an assertion from an epistemic point of view. In contrast, the account of a warranted wedding in terms of the authority of the official to wed people is not supposed to capture the epistemic condition required to warrant the wedding from an epistemic point of view. So, it may make sense to ask an official who is warranted to wed people – in the sense of the non-epistemic authorization to do it – whether she also has the epistemic warrant for doing it. This may require that she knows, or is certain, that she has this authority. More generally, we might think that the epistemic condition that an official has to satisfy for her action of wedding people to be epistemically warranted is not specific to weddings. It is the epistemic condition governing rational action in general. On the contrary, consider the case of assertion. If the assertion is epistemically warranted, it is unclear why the assertor should satisfy another epistemic condition for her assertion to be epistemically rational. Requiring the assertor to satisfy a further second-level epistemic condition would be analogous to requiring the official to satisfy a further second-level (non-epistemic) authorization to act on the first-level authorization to wed people, before being warranted in wedding people. This is absurd.

Likewise, it seems that if the assertor already satisfies the epistemic condition for epistemically warranted assertions, what we would be asking for, in asking her to be in a good enough epistemic position with respect to this first-level epistemic warrant, would be for her to have a further second-level epistemic warrant warranting her to assert on the first-level epistemic warrant. But if this second-level epistemic warrant is required, the first-level epistemic warrant does not warrant the assertion, after all.³⁹ Once the analogy is made correctly, it cannot show that appropriate challenges in terms of certainty are compatible with the knowledge account.

In addition, suppose that we appeal to a distinction between, on the one hand, whether the assertion is (epistemically) warranted, and, on the other hand, whether the subject is in a good enough epistemic position with respect to whether the assertion is (epistemically) warranted. Arguably, this strategy must rely on the claim that, in some cases, but not all, due to specific circumstances, for example, that it is important to make a warranted assertion, the subject should be certain that she

³⁹See chapter 9 section 3.1 for a related discussion concerning the epistemic norm of action and a more developed defence of the claim that it is spurious to distinguish a primary epistemic norm of action and a derivative secondary epistemic norm.

is making a warranted assertion.⁴⁰ But if, so, in cases where it is not important to make a warranted assertion, the challenge in terms of certainty should strike us as inappropriate. Still, acknowledging that one does not meet the standards of certainty always sounds like a retraction of the assertion, whatever the context. Suppose, for example, that a subject respond to a challenge in terms of certainty by acknowledging that she is not certain that *p*, but still maintains that *p*, and says “I am not certain that *p* (/ it is not certain that *p*), still, *p*”. It clearly seems that there is something wrong in her assertion. A much more natural reaction to a challenge in terms of certainty, when the standards of certainty are not met, is to retract the assertion, even when the circumstances are such that it is not really important that the assertion be warranted. For example, it is not unnatural to answer: “You’re right, I am not certain that *p*, but only that *p* is probable enough. But who cares? That is not really important.” The fact that this is a natural answer is surprising if certainty is not the norm of assertion. Why are we willing to retract our assertions when certainty is not met even in contexts where it is not important that the assertion be warranted? This is at odd with the claim that challenges in terms of certainty merely challenge whether the assertor is in a good enough epistemic position regarding whether her assertion is epistemically warranted.

Finally, the strategy appealing to a distinction between the epistemic warrant for assertion and a good enough epistemic position with respect to whether the assertion is warranted is dialectically bad for the advocate of the knowledge account. For if it worked, it could also be argued that challenges in terms of knowledge do not concern the epistemic warrant of assertion, but whether the subject is in a good epistemic position with respect to whether the assertion is epistemically warranted. A defender of the truth account, for example, might say that challenges in terms of knowledge merely express a requirement that the subject be in a good enough epistemic position with respect to whether the assertion is epistemically good, that is, on this view, true. As a result, this strategy undermines the argument from challenges proposed by the advocate of the knowledge norm in the first place.

The second point made by Turri concerns the meaning of ‘certain’. For the sake of argument, let us grant that there is no tension between the knowledge account of assertion and the fact that we appropriately challenge assertions in term of certainty, because the latter challenge does not concern the warrant for assertion, but whether the subject is in a good epistemic position with respect to this epistemic warrant. It remains to explain why this latter challenge is put in term of certainty rather than, for example, in terms of knowledge that one knows. Asking “Are you certain that *p*?” is not asking “Are you certain that you know that *p*?”, nor “Do you know that you

⁴⁰For example, in a parallel debate, Williamson (2005) argues that knowledge is the norm of action but that, when the stakes rise, we should know that we know (that we know...) the target proposition before acting. This is so because, according to Williamson, when the cost of inappropriateness is high, we should know (that we know...) that the premise is appropriate (that is, known) (see chapter 9 section 3.1 for a discussion of this approach). On this approach, a challenge in terms of something stronger than knowledge (certainty or knowledge that one knows) is misplaced when the cost of inappropriateness is weak. Such a challenge makes sense only when the cost of inappropriateness is strong. The same should go for assertion if we adopt this kind of explanation.

know that p ?”. Turri’s reply is that ‘certainty’ should be understood in terms of knowledge that one knows. However, is it plausible to think that S is certain that p if and only if S knows that she knows that p ? It does not seem so. The main way of acquiring certainty about p is to look for evidence concerning p , not to look for evidence regarding whether one knows that p . Further, you may know that you know that p without being certain that p , if knowledge does not entail certainty. As a consequence, Turri’s proposal fails to satisfactorily explain challenges in terms of certainty.

6.3.3 Unified explanation of Moorean sentences and appropriate challenges

Benton (2011) points out that a satisfying account of assertion should be able to explain in a unified way the Moorean data and the appropriateness of challenges. Indeed, appropriate challenges “can elicit a de facto Moorean paradox within a conversational context”. Consider:

A: It is snowing.

B: How do you know?

A: Oh, I don’t.

B: Huh??

A: Still, it’s snowing

Benton (2011, 686) argues that this favours the knowledge account because only this account is well enough positioned to explain this relation between Moorean data and challenges in terms of knowledge.

However, the certainty account of assertion does as well in explaining why challenges in terms of knowledge may elicit de facto Moorean paradoxes. Indeed, according to the certainty account of assertion, the same thing explains why utterances of Moore-paradoxical sentences sound paradoxical, and why we can challenge an assertion by using ‘know’, namely, that a necessary condition for asserting that you know that p is that you know that you know that p , and hence, by factivity of knowledge, that you know that p . Further, it is not difficult to imagine a conversation with ‘certain’ eliciting a de facto Moorean paradox in terms of certainty:

A: It is snowing.

B: Are you sure?

A: Oh, I am not.

B: Huh??

A: Still, it’s snowing

While the certainty account has a unified explanation for all these data, the knowledge account does not. Therefore, the certainty account is superior in this respect.

6.3.4 Parenthetical uses

Other linguistic data that are worth considering have to do with parenthetical uses. First, consider parenthetical uses expressing the assertor's mental state. When we want to express a mere belief in what we assert, we can use 'believe' in a parenthetical position:

- (1) It is, I believe, raining.
- (2) It is raining, I believe.

However, as Benton (2011) notes, it is striking that we cannot (or do not) use 'know' parenthetically in the same way. Consider:

- (3) ??It is, I know, raining
- (4) ??It is raining, I know.

This use of 'know' in parenthetical position seems redundant. Benton argues that this constitutes a further argument in favour of the knowledge account. Indeed, on this account, in asserting that you know that *p*, you already represent yourself as knowing that *p*. But if so, it is redundant to parenthetically use 'know' to express your knowledge that *p* when you already assert that *p*.

However, Benton fails to note that, as 'know', 'certain' cannot be (or is not) used in this way without redundancy. Consider:

- (5) ??It is, I am certain (/ I am sure), raining
- (6) ??It is raining, I am certain (/I am sure).

The certainty account of assertion can explain why 'certain' *and* 'know' are not parenthetically used to express the assertor's mental state without redundancy. Indeed, on this view, you satisfy the epistemic norm of assertion only if you are certain of what you assert and only if what you assert is certain for you, and thereby, only if you know what you assert. In contrast, the knowledge account of assertion can *merely* explain why 'know' is not parenthetically used in this way without redundancy. Therefore, the considerations regarding parenthetical use expressing the assertor's mental state ultimately speak in favour of the certainty account.

A possible worry is that one might think that if certainty is the norm of assertion, then, as knowledge is weaker than certainty, the parenthetical use of 'know' should have the same effect as the parenthetical use of 'believe' in hedging the assertion. However, not only the use of 'know' in parenthetical position does not hedge the assertion, but, as Blaauw (2012) notes, it can have a reinforcing effect. Consider the following situation (Blaauw, 2012, 106):

LAZY. John is having a fight with his wife Jill. Apparently, as Jill brings forward repeatedly during their heated conversation, John is very lazy; a point that Jill supports with ample evidence. At one point, exasperated, John asserts,

(7) I am very lazy, I know!

What John says sounds natural, and he could also have said:

(8) I am, I know, very lazy!

Of course, it is striking that the use of ‘believe’ in parenthetical position does not have this reinforcing effect, but, rather, typically hedges the assertion. Blaauw takes these considerations to favour the knowledge account. Indeed, on this account, if you already express your knowledge that *p* in asserting that *p*, then by saying that you know that *p*, you reinforce what is already expressed. However, if certainty, rather than knowledge, is the epistemic norm of assertion, then it is unclear why the parenthetical use of ‘know’ can have this reinforcing effects.

There is a reply. Assume the certainty norm of assertion. By asserting that you believe that *p*, you represent yourself as certain that you believe that *p*, but not as certain that *p*. Thus, you do not represent yourself as satisfying the epistemic norm for asserting that *p*, which is why the parenthetical use of ‘believe’ can hedge the assertion. In contrast, by asserting that you know that *p*, you represent yourself as certain that you know that *p*, and by factivity of knowledge, as certain that *p*. As a result, you represent yourself as satisfying the epistemic norm of assertion. This is why the parenthetical use of ‘know’ does not hedge the assertion.

In sum, the set of considerations having to do with parenthetical use are more easily explained by the certainty account than by the knowledge account.

6.3.5 Six further arguments

Six further considerations can be proposed in favour of the certainty account of assertion.

First, Benton (2014) notes that “a standard response when one feels not well-positioned to assert, in reply to a prompt like “Is it the case that *p*?”, is to answer “I don’t know.””. According to him, this speaks in favour of the knowledge account. Indeed, “the query was about *p*, not about whether one knows that *p*”, and thus, the fact that it is appropriate to answer “I do not know that *p*” has to be explained. The knowledge account explains why this answer is appropriate. On this view, the reply is appropriate because it is appropriate to decline a prompt to assert by saying that one does not meet the epistemic norm for assertion, and knowledge is the norm of assertion.

Still, this line of argument also favours the certainty account of assertion. Suppose that you unconfidently believe that *p*. A perfectly correct response to “Is *p*?”, is “I am not sure.” Further, note that the certainty account explains why “I do not know” is always an appropriate answer, because knowledge is a necessary condition for epistemic certainty. In contrast, the knowledge

account cannot explain why “I am not sure” is always an appropriate answer, as, according to this view, certainty is not required for assertion. As a consequence, considerations concerning typical ways of declining prompts to assert favour the certainty account over the knowledge account.

Second, the certainty account of assertion is in a good position to do justice to the communicative and social functions of assertions. Arguably, the epistemic norm of assertion should not be determined independently of the effects of assertions on the audience. This has led some philosophers to argue in favour of a knowledge norm of assertion regarding the hearer of the assertion. For example, García-Carpintero (2004, 156) argues that one must assert that *p* only if one’s audience comes thereby to be in a position to know that *p*. But for one’s audience to be in a position to know that *p* on the basis of an assertion, it is plausible to think that the assertor must be in a better epistemic position than knowledge, at least if testimonial transmission weakens the initial epistemic warrant (Stanley, 2008). Therefore, something stronger than knowledge may be required to warrant the assertion from the assertor’s point of view.

Third, the certainty account can accommodate the intuitive idea that the warrant for assertion varies with the context (and, in particular, with the audience), while explaining the invariant data having to do with ‘know’, and this, without endorsing contextualism about ‘know’.⁴¹

In other words, the certainty norm is the only norm which is compatible with these three independently plausible claims:

- A. The warrant necessary and sufficient for asserting *p* is context-sensitive.
- B. The linguistic and conversational data (regarding Moorean sentences, challenges, etc.) about ‘certain’, ‘know’ and ‘believe’ are invariant.
- C. The meaning of ‘know’, or ‘belief’, is not context-sensitive.

For example, if knowledge is the norm of assertion, we can account (partially) for B, but we must reject either A or C. If belief is the norm of assertion, we can account (partially) for B but we must reject A or C. If we adopt a warrant account – in which the warrant is context-sensitive, sometimes weaker and sometimes stronger than knowledge – we can accept A and C, but we have difficulties with B. For example, in cases in which the warrant is supposed to be weaker than knowledge, it should be felicitous to say “*p* but I do not know that *p*”. But saying this seems always odd.

In contrast, if certainty is the norm of assertion, we can adopt A, B, and C. Indeed, according to the certainty account of assertion, the flexibility is limited. The warrant inflexibly requires knowledge. This explains the invariant paradoxical flavour of asserting Moorean sentences involving ‘know’, the fact that it is always appropriate to prompt and challenge assertions by using ‘know’, and why using ‘know’ in parenthetical position is always either redundant or reinforcing.

⁴¹The problem of combining the flexibility of the epistemic warrant for assertion and the data concerning knowledge is raised by Benton (2014). See DeRose (2009) for an argument in favour of contextualism about ‘know’ from the claims that the warrant for assertion varies across contexts and that knowledge is the norm of assertion

Fourth, the certainty account is compatible with the idea that a warranted assertion may require the assertor to have current access to reasons in favour of what she asserts, in particular, in discursive contexts.⁴² In contrast, suppose that we can know that *p* without having current access to reasons for *p*. This may happen, for example, if we remember that there are sufficient reasons in favour of *p*, without remembering these reasons. If so, the claim that an assertor must (sometimes) be in a position to defend her assertion speaks against the knowledge account, because knowledge is not always sufficient for that.

Fifth, the certainty account of assertion explains what is wrong with a basic argument in favour of the idea that knowledge implies certainty, namely, the fact that uttering

- (1) I know that *p* but I am not certain that *p*

sounds incoherent. The certainty account of assertion provides a pragmatic explanation of the apparent incoherence. Roughly speaking, given the certainty norm of assertion, by asserting that you know that *p*, you represent yourself as being certain that you know that *p*, and by the factivity of knowledge, as being certain that *p*. But this contradicts the second half of your assertion. It is worth noting that third-person ascriptions do not share this oddity (e.g. “*S* knows that *p* but, being a cautious fellow, he is somewhat uncertain that *p*”). Therefore, the oddity of “I know that *p* but I am not certain that *p*” has to do with a pragmatic feature of the utterance, and not with a semantic feature (see Stanley (2008)). It is worth noting that the account rightly predicts that the following assertions will also sound incoherent:

- (2) *p* but I am not certain that I know that *p*.
 (3) *S* knows that *p*, but I am not certain that *p*.
 (4) *S* knows that *p*, but I do not know that *p*.

Sixth, there may be cases where someone knows that *p* without believing that she knows that *p*. As Hawthorne (2004, 50) notes, in such cases, “it is at best unclear whether that person is entitled to flat-out assert that *p*”, which casts doubt on the fact that knowledge is always sufficient for assertion. In contrast, arguably, if you are certain that *p* (and *p* is certain for you), you do not doubt that *p*, or that you know that *p*.

These six arguments added to the previous considerations regarding the unified explanation of Moore-paradoxical sentences, appropriate challenges and parenthetical uses, render very plausible that the certainty account, rather than the knowledge account, is the right account of the epistemic norm of assertion.

⁴²See Gerken (forthcoming).

6.4 Defence of the certainty account

6.4.1 Defence against Pritchard (2008)

Against the certainty norm of assertion, Pritchard (2008) raises the worry that even if ‘certain’ is context-sensitive, it remains unclear that we meet the standards of certainty in many ordinary contexts. First, he considers the following case (originally from Lackey (2007)):

JENNY. Jenny steps off the train in an unfamiliar town and asks the first person that she meets for directions. Suppose further that this person is indeed knowledgeable about the area and communicates this knowledge to Jenny, who promptly heads off to where she needs to go.

Pritchard writes:

Wouldn't we naturally describe such a case as one in which the informant's knowledge is transmitted to Jenny ? (...) Would it be acceptable for Jenny to assert what she knows? Well, I don't see why we would criticise her for making such an assertion. Notice, though, that given the way that Jenny formed her belief the level of evidential support that she possesses for her belief is bound to be very low. Moreover, if she is remotely conscientious in how she assigns her degrees of confidence to propositions that she believes, then she is bound to have a low level of confidence in the target proposition too. With both these points in mind, it is hard to see why Jenny (or anyone else for that matter) would regard her as being certain of what she believes, whether the certainty in question is of the subjective or epistemic variety.

The challenge for the certainty account would then be to explain why Jenny can appropriately assert what she knows, although she is uncertain of it. The dilemma would be the following. Either the advocate of the certainty account must say that Jenny cannot appropriately flat-out assert the relevant proposition, which, according to Pritchard, would contradict our intuitions and would amount to conceding too much to the sceptic about assertion, or the assertion is taken to be appropriate, and hence, the standards of certainty met, which is implausible. According to Pritchard: “no-one would surely describe Jenny as certain of what she knows”.

Let me start by the second horn of the dilemma. Why can't we describe Jenny as certain? She is described as promptly acting on what she has been told, which suggests that she acts unhesitatingly, and hence, that she is in some sense certain. All that one is told in favour of the fact that she is uncertain is that she has a low degree of confidence. But this does not necessarily imply uncertainty. Consider, for example, the degree of certainty that I have for the proposition that I exist, and compare it with the degree of certainty that I have for the fact that, say, yesterday it was raining. My degree of confidence in the latter proposition is very low as compared with my

degree of confidence in the former proposition. Still, I am in no way uncertain that it was raining yesterday.

Consider the first horn of the dilemma. Suppose we say that Jenny is uncertain, and thus, conclude that she inappropriately asserts the proposition in question. First, contrary to what Pritchard claims, the idea that her flat-out assertion would be inappropriate is rather intuitive. When we flesh out in more detail the supposition that Jenny is uncertain, it becomes more difficult to grant that she can appropriately make a flat-out assertion. Suppose, for example, that we say that she is somewhat uncertain because she feels doubtful about the truth of the proposition, or because she hesitates to act on it. Now, imagine that someone asks her for directions, and unhesitatingly, she flat-out asserts the proposition in question. It seems clear that we would regard her flat-out assertion as inappropriate. If Jenny has an uncertain belief, we would expect her to hedge the assertion, by saying something like: “This is the right direction, I believe”. Alternatively, it would be very natural for her to say “I was told that this is the right direction”. By making clear how Jenny is uncertain of what she believes, we also make it clear that a flat-out assertion from her would be inappropriate. Thus, it is no contrary to intuition to say that Jenny should not flat-out assert what she is uncertain of. (Of course, we may have the intuition that Jenny should not be blamed for her flat-out assertion. Perhaps it was not really important to make a warranted assertion, but this is not to say that her assertion was warranted.)

If, as Pritchard seems to think, it is most of the time appropriate to flat-out assert p even when one merely has an unconfident belief as to whether p , why do we have a linguistic tool to mark, on specific occasions, that the assertion is made on the basis of an uncertain belief? The fact that the assertion in question is made on the basis of an uncertain belief should go without saying. After all, according to Pritchard’s hypothesis, in most cases we do not meet the standards for certainty. Further, if in most cases assertions do not express more than uncertain beliefs, why does the parenthetical use of ‘I believe’ hedge an assertion rather than reinforce it?

Second, it is unclear why acknowledging that Jenny cannot appropriately flat-out assert the proposition that she unconfidently believes would lend toward a problematic scepticism about assertion. Most of the propositions we believe, even on the basis of testimony, are more strongly supported than that of Jenny, and are confidently believed.

The second way in which Pritchard argues that, in most context, the threshold for certainty is not met, is by arguing that “it is in fact very easy to get people to concede that they are not certain of something that they believe, even when no additional practical considerations are being raised”. This putative fact is problematic, indeed, if we want to maintain that we are certain of the things we appropriately assert, and that most of our assertions are in fact appropriate. The most plausible explanation why people easily concede that they are not certain of something they believe is precisely that they are not certain of it.

In favour of the claim that we are “rarely willing to explicitly grant that we are certain once the

issue is pressed”, Pritchard proposes the following case:

CAR PARK. [M]y wife and I are heading out of the shopping arcade and I stride purposively towards the part of the car park where I believe our car to be. Nothing in my behaviour indicates any doubt on my part on this score and, indeed, I do know that my car is parked at the relevant location. Nevertheless, were my wife to ask me whether I was certain that it was parked there—whether there was any possibility that I could be wrong—I would be unlikely to say ‘yes’. Instead, I would say something like ‘well, that’s where I seem to recall parking it and I’m pretty good at remembering this sort of thing’, or something along those lines.

I must say that I do not share Pritchard’s intuition here. It seems to me really implausible that an ordinary subject would reply this sort of things. This may be because the case is underdescribed. In particular, it is unclear with which degree of vivacity the subject is supposed to remember where he parked the car, or to what extent he paid attention when he parked it. Undoubtedly, there are cases where we are unsure, for example, when we do not remember very well. But in normal cases, the most natural reply is that we have no doubt that the car has been parked in the location in question.

Finally, note that the case proposed by Pritchard does not concern an assertion, but an action. It might be that the epistemic norm of action is weaker than the epistemic norm of assertion. Consider the same case, but, instead of merely acting on the target proposition, suppose that the subject also asserts “My car is parked there.” Then, suppose that the subject’s wife asks whether he is certain. Would the subject be likely to say “No, I am not”? It does not seem to me, at least if the subject remembers where he parked her car and makes the assertion on this basis. On this point, I can just encourage the reader to check that folk people are willing to insist that they are certain of where they have parked their car, what they have eaten at lunch, etc.

6.4.2 Defence against Turri (2010)

Against the certainty account of assertion, and in favour of the knowledge account, Turri (2010, 461) writes:

Certainty features in appropriate challenges but *not* in appropriate prompts. And certainty does not figure in as full a range of appropriate challenges as knowledge does. This demonstrates that knowledge is more intimately related to assertion than is certainty.

These considerations are far from decisive. Consider how we prompt assertions. First, it is false that certainty does not figure in appropriate prompts. For example, we may prompt an assertion by saying “Is it certain that p?”, “Is someone certain about what this is?”, or “Is there any certainty

about whether p ?” In a context where someone seems to act on p , we can prompt an assertion by saying “Are you sure that p ?”.

Turri recognizes these points but notes that it is far more natural to use ‘know’. However, this can be explained, if certainty is the norm of assertion, by the fact that using ‘certain’ to prompt an assertion leads to require too much of the subject. Indeed, such prompts in term of certainty require that the subject be certain that she is certain that p before asserting that p . To see this, note that by asking S whether she is certain that p , we are asking S to assert that p only if she can tell that she is certain that p , and thus, if certainty is the norm of assertion, only if she is certain that she is certain that p . The speech act prompted by the use of ‘certain’ would then be stronger than an assertion, perhaps, as Turri (2010, 461n8) suggests, it would rather be an act of guaranteeing. But this is perfectly coherent with, and explained by, the idea that certainty is the norm of assertion.

In contrast, note that if we grant Turri’s particular account of certainty, according to which certainty amounts to knowing that one knows, and the claim, endorsed by Turri, that knowledge is the norm of assertion, we are led to the view that in using ‘know’ to prompt an assertion, we are in fact requiring the subject to be certain that p before she asserts that p .⁴³ Indeed, by such prompts, the subject is asked to be in a position to assert that she knows that p , and hence, by the supposed knowledge norm of assertion, she is asked to know that she knows that p , before asserting that p . Hence, if certainty is knowledge of knowledge, the subject is asked to be certain. If an appropriate assertion responding to a prompt in term of knowledge in fact requires certainty from the subject, one may think that prompts in terms of knowledge do not count in favour of the knowledge account, after all, but rather count in favour of the certainty account.

In addition, on the supposition that the epistemic norm of assertion is certainty, one can easily explain why it is appropriate to prompt assertions by using ‘know’, even if one rejects the claim that certainty amounts to knowledge that one knows. If you can appropriately assert that you know that p , then, given the certainty account of assertion, you are certain that you know that p . By the factivity of knowledge, you are certain that p . Thus, by asking someone whether she knows that p , we may ask whether she is in a position to assert that p , even if certainty is the norm of assertion.

The claim that we prompt assertions by citing knowledge, which is, according to the certainty account, a mere necessary condition for warranted assertions, is in line with the other natural ways in which we prompt assertions. We may prompt an assertion by asking “What do you think? p or not p ?”. Some might think that something weaker than a flat-out assertion is prompted in this way, for example, merely an hedged assertion.⁴⁴ However, this is not necessarily the case. For

⁴³Turri (2010, 459) merely writes that certainty “requires” knowing that one knows, so perhaps he thinks that knowing that one knows is not sufficient for certainty. However, he strongly suggests that knowing that one knows is sufficient for certainty, because he argues that challenges in terms of certainty amounts to challenges regarding whether one knows that one knows. This could not be the case if more than knowledge that one knows were required for certainty.

⁴⁴See Benton (2014).

example, when asking for an advice to a doctor, it is not unnatural to say “What do you think Doctor? Should I take this medicine?”. The reply we expect is just a flat-out assertion: “Yes, you should” or “No, you should not”. It is also frequent to prompt an assertion by saying things like “Is it true that p?”, or “Is p?”. Therefore, an advocate of the knowledge account must also grant that we can prompt assertions by citing epistemic conditions weaker than the one involved in the supposed epistemic norm of assertion.

Next, consider appropriate challenges. It seems true that certainty does not figure in many challenges. But this does not show that it could not. According to Turri, however,

We needn't explain the propriety of 'You aren't certain of that!' or 'You're not absolutely sure of that!', because intuitively those challenges are out of line. They're too harsh, too demanding. (Turri, 2010, 460)

While I agree that the second challenge seems out of line, I do not see the first one as too harsh. Consider the following case:

STATION. You are in a railway station in an unfamiliar town, with Bill, and you both overhear a conversation between Jenny and Jill. Jenny is asking Jill where the Museum is. Jill answers: “I believe that the Museum is in the city centre”. Later, you meet John, who asks Bill where the Museum is. Bill flat-out assert: “It is in the city centre”. You know that the only information Bill has comes from what he heard from Jill and Jenny's conversation, and you say: “But, Bill, you aren't certain of that!”.

It does not seem to me that the challenge “You aren't certain of that” is too harsh, or out of line.

Finally, it is worth repeating here that, given the certainty norm of assertion, knowledge challenges are always sufficient. According to the certainty norm, a subject can assert that she knows that p only if she is certain that she knows that p, and hence, by factivity of knowledge, only if she is certain that p. So, one can equivalently challenge an assertion that p by asking whether the assertor is certain that p or whether she knows that p. But then, even if certainty is the norm of assertion, certainty need not figure in any challenge at all.

In sum, first, we may think that the fact that we prompt assertions by using ‘know’ rather than ‘certain’ in fact favours the certainty account over the knowledge account. And, second, although ‘certain’ does not figure in as full a range of challenges as knowledge does, which can be explained by the equivalence of challenges in term of knowledge and those in term of certainty, that does not mean that it cannot so figure.

6.5 The certainty norm of assertion and ‘knowledge’ ascriptions

So far, I have presented and defended the certainty account of assertion. In the remaining of this chapter, I shall show that if we adopt this account, we can fully account for why ‘knowledge’ ascriptions shift with practical factors, in an orthodox framework. We can do this, I will argue, in a better way than by endorsing the simple doxastic account or the Gricean account based on the rule of relevance.

6.5.1 Why ‘knowledge’ ascriptions shift with practical factors: the role of the certainty norm of assertion

How can we explain in an orthodox framework that ‘knowledge’ ascriptions shift with practical factors? A partial explanation is provided by the doxastic approach. When stakes rise, epistemic anxiety rises, and the subject, or the attributor, can lose her belief, and, thereby, her knowledge. However, there are cases where the (intuitive) doxastic situation of the subject and that of the attributor are not at issue. The explanation that I propose for these cases is based on the certainty account of assertion, and runs like this.

Suppose that we adopt the certainty norm of assertion. According to this norm, you should assert that p only if p is epistemically certain for you and you are subjectively certain that p . Suppose that certainty (epistemic and subjective) is context-sensitive. That means that what counts as certain depends on the context. As a result, whether it is epistemically appropriate for you to assert that p depends on the context. The proposition that S knows that p is not different in this regard. So, it is not surprising that the appropriateness of an assertion “ S knows that p ” shifts with the practical context.

In addition, note that the appropriateness of an assertion “ S knows that p ” will vary with the appropriateness of an assertion “ p ”. Indeed by an assertion “ S knows that p ”, you represent yourself as certain that S knows that p , and hence, by factivity of knowledge, as certain that p . Thus, if you do not meet the epistemic norm for asserting that p , you do not meet the epistemic norm for asserting that S knows that p .

In chapters 5, I have considered the main cases showing how practical factors may impact on the intuitive appropriateness of ‘knowledge’ ascriptions, and I have shown that the doxastic approach can handle a certain number of them, but not all. I shall now show that the certainty account of assertion, supplemented with a doxastic approach, can account for all the main cases.

Recall that the main cases are the following:

BANK CASES

Bank Case A (Low Stakes). My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays. I reply, No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon.

Bank Case B (High Stakes). My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, Banks do change their hours. Do you know that the bank will be open tomorrow? Remaining as confident as I was before that the bank will be open then, still, I reply, Well, no. I'd better go in and make sure.

AIRPORT. Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds Yes I know it does stop in Chicago. It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute. Mary and John agree that Smith doesn't really know that the plane will stop in Chicago. They decide to check with the airline agent.

GOOD EVIDENCE AIRPORT. The same as AIRPORT, except that Mary and John now have sufficient evidence and believe that the flight has a layover in Chicago. They still deny that Smith knows.

SMITH AIRPORT. The same as AIRPORT except that it is clear to Smith that Mary and John believes that the flight has a layover in Chicago on the basis of the flight itinerary, and that Mary and John agree that they know. Smith judges that they should not say that they know that the flight has a layover in Chicago.

First, consider BANK CASES. According to most orthodox epistemologists, in the high stakes case,

if Keith's confidence is not shaken, he knows that the bank is open. However, it remains to explain why, in this case, it is inappropriate for him to say that he does not know. According to the certainty account of assertion, if Keith says that he knows that the bank is open, he represents himself as certain that he knows that the bank is open. Hence, by factivity of knowledge, he represents himself as certain that the bank is open on Saturday. However, given the practical context associated with the proposition that the bank is open on Saturday, the standards for certainty are high. It is not epistemically certain for him that the bank is open, and he is not subjectively certain that the bank is open. Therefore, he is not in a position to assert that the bank is open, and hence, that he knows that the bank is open. This is why it is inappropriate for him to assert that he knows that the bank is open.

A worry might come from the fact that this approach does not explain why it may seem appropriate for Keith to say that he does not know that the bank is open, but merely why it is inappropriate for him to say that he knows that the bank is open.

However, it is crucial to realise that the high stakes scenario is underdescribed. We should be careful to distinguish two different possible high stakes bank cases.

Suppose it is assumed that Keith is somewhat confident that the bank is open, and that he believes (without being certain) that he knows that the bank is open. If so, it is unclear that it is appropriate for her to assert that he does not know. He should rather say something like "Well, I believe (that I know) that the bank is open but it is not certain (/ I am not certain)". This is the case which is explained by the certainty norm of assertion.

Now, suppose it is assumed that Keith believes that he does not know that the bank is open. In this case, we can appeal to the doxastic approach, and explain why Keith believes that he does not know by appealing to the fact that his belief that the bank is open is shaken by the perception of high stakes. In this case, the explanation why Keith can appropriately say that he does not know is straightforward.

In sum, the fact that it can be appropriate for Keith to assert that she does not know in the high stakes case does not compromise the approach appealing to the certainty account of assertion. Indeed, according to the orthodox epistemologist, there is one interpretation of the bank cases in which it is assumed that Keith knows and believes that he knows. Under this interpretation of the case, it is not appropriate for Keith to say that he does not know, and hence, there is no false, although appropriate, 'knowledge' denial to explain.

It pays to note that the fact that it is not inappropriate for Keith to say "Well, I believe (that I know) that the bank is open but it is not certain (/ I am not certain)" is actually in tension with shifty views of knowledge. Indeed, according to shifty views, in the high stakes case, Keith should say "I do not know" because such an assertion is true (in this context). But we should take note of the fact that the data to be explained are not that *it is appropriate* for Keith to say "I do not know". Rather, the data to be explained is that *it can be appropriate* for Keith to say "I do not

know”, and that *it can be appropriate* for Keith to say “Well, I believe (that I know) that the bank is open but it is not certain (/ I am not certain)”. But the fact that *it can be appropriate* for Keith to say “Well, I believe (that I know) that the bank is open but it is not certain (/ I am not certain)” is at odd with shifty views of knowledge, since if these views are correct, Keith obviously does not know.

Second, consider AIRPORT. The intuition is that it is inappropriate for high stakes Mary and John to assert that Smith knows that the flight stops in Chicago. This can be explained by the fact that they do not meet the certainty norm of assertion with respect to the proposition that the flight stops in Chicago, so they cannot assert that Smith (or themselves) know(s) that it stops in Chicago.

Now, as previously, we may say that this does not explain why it may seem appropriate for Mary to say “Smith does not really know”. However, as previously, we must be careful to distinguish two possible cases. Either Mary does not believe that she knows the target proposition (despite Smith’s evidence) or she does. If she does not, because her confidence is shaken, we can explain along the same lines as the doxastic approach why she appropriately denies ‘knowledge’ to Smith. It would be incoherent for her to acknowledge that Smith is in a position to know the target proposition while herself not believing this proposition. However, if Mary believes that she knows the target proposition, it becomes unclear that she could appropriately assert that Smith does not know. In such a case, it seems, she should assert “I am no sure!”, or “It is not certain that the flight stops in Chicago!”.

Third, consider GOOD EVIDENCE AIRPORT. Is Mary’s denial that Smith knows appropriate? Here, again, there are two possible cases. If Mary’s confidence has been shaken by the perception of high stakes (but is now stable given the new amount of evidence), we can appeal to the hindsight bias to explain in the same way as the doxastic approach why biased Mary asserts “Smith does not know that p”. Mary ascribes to Smith the same need-for-closure as she has, and thus, she takes Smith to need the same amount of evidence as she has to have an epistemically well-formed belief. If, in contrast, Mary is no biased, and she takes Smith’s evidence to be sufficient for knowledge, intuitively, it would not be appropriate for her to deny ‘knowledge’ to Smith.

Finally, consider SMITH AIRPORT. Mary believes the proposition on the basis of Smith’s evidence, and she asserts that she knows it. Smith is aware of Mary’s practical situation, and he assesses whether she knows the target proposition and whether her assertion is appropriate. Clearly, it would not be appropriate for Smith to say that Mary does not know the proposition, and it seems appropriate for Smith to say that, indeed, she knows. Still, Smith could also appropriately say that Mary should not assert that she knows. Appealing to the certainty norm of assertion explains this. Given Mary’s practical context, she is not certain that she knows the target proposition, and as a result, she should not assert that she knows it, even if this is true.

In sum, we can explain why ‘knowledge’ ascriptions shift with practical factors in an orthodox

framework by appealing to two different approaches, depending on how we flesh out the details of the case. When what is at issue is the (intuitive) doxastic situation of the attributor or that of the subject of attribution, the doxastic approach does well. When what is at issue is merely the (intuitive) inappropriateness of ‘knowledge’ ascriptions, we must appeal to the certainty account of assertion.

The fact that the certainty account of assertion explains why ‘knowledge’ ascriptions shift with practical factors, even when the doxastic states of the attributor or that of the subject are not at issue, bolsters the case in favour of the certainty account of assertion.

6.5.2 Why a mixed approach outperforms the simple doxastic approach

It is not difficult to see why an approach appealing to the certainty norm of assertion, supplemented with the doxastic approach, is better than the simple doxastic approach. What is problematic with the simple doxastic approach is that it aims at explaining all our intuitions concerning the appropriateness of ‘knowledge’ ascriptions in terms of the attributor’s (sometimes mistaken) intuitions regarding whether the subject (rationally) believes the proposition. On this approach, we do not have the resources to account for cases in which an unbiased assessor has the complex intuition that a subject knows a certain proposition but, still, that she herself (or, the subject) should not assert that the subject knows it, due to the practical situation.

In particular, recall what is problematic with Nagel’s proposal. First, this proposal cannot explain cases in which it is postulated that the confident high stakes subject has formed her belief in an epistemically impeccable way on modest (but knowledge-level) evidence, although it is still inappropriate for her to say that she knows. Second, there are cases where assuming the explicit presence of compromising-accuracy factors, such as pressure of time, renders the ‘knowledge’ ascription intuitively more appropriate. Third, it is unclear why the assessor should see the confident high stakes subject as not knowing, rather than as merely inappropriately asserting that she knows. Finally, it is problematic to assume that knowledge is incompatible with epistemic anxiety.

The approach appealing to the certainty norm of assertion avoids these problems. First, it can explain the case of a rationally confident high stake subject inappropriately asserting that she knows on the basis of the fact that her assertion is unwarranted, although true. Second, it can explain why postulating the presence of a compromising-accuracy factor, such as time pressure, can intuitively make a ‘knowledge’ ascription more warranted. Suppose that the presence of this compromising-accuracy factor is insufficient to make the belief unjustified. Still, the presence of this factor can lower the standards for certainty. If so, it becomes easier to appropriately (self-)ascribe ‘knowledge’. Third, this approach makes room for the (complex) intuition that a confident high stakes subject can know something on the basis of modest evidence although she should not say that she knows it, due to her practical situation. This is, I suppose, the intuition that orthodox epistemologists

have concerning (rationally) confident high stakes Keith. This was what the Gricean account based on the rule of relevance aimed at explaining. Fourth, the approach appealing to the certainty norm of assertion does not assume that knowledge is incompatible with epistemic anxiety.

6.5.3 Why a mixed approach outperforms the Gricean approach

The mixed approach that I propose is also superior to the Gricean account which is based on the rule of relevance. Indeed, it cannot be accused of being *ad hoc*. I have argued that the Gricean approach is *ad hoc* at least in that it does not explain, but merely assumes, that ‘know(s) that p’ ascriptions convey a good enough epistemic position with respect to p (rather than an insufficient epistemic position with respect to p). Although it seems true that ‘know(s) that p’ ascriptions convey a good enough epistemic position with respect to p, given the rule of quantity (assert the stronger!), this is surprising. Given that S does not assert the stronger “I am certain that p”, why is it that by asserting “I know that p”, she does not convey that she *merely* knows that p, and hence, that she is in an insufficient epistemic position with respect to p (supposing that more than knowledge is required in the context)?

In contrast, appealing to the certainty norm of assertion is not *ad hoc* because this norm applies to assertions in general, and not merely to ‘knowledge’ ascriptions.

Moreover, appealing to the certainty norm of assertion can explain the conversational and linguistic data that are at odd with the Gricean approach.

First, recall that the Gricean account does not explain why, in contexts where the goal of the conversation is to make a practical decision, ‘knowledge’ ascriptions are always conversationally relevant and decisive. Indeed, if the practical relevance of ‘knowledge’ ascriptions is merely due to a Gricean implicature, then there should be cases where the irrelevance of such ascriptions should clearly appear. In particular, if knowledge is not the norm of action, as the Gricean approach assumes, in contexts in which knowledge that p is not sufficient to act on p, it should be possible to challenge the relevance of claims to know that p. But we never challenge the relevance of such claims when it comes to make a decision (when p is relevant to action).

In contrast, given some plausible assumptions, the approach that I propose can offer an explanation why ‘knowledge’ ascriptions are bound to be always decisive and relevant when it comes to make a practical decision, even if knowledge is not the norm of action. The explanation is the following:

1. Certainty is the epistemic norm of assertion (assumption).
2. The epistemic norm of action does not require more than certainty (assumption).
3. If it is epistemically appropriate to assert that p, then it is epistemically appropriate to act on p (by 1 and 2).

4. If it is epistemically appropriate to assert that p , then it is epistemically appropriate to assert the obvious semantic implications of p (assumption).
5. p is an obvious semantic implication of S knows that p (assumption).
6. If it is epistemically appropriate to assert that S knows that p , then it is epistemically appropriate to assert that p (by 4 and 5).
7. If it is epistemically appropriate to assert that S knows that p , then it is epistemically appropriate to act on p (by 3 and 6).

The fact that if it is epistemically appropriate to assert ‘ S knows that p ’ then it is epistemically appropriate to act on p explains why, even if knowledge is not the norm of assertion, claims to know that p are always relevant and decisive to the question whether to act on p . This explanation assumes that the epistemic norm of action is not stronger than the epistemic norm of assertion. Therefore, unless we opt of an absolute certainty norm of action, this explanation can be accepted. (In chapter 10, I argue in favour of a certainty norm of action, similar to the certainty norm of assertion.)

A related issue has to do with the communicative and social functions of ‘knowledge’ ascriptions. Many philosophers have emphasized that when someone asserts ‘I know that p ’, she does not merely describe a state of affairs. She also gives her word for p .⁴⁵ More generally, it has been argued that the act of ascribing ‘knowledge’ often serves to direct action, to give assurance, etc. Arguably, the fact that ‘knowledge’ ascriptions can have these communicative functions should be explained (partially) by the semantics of ‘know’. Thus, for example, Lehrer writes:

Even if Austin is correct in declaring that when I say, ‘I know that S is P ’, I give others my authority for saying that S is P , the performance of this act is perfectly consistent with describing oneself as accepting with complete justification that S is P . Indeed, the assumption that I am making such a descriptive claim about myself when I say, ‘I know that S is P ’, helps to explain the way in which I give my authority for saying that S is P by saying, ‘I know that S is P ’. If I were not claiming to be completely justified in accepting that S is P when I say, ‘I know that S is P ’, then why in the world should my saying, ‘I know that S is P ’ be taken as giving my authority for saying that S is P ? It might be more reasonably be taken as an expression of opinionless agnosticism. (Lehrer, 1990, 28)

In this passage, Lehrer seeks to explain the Austinian function of ‘knowledge’ ascriptions by referring to the semantics of ‘know’, and, on his view, to the fact that someone knows that p only if she has complete justification for p .

⁴⁵See Austin (1959).

However, plausibly, the epistemic standards that we must meet with respect to p for appropriately give others our authority as to whether p , to recommend acting on p , etc., vary with the context. Now, if knowledge requires complete justification in the sense of absolutely maximal justification, then we can explain why ‘knowledge’ ascriptions can fulfil these communicative functions, but scepticism threatens. On the other hand, if knowledge does not require complete justification in the sense of absolutely maximal justification, then there are cases where the epistemic standards for knowledge are insufficient to appropriately give others our authority (think about high stakes cases). If so, we cannot explain the fact that ‘knowledge’ ascriptions have these communicative functions on the basis of the fact that knowledge requires complete justification.

It is then sometimes argued that the fact that ‘knowledge’ ascriptions have these functions speaks in favour of an unorthodox view of knowledge, in which the epistemic standards of knowledge shift with the context.⁴⁶

We should note, however, that an orthodox (non-sceptical) semantics of ‘know’, associated with the certainty account of assertion, easily explains why ‘knowledge’ ascriptions can fulfil these functions. Again, according to the certainty account of assertion, ‘knowledge’ ascriptions represent the ascriber as certain that she knows that p , and by factivity of knowledge, as certain that p . Thus, these ascriptions can serve as recommendations to act on p , to give assurance that p , etc., in all contexts.⁴⁷

The second problem for the Gricean approach is that the pragmatic implication that it assumes is not cancellable. Indeed, on this view, in the high stakes bank case, the pragmatic meaning of an assertion “I know that the bank is open” is supposed to be that Keith’s epistemic position is good enough to come back without checking. However, suppose that Keith says:

- (1) I know that the bank is open, but I do not mean that we should not check. The stakes are high.

This seems incoherent. More generally,

- (2) S / We / I know(s) that p but we should check further

sounds incoherent. As Cohen (1999, 58-60) writes, “to my ear, it is hard to make sense of that claim”. We can think that the issue is pragmatic.⁴⁸ But the idea that we should explain the oddity of (1) and (2) by the fact that an assertion “I know that p ” pragmatically imparts in a Gricean way that we should not check whether p is problematic. Indeed, on this approach, (1) and (2) should just be instances of cancellation of this implicature. Hence, they should not sound odd.

In contrast, the approach appealing to the certainty account of assertion easily explains why the pragmatic meaning associated with an assertion “I know that p ”, namely, that I take myself

⁴⁶See Henderson (2009, 2011); Lawlor (2013) and Gerken (forthcoming) for discussion.

⁴⁷In chapter 8 section 3.2 I discuss another proposal, from Gerken, and I argue that it is *ad hoc*.

⁴⁸See Rysiew (2001, 495), following Sosa.

to be certain that p , and hence, in a good enough epistemic position with respect to p to act on p , is not cancellable. Thus, this approach explains assertions of (1) and (2) are bound to sound incoherent. Again, assuming that the epistemic norm of action is not stronger than the epistemic norm of assertion, if you represent yourself in a good enough epistemic position to assert that you know that p , and thereby, in a good enough epistemic position to assert that p , you represent yourself in a good enough epistemic position to act on p . Hence the oddity of asserting in the same breath “I know that p ” and “we should check further”. The former is appropriate only if you meet the norm of assertion with respect to p , and thus, only if you meet the epistemic norm of action with respect to p , but the latter is true (supposing that there are no other non-epistemic reasons to check) only if you are not in a good enough epistemic position to act on p .

Finally, note that

- (3) S knows that p but she should check because she is not certain that p (/ because she does not know that she knows that p)

is not obviously incoherent. At least, it does not sound as incoherent as (1) and (2). This asymmetry speaks in favour of the idea that the explanation of the incoherence of asserting (1) and (2) must be pragmatic and related to the epistemic norm of assertion.

6.6 Conclusion

In this chapter, I have argued in favour of the certainty account of assertion. I have shown that certainty is often achievable and that this account explains in a unified way all the conversational and linguistic data that can be found in the literature. Further, this account satisfies many theoretical desiderata. On the basis of this account, I have proposed a mixed orthodox explanation of patterns of ‘knowledge’ ascriptions. This explanation invokes the certainty account of assertion when the assertor is seen as knowing (or justifiably believing) that the subject of attribution knows the proposition, although, given the practical context, it is inappropriate for the assertor to assert that the subject knows. This explanation invokes the doxastic situation of the ascriber or that of the subject, when the assertor, or the subject, is seen as not believing the relevant proposition. I have emphasized that this mixed approach can explain all the relevant cases, and that it does better in this regard than the simple doxastic approach and the Gricean account appealing to the rule of relevance. We can conclude that the patterns of ‘knowledge’ ascriptions that are presented as a challenge to the orthodox account of knowledge can actually be explained in an orthodox framework. Thus, the second challenge can be met. In the next chapter, I will start addressing the third challenge, concerning the intuitive normative relation between knowledge and action.

Part III

Knowledge and rational action

Chapter 7

The epistemic evaluation of actions

7.1 Introduction

In the previous chapter, I have shown that we can explain in an orthodox framework why ‘knowledge’ ascriptions shift with practical factors by appealing partially to a psychological approach and partially to the certainty account of assertion. A different challenge for the orthodox view of knowledge comes from an argument based upon a supposed normative connection between knowledge and action. This argument has two premises. The first one states that the epistemic standards that a subject must satisfy with respect to a certain consideration to be rational to act on it varies with practical factors. The second one is that knowledge (or epistemic justification to believe) is the epistemic norm of action. The conclusion is that knowledge (or epistemic justification to believe) varies with practical factors.

In this chapter, I lay the groundwork for discussing this argument. I start by presenting the idea of evaluation of action in general. I identify three levels of evaluation of action, two of which involving a consideration of the agent’s epistemic position. I show how these levels of evaluation may be related to the notions of justification, rationality and excuse. Next, I consider the notion of epistemic evaluation of action. I show that this notion is equivocal because there are at least six things that can be evaluated by considering the subject’s epistemic position. Finally, I present the main accounts of the epistemic norm of action and I identify four ways in which the orthodox philosopher might wish to reply to the argument from the supposed normative connection between knowledge and action. First, she might deny the supposed normative connection between knowledge and action. Second, she might grant that practical factors can affect whether the subject is in a good enough epistemic position to act on a certain consideration, and that knowledge is the epistemic norm of action, while maintaining that, correctly understood, these claims are compatible with the orthodox account of knowledge. Third, she might argue that although there is a normative

connection between knowledge and action, there are several epistemic norms of action. Fourth, she might say that there is a normative connection between knowledge and action in that knowledge is necessary, but not sufficient, for action.

7.2 The evaluation of action

7.2.1 The normative and the descriptive

There is a fundamental distinction between the normative and the descriptive. Descriptive statements attempt to state what is the case. Normative statements attempt to state what ought to be the case, or what one ought to do (believe, feel), or what it would be good to do (believe, feel).¹

Noting a possible ambiguity in the use of the expression ‘normative statement’ is a good way of seeing the distinction between the descriptive and the normative. Indeed, the expression ‘normative statement’ may be used to refer either to descriptive statements referring to norms (following von Wright, we may call these statements ‘norm-propositions’), or to statements constitutive of norms. As von Wright (1963, viii) writes:

Tokens of the same sentences are used, sometimes to enunciate a prescription (that is, to enjoin, permit, or prohibit a certain action), sometimes again to express a proposition to the effect that there is a prescription enjoining, or permitting or prohibiting a certain action.

To illustrate, suppose that the king is a legitimate authority, that is, a source of norms, and says: “You ought to join the army”. Hence, there is a norm to the effect that you ought to join the army. By saying this, the king makes a prescription that may be rephrased as an imperative (“Join the army!”) that has a normative force to the extent that the king is the authority. Now, suppose that the king’s servant says: “You ought to join the army”. The servant is not a source of norm. Hence, his statement cannot be rephrased with an imperative form. Rather, it may be rephrased with the description that there is a norm according to which you ought to join the army (e.g. “The king said you ought to join the army”). While descriptions can be true or false, a norm (an order, or a prescription) cannot be true or false. However, in contrast to descriptive claims, norms have a certain force.

The fact that norms have a certain force can also be seen in the distinction between normative and non-normative evaluations. We use standards to assess actions, persons, beliefs, etc., and we say that these actions, persons, beliefs, etc., are correct or incorrect, right or wrong, good or bad, depending on whether they conform to (or comply with) the standards in question. Still, not every

¹A traditional way of drawing this distinction appeals to the distinction between “fact” and “value”, or to the Humean idea that there is an “is-ought gap” in that no ought-claim can be logically derived from an is-claim. See Pidgen (2013).

assessment (or underlying standard) is normative. For instance, when we say that a tree has bad roots, we refer to a standard of roots that may have to do with the proper functioning of roots. We may think that this evaluation is descriptive, or not truly normative.² It is a mere claim about the fact that, for example, given its roots, the tree won't complete its life-cycle. This does not mean that the state of affair in question is disvaluable, or that one should or ought to do something about it. In contrast, when we say that lying is morally bad, this seems to have normative implications. If moral requirements are not optional, that means, for example, that if you are lying, there is a force commanding you to stop lying. If you do not correctly respond to this force, you can be criticized or blamed. To say that lying is morally bad means that there is a normative reason not to lie.³

In the same way as there are normative and a non-normative senses of 'good' or 'bad', there are normative and non-normative senses of 'ought', 'right' and 'wrong', 'correct' and 'incorrect', 'reason' and 'norm'.

7.2.2 Normative reasons and what you ought to do

Norms are standards that one ought to satisfy. If an action satisfies the norm, the action is right with respect to this norm. A consideration that shows that an action (a belief, a emotion) is right or conform to a norm, or that explains why this action (belief, emotion) is right or conform to a norm, is a normative reason.⁴ Some philosophers also say that a normative reason to act (believe, feel) is a consideration that makes it right to act (believe, feel).⁵ For example, consider the evaluation of beliefs. Suppose that truth is the norm of beliefs. If so, a belief is right if and only if it is true. Then, as a normative reason to believe that p is a consideration which shows that believing p is right, this is a consideration that shows that p is true. For example, if p is certain given q , q is a reason to believe p . In showing that p is true, q shows that believing p is right.

There are several normative sources, and hence, different types of reasons. There are aesthetic reasons, morals reasons, legal reasons, prudential reasons, etc. They can give rise to different evaluations of the same action. Thus, lying may be prudentially right but morally wrong. In order to determine what a subject ought to do, all things considered, we consider the normative reasons that apply to her and what these reasons favour, overall. We say that a subject S has sufficient or *pro toto* reason to ϕ if S has a reason to ϕ that is not defeated by a reason for S not to ϕ .

A reason can be defeated because it is outweighed by an opposite reason. Whether or not a reason outweighs another reason is a matter of balance between their respective normative forces. For instance, suppose that you know that there is a moral norm saying that you ought not to lie. Suppose also that you know that lying will make you happy. You have a reason to lie: this will

²See Broome (2013, 11-12).

³Parfit (2011, 39-40) calls the normative goodness the "reason-involving goodness".

⁴See Dancy (2000, 2, 7).

⁵See Alvarez (2010, 13-14).

make you happy. Happiness is valuable, and we may suppose that you have a reason to do what has value, and hence, that which will make you happy. You also have a reason not to lie: this is morally wrong. Whether or not you, overall, ought to lie depends on whether it is more important to be morally right or to be happy. In some cases, what is more important is a matter of dispute.

According to some philosophers, a reason can also be defeated because there is a second-order reason to exclude it (not to act on it or not to treat it as a reason).⁶ For example, it may be argued that as a soldier, you should obey the order to steal a truck even if the balance of reasons speaks in favour of not stealing it. Exclusionary reasons are reasons to ignore the balance of reasons. However, whether there are such exclusionary reasons is controversial.

A reason can also be defeated because it is cancelled, or undermined. A reason is cancelled if an enabling condition is absent. For instance, suppose that I promise under constraint to do something. The fact that I promised to do it is a reason to do it, but the fact that I was constrained cancels the normative force of the promise. A canceling condition is not necessarily a reason for something.

When a reason is outweighed, or when it is excluded, it is still a (*pro tanto*) reason. That is, this reason continues to have a normative force.⁷ When a reason is cancelled, the consideration in question no longer counts as a normative reason, since it no longer has normative force. However, in contrast to outweighed reasons, and similarly to undermined reasons, excluded reasons do not count when it comes to determine what one ought to do.

In sum, as there are different normative sources, the same action may be evaluated in light of different norms. Whether an action is, overall, what ought to be done, depends on the relative importance of these different norms, and more specifically, on the strength of the reasons that apply.

7.2.3 Deontic and evaluative normativity

We must also distinguish between deontic assessments, having to do with obligations, interdiction and permission, and evaluative assessments, invoking what is bad, good, or indifferent.⁸

To illustrate, consider the normative ‘ought’. It has an agential sense (“ought to do”) and a non-agential sense (“ought to be”). The agential sense is related to the notions of will, reason and obligation in a way in which the non-agential sense is not. For instance, it is related to some plausible principles, such as “ought implies can”, or “it must be possible to do or not to do what one ought to do”. In contrast, the non-agential sense rather refers to an ideal that may be inaccessible. We may say that if the ideal is not accessible, there is merely a reason to want or desire it, but no reason to adopt an intention to realize it. Further, if there are supererogatory actions, actions

⁶See Raz (1999, 35-40).

⁷See Dancy (2004, 29) for the idea that outweighed reasons explain residual duties and regrets, and Raz (1999, 43-45) for the idea that exclusionary reasons can explain some mixed evaluations.

⁸See Tappolet (2014).

that would be best but that are not required, it is clear that we should distinguish requirements and reasons (what you ought to do, what you have (most) reason to do) from values (what would be best to do). There is no consensus on whether and how evaluative and deontic concepts and assessments relate to one another.⁹

7.2.4 Deontic norms

Let us consider deontic norms. We must distinguish between the character of a norm, its conditions of application and its content.¹⁰ The character of a norm concerns whether it is an obligation, a permission, or an interdiction. Its content is what the norm prescribes (or authorizes). Its conditions of application are conditions that must obtain for the norm to prescribe (or authorize) its content. Thus, for example, the content of the deontic norm according to which you ought to take an umbrella if it rains is that you take an umbrella, and its conditions of application include that it rains.¹¹ Often, the word ‘content’ is also used to refer to the conditions of application of the norm.

Further, the word ‘condition’ is often used not for the conditions of application thus defined but, instead, for the conditions under which the norm is in force. These are two different things. A norm can apply, or be in force, even if its content is not prescribed because the conditions of application of this norm do not obtain. For example, suppose that you are in the UK. You are under the authority of a norm according to which you ought, if you drive, drive on the left side of the road. This norm is in force even if you are not driving. If you are in the UK and you are not driving, the condition of application of this norm is not realised, but the condition under which the norm applies is realised.

To see this point, it is useful to consider the logical structure of deontic norms and to distinguish conditioned norms, hypothetical (or conditional) norms, and categorical (or absolute) norms.

The logical structure of an absolute norm is the following (where ‘O’ stands for an obligation, ‘ φ ’ for the content of the norm, and ‘p’ for a condition of application):

- O (φ)

The logical structure of an hypothetical (or conditional) norms is the following:

⁹Concerning the relations between evaluative and deontic concepts see Tappolet (2014).

¹⁰See von Wright (1963, 71-73).

¹¹It is worth noting that if most conditions of application of a certain norm are specific to this norm, some are not. For example, the condition of application that there is an opportunity for doing what the norm prescribes (or authorizes) applies to all norms, whether categorical or hypothetical. Thus, besides the specific condition of application that it rains, the conditions of application of the hypothetical norm according to which you ought to take an umbrella if it rains also include as a condition of application that you do not already have an umbrella and that there is an action available to you such that by performing it you will have an umbrella (ought implies can). As the specific condition of application following from the general condition of application that there is an opportunity for doing what the norm prescribes (or authorizes) may be derived from the content of the specific norm, it need not figure explicitly among the conditions of application. See von Wright (1963, 74-75).

- $O(\varphi \text{ if(f) } p)$

The logical structure of a conditioned norm is the following:

- If X then $O(\varphi)$ ¹²

Kantian categorical imperatives are examples of categorical norms. The norm according to which you ought to take your umbrella if it rains is an example of hypothetical norm. The laws of a state are examples of conditioned norms. They apply to you if you are in their domain of jurisdiction.

7.2.5 Objective and subjective reasons

When we assess an action, we rely on the different norms we think apply in the situation. We balance the reasons the subject has and we determine what she ought to do all things considered.

There are several debates concerning the necessary conditions for there being a normative reason for a subject to do something (and for a standard to be normative for a subject). One of them concerns the role played by the subject's desires.¹³ According to one kind of theories, sometimes called 'normative existence internalism about reasons', a necessary condition on normative reasons is that a certain (counterfactual) motivational fact about the agent obtains (e.g., that the agent has such or such desire, or would have such or such desire in such or such circumstances).¹⁴ According to another kind of theories, sometimes called 'normative existence externalism about reasons', the subject's (counterfactual) desires are typically irrelevant. Rather, the normative reasons have to do with the value, or the goodness, of the state of affair that they favour independently of the subjects (counterfactual) desires.¹⁵ (Note that on the buck passing account of values, the value of a state of affair is nothing more than its propriety of having reason-giving (natural) properties.¹⁶ Thus, on this account, to say that a state of affair has value just amounts to saying that its natural properties give us a reason to promote it, or to have a pro-attitude toward it.)

¹²There are different kinds of conditioned norms depending on how we understand the conditional (as a material, subjunctive, or strict conditional). For more on this point, see Feldman (1983). Broome (2013, 134-135) calls the conditioned norms "requirements that are conditional in their application", and hypothetical norms "requirements that are conditional in their content".

¹³For a general presentation, see Finlay and Schroeder (2012); Lenman (2009).

¹⁴This kind of view is defended in particular by Williams (1980). This kind of theory splits into two species. On the first one, sometimes called 'the desire-based theory of reasons' or 'the Humean theory of reasons', the motivational (counterfactual) fact is constitutive of the existence of the reason and explains its normative force. This approach is reductionist if it says that we can explain the normative force of reasons by reducing it to the natural force of desires. For a recent defence of a reductionist approach, see Schroeder (2007) and for criticisms, see Scanlon (1998) and ?; Parfit (2011). On the second one, sometimes called the 'rationality-based theory of reasons', the motivational (counterfactual) fact is merely a necessary condition on normative reasons, but it does not explain the existence or the normative force of the reason (see, e.g., Korsgaard (1986)).

¹⁵The view that, typically, there is no necessary condition on normative reasons arising from the subject's (counterfactual) motivational set is sometimes called 'the valued-based theory of reasons'. For a defence, see Scanlon (1998), Parfit (2001), Parfit (2011), Alvarez (2010).

¹⁶See Scanlon (1998).

I will focus on a different debate, concerning the role played by the subject's epistemic position (or perspective). Some philosophers think that, typically, or at least in some cases, there is no necessary epistemic condition on a subject's normative reasons. On this view, whether p is a normative reason for S to φ , or whether S ought to φ , does not typically depend on S 's epistemic position with respect to p . Other philosophers think that the normative reasons there are for S depend on her epistemic position. If S ought to φ , then the things that are in S 's epistemic perspective must favour, overall, φ -ing. To evaluate whether S ought to φ , we must consider S 's epistemic perspective.

To illustrate, consider the following case:¹⁷

MINERS. Ten miners are trapped either in shaft A or in shaft B, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed.

Suppose that the miners are in fact in shaft A. Blocking shaft A is the materially right action. An action is materially right when it is the action favoured by all the facts.¹⁸ These facts are often called "reasons there are" in contrast to "reasons one has", or "objective reasons" in contrast to "subjective" reasons.

Philosophers disagree on what reasons are. Some philosophers think that what is often called "subjective reasons" are not reasons, and others that what is often called "objective reasons" are not reasons.¹⁹ Yet others think that the notion of reason is ambiguous, and that there are objective and subjective reasons.²⁰ In what follows, for ease of exposition, I will use the terminology of subjective and objective reason, but I leave open whether what is thus called really is a reason.

Those who think that normative reasons do not typically depend on the subject's epistemic position will presumably say that, in MINERS, what the subject ought to do is to block shaft A. According to them, the normative reasons are the objective reasons. The objective reasons determine what the subject ought to do.

However, some may have the intuition that, in MINERS, what the subject ought to do is to block neither shaft. According to this intuition, what you ought to do is the formally right action, the action that is right given your limited perspective. The considerations that are determined by the subject's perspective are often called "subjective reasons" in contrast to "objective reasons", or

¹⁷This case is discussed in particular by Kolodny and MacFarlane (2010, 1-2) and Parfit (2011, 159).

¹⁸The materially right is what Parfit (2011, 32) calls the "right or wrong in the fact-relative sense".

¹⁹Thus, those who think that reasons are facts independent from the subject's epistemic perspective often use "reasons one has" to mean objective reasons (see, e.g., Parfit (2011) and Alvarez (2010, 22-23)). They think that subjective reasons are apparent reasons, and hence, that they are not always reasons.

²⁰See Schroeder (2008).

“reasons the subject has” in contrast to “reasons there are”. According to those who think that the subject ought to close neither shaft, the normative reasons are the subjective reasons.

There are different accounts of subjective reasons. It is common ground that subjective reasons are accessible to the subject. They are provided by the subject’s epistemic position or (some of) her mental states.

First of all, it is important to note that we may say that subjective reasons are provided by (some of) the subject’s mental states without endorsing the claim that subjective reasons are mental states (although these claims are compatible). It is disputed that reasons are mental states.

For example, Broome (1999) has argued that even if the mental states of a subject may require something of her, that does not mean that these mental states are reasons. Firstly, the idea that the relation between a mental state and what it requires is a reason relation leads to bootstrapping. To illustrate, suppose that the fact that you believe that p , and that p entails q , is a (subjective) reason for you to believe q . Then, the fact that you believe that p and that p entails p should also be a (subjective) reason for you to believe p . The mere fact that you have a belief would generate a (subjective) reason in favour of this belief. However, we cannot generate reasons in this way.

Secondly, you can correctly respond to the alleged subjective reason that you believe that p and that p implies q in different ways. You can either believe q , or drop the belief that p , or the belief that if p then q . Your beliefs (and more generally your mental attitudes and actions) are sensitive to requirements of coherence. These requirements have wide scope, like conditional norms. They tell you, for example, that you ought [if you believe that p and that p implies q , then believe that q]. But they do not tell you that if you believe that p and that p implies q , then you ought believe q . Indeed, if q is absurd, it seems it would be more reasonable for you to retract your belief that p or your belief that p implies q rather than believing q . Thus, the requirements that apply to you in virtue of the fact that you have certain mental attitudes do not give you reasons. They do not allow to detach a reason for a specific attitude or action. For example, even if you believe that you ought to kill someone, and that it is incoherent to believe that you ought to φ without intending to φ , that does not imply that you have a reason to intend to kill someone. Perhaps, instead, you have a reason not to believe that you ought to kill someone.

If we think that subjective reasons are provided by (some of) the subject’s mental states but do not consist of mental states, we may say that the mental states are mere enabling conditions. They are not themselves subjective reasons, but necessary conditions for having subjective reasons.²¹

Although it is common ground that subjective reasons are considerations accessible to the subject, there are many ways of understanding the notion of accessibility. According to internalist approaches of subjective reasons, what is accessible to a subject is what is provided by (some of) her non-factive mental states, or what she can know *a priori* on the basis of mere reflection. On this

²¹See Williamson (2000, 197) and Littlejohn (2012, 93). On the idea of enabling conditions, see Audi (2001, 26) and Dancy (2004, 38-39).

kind of views, a mere difference in the external world cannot make a difference on the subjective reasons, and if subjective reasons determine what you ought to do, a mere difference in the external world cannot make a difference on what you ought to do.

According to externalist approaches of subjective reasons, external facts may be accessible, and thus, they can be relevant to the set of subjective reasons (and if subjective reasons determine what you ought to do, there are relevant to what you ought to do). According to a view called “the factoring account”, for example, subjective reasons are facts that you possess.²² On some versions of this type of view, the subject must know the relevant proposition in order to possess it as a reason. On other versions of this type of view, it is sufficient for the subject to be in a position to know the reason-giving fact in order to possess it as a reason (it is not necessary for the subject to believe it).²³ If our subjective reasons are what we are in a position to know, presumably, subjective reasons are not mental states.²⁴

Besides the distinction between internalist and externalist views about subjective reasons, we must distinguish normativist and non-normativist accounts of the accessibility relation.²⁵ According to normativist accounts, the relation that must hold between a subject and the relevant consideration for this subject to have a subjective reason involves a normative ingredient. According to non-normativist views, this is not the case.

Let me give two examples. A non-normativist about the accessibility relation advocating for the factoring account may say that subjective reasons only require true beliefs. Alternatively, she may say that subjective reasons require knowledge, but that knowledge is a non-normative notion.²⁶ Yet another possibility is to say that there are non-normative factive mental states, such as perceiving, that provide us with subjective reasons. An advocate of a normative notion of accessibility rejecting the factoring account may say that the subject must only be justified in believing the relevant consideration to have it as a subjective reason.²⁷

As said above, there is no consensus on whether the normative reasons – those which determine what we ought to do – are the objective reasons or the subjective reasons.²⁸ We might be tempted to opt for a kind of normative pluralism in which both objective reasons and subjective reasons determine what one ought to do. However, on the one hand, we do not and cannot weigh together

²²See Schroeder (2008) for a criticism of the view and Lord (2010) for a defence.

²³See Lord (2010) for the claim that subjective reasons are facts we are in a position to know and Gibbons (2013) for the claim that subjective reasons include facts we are in a position to know.

²⁴Although see Gibbons (2013, 189) for the idea that being in a position to know is a mental state.

²⁵See Sylvan (2016).

²⁶See Sylvan (2016).

²⁷One might think that such a view is doomed to a regress or a circle. The notion of having a reason is analysed in terms of justification, but, it seems, the notion of justification should be analysed in terms of having (sufficient) reasons. However, there are various ways of avoiding these problematic consequences. For example, we may say that the notion of justification should not be always analysed in terms of (sufficient) reasons-had (see Sylvan (2016) for further discussion).

²⁸While in ethics, it seems popular to think of what you ought to do in terms of objective reasons, in epistemology, it is more popular to think of what you ought to believe in terms of subjective reasons.

objective and subjective reasons as we weigh, say, prudential reasons and aesthetic reasons.²⁹ On the other hand, if we accept that there are two senses of ‘ought’, one subjective and one objective, we are led to ask which ought is the normative one. If we say that only one of them is normative, then we grant that the other one is irrelevant to determine what one ought to do in the normative sense. If we say that both oughts are normative, we will have to say that, in situations where the two oughts gives different verdicts, we ought to φ and we ought not to φ . But if one thinks, as the cognitivist, that normative statements express normative propositions that can be true or false, as the proposition that you ought to φ and you ought not to φ is a contradiction, it cannot be true. Alternatively, if one thinks, as the non-cognitivist, that normative statements express pro-attitudes toward a certain action (or attitude), that means that we have incoherent attitudes. As a result, pluralism does not look like a promising option.³⁰

In this chapter, my aim is to have a better grasp of the notions of epistemic norm of action, epistemic reason to act, and epistemic evaluation of action. These notions figure in various ways in versions of the argument against the orthodox account of knowledge from the (supposed) normative connection between knowledge and action. The distinction between objective and subjective reasons provides us with a first way of understanding these notions. Indeed, in general, the distinction between objective and subjective reasons is a distinction between things that need not be epistemically accessible to the subject and things that need to be epistemically accessible to the subject. Thus, one natural way of understanding the role of epistemic considerations when evaluating actions is related to this distinction. As epistemic considerations play a role when it comes to subjective reason, one natural way of understanding these notions is to say that an epistemic evaluation of action is an assessment of the subjective reasons that the agent has, epistemic reasons to act are subjective reasons, and the condition of application (or the satisfaction condition) of the epistemic norm of action is the epistemic condition on subjective reasons.

As there is no consensus on whether or not subjective reasons are normative considerations, though, there are two possible views associated with this way of understanding these notions of epistemic evaluations, epistemic reasons to act, and epistemic norm of action. If the subjective reasons are the normative reasons, the epistemic evaluation of action understood in this sense has normative consequences, and the underlying epistemic norm of action is relevant to determine what a subject ought to do. If the normative reasons are the objective reasons, however, the epistemic norm of action is not a normative standard. At least, it can be relevant to determine the rationality of the action, as I will now explain.

²⁹See, e.g., Dancy (2009).

³⁰See Gibbons (2013).

7.2.6 Subjective reasons, motivating reasons and rationality

Subjective reasons are considerations accessible to the subject. As such, they can get a grip on the subject, guide her, and figure in rationalizing explanations of her actions *in the circumstances in which the subject actually is*. Arguably, objective reasons are considerations that can *in principle* get a grip on the subject and rationalize her actions, but not all objective reasons can guide the subject in her actual circumstances, given that she can be ignorant of some of them. (A different question, that I leave aside because my interest is only on the epistemic condition on subjective reasons, is whether we should assume the Humean view of motivation according to which for a consideration to get a grip on the subject, she must have an appropriate desire.)

In brief, there is a special connection between subjective reasons and motivation to act.³¹ When a subject acts on a consideration which is a subjective reason, this subjective reason is her motivating reason. When a subject acts on a consideration that is not a subjective reason, this motivating reason can explain the action but it cannot make it rational.

Some philosophers identify subjective and motivating reasons, and hence, they say that there are (rationally) good and bad subjective reasons, those that make the action rational and those that do not.³² However, I prefer to say that subjective reasons are the considerations that can make the action rational, and that motivating reasons give the rationalizing explanation of the action. The action is rational only when the motivating reasons are also subjective reasons, that is, when they are rationally good. The action is irrational if the motivating reasons, the reasons for which the action is performed, are not subjective reasons. Some subjective reasons may not motivate, in which case they are not among the reasons for which the subject act.

The distinction between subjective and (rationally good) motivating reasons is required in order to acknowledge the fact that to evaluate an action, it is not sufficient to cite the subjective reasons the agent has. We must identify the reason upon which she acted. Subjective reasons are not always causes of the action, or they do not always cause the action in the right way. A subject may do the rational thing irrationally. The rational thing to do for a subject is determined by her subjective reasons. Whether a subject rationally does the rational thing depends on her motivating reasons.

To illustrate, consider MINERS and assume that the rational thing to do for the subject is to block neither shafts. There are two ways in which the subject can do the rational thing irrationally, although as a result of being motivated by a reason that can explain her action. First, she can do it for a rationally bad reason. For example, suppose that the subject indeed blocks neither shafts. However, suppose that she does it because although she intended to block one shaft, someone promised her a huge amount of money for not doing so. This subject has most subjective reason to block neither shafts and she does it, but she does not do it for the reasons that make it the rational thing to do. Second, a subject can do the rational thing to do for a rationally good reason without

³¹See Gibbons (2013, 135-136).

³²See Gibbons (2013, 31-32).

appreciating that it is rationally good. For example, suppose that the subject blocks neither shaft for a sufficient subjective reason, say, that this action has the best expected utility. However, suppose that she thinks that this consideration does not, or is not sufficient to, make the action rational. The subject is motivated by a rationally good reason, but not in a good way, rationally speaking.

To capture the distinction between doing the rational thing and doing the rational thing rationally, it is standard to distinguish *ex ante* rationality and *ex post* rationality.³³ In both of the examples above, the subject's action is *ex ante* rational but not *ex post* rational. The subject does the rational thing for a reason but she does not act rationally. Either she acts for a rationally bad reason (her motivating reason is not a sufficient subjective reason), or she acts for a rationally good reason without appreciating that it is rationally good.

In sum, subjective reasons determine the rational thing to do, that is, *ex ante* rationality. Motivating reasons determine whether the subject rationally does the rational thing, that is, *ex post* rationality. By using the expression 'epistemic evaluation of action', we may thus refer to an evaluation of action appealing to the epistemic condition that a subject must satisfy with respect to the considerations that make her action the rational thing to do. If the subject does not satisfy the relevant epistemic condition, her motivating reasons cannot be rationally good. Indeed, if a certain action is not the rational thing to do, it cannot be done rationally.³⁴ And if the subject's motivating reasons are not rationally good, the subject does not perform the action rationally.

If the subjective reasons are not the normative reasons, the epistemic evaluation of action thus understood has no normative force, but it is still an evaluation of the rationality of the action. This approach should then assume that rationality is not normative.³⁵ If the subjective reasons are the normative reasons, the epistemic evaluation of action thus understood is normative. In the first case, we may say that the epistemic norm of action is derived from a necessary epistemic condition on the considerations that can make an action rational. In the second case, we may say that the epistemic norm of action derives from a necessary epistemic condition on normative reasons.

7.2.7 Justification, rationality, excuse

Subjective reasons are considerations that determine the rational thing to do. That is, these are considerations upon which a rational subject can act rationally. Normative reasons are considerations that determine what you ought to do. If the normative reasons are the subjective reasons, what you ought to do just is what it is rational for you to do. Often, philosophers use the term 'justified' in relation to what you ought to do, and call normative reasons 'justifying reasons'. On

³³Regarding epistemic rationality, philosophers often distinguish between propositional justification and doxastic justification.

³⁴The claim that we can rationally do something only if it is the rational thing to do is parallel to the claim that doxastic justification requires propositional justification. See, e.g., Littlejohn (2012, 7).

³⁵See Kolodny (2005).

this terminology, if you ought to do what it is rational to do, rationality and justification come to the same thing. If not, you can rationally do the rational thing without being justified.

However, the terms ‘justified’ and ‘rational’ are sometimes used in evaluations that do not directly bear on objective reasons or subjective reasons. To see this, consider the following case, from Skorupski (2010, 126):

FIRE ALARM. Jane is the director of a nuclear power station. She hears the fire alarm, forms the belief that a fire is being signaled and immediately presses the button for evacuation. She has forgotten that at the beginning of the month she ordered that the alarms should be tested this morning—by being rung at this precise time to a specific pattern, careful checks having previously been made to establish that there was no fire. She hasn’t forgotten this in the sense that she is unable to recall it but only in the sense that she is preoccupied when the bell rings, and in her hurry to take prompt action ‘fails to bring it to mind’. If she had recalled it she would have noted the ringing pattern and formed the belief that there was no fire.

We may agree that, to use Skorupski’s terms, Jane’s epistemic state “warrants the belief that it is most unlikely that there is a fire”. To say that Jane’s epistemic state warrants this belief is to say, roughly, that Jane has a sufficient subjective reason to believe that it is most unlikely that there is fire.³⁶ Hence, we may think that Jane has a sufficient subjective reason not to press the button for evacuation.

If Jane has sufficient subjective reason not to press the button for evacuation, we should grant that the rational thing to do for Jane is not to press the button, at least if ‘rational’ is understood in terms of responses conform to subjective reasons. However, intuitively, Jane should not be blamed for what she does (and believes). To convey the idea that Jane should not be blamed, it not unnatural to use the term ‘justified’ or ‘rational’. For example, Skorupski writes:

Was Jane’s actual belief, that someone was indicating the presence of fire, justified? Or could we criticize her for carelessness or absence of mind? To answer that question, we would have to examine the actual circumstances in detail. One crucial issue would be how much else she had to think about, and how important it was. Again, did her secretary remind her that morning of her earlier decision? If the circumstances were sufficiently pressured we might conclude that Jane’s belief that a fire was being indicated was perfectly ‘understandable’. She was not criticizable — in the circumstances, the memory lapse could easily have occurred to anyone. In which case her belief was, in the present sense, justified though not warranted. (2010, 126).

³⁶The notion of warranted reason is the way in which Skorupski fleshes out the notion of subjective reasons. He distinguishes the reasons there are for someone (the equivalent to what I have called ‘objective reasons’) and the reasons she has. The reasons a person has are her warranted reasons and they are distinguished from “the reasons he takes himself to have” (the equivalent to what I have called ‘motivating reasons’). See Skorupski (2010, 108).

If we say that Jane's belief and action are justified, and that what Jane ought to believe and to do is either a matter of objective reasons or subjective reasons, then justification is not a matter of normative reasons. Jane's belief and action are neither conform to her objective reasons nor to her subjective reasons. Justification thus understood seems merely to imply that Jane is not blameworthy, that what she did was understandable, or what could be expected of her. As Skorupski notes,

we need a concept of justification distinct from that of warrant, though related to it. *It will take account of what normative features in a concrete situation in which someone finds themselves they could be expected to notice.* For example, in a law court the jury may need to judge whether my beliefs on some matter were reasonable ones for me to have at the time in question. It is asking whether my beliefs were justified; its question is not settled by determining whether or not they were warranted. (2010, 111, my italics).

Thus, if what you ought to do (and believe) is determined by the (warranted) normative features of your concrete situation, and we say that you can rationally or justifiably fail to notice some of them (even if they are accessible), then what we call 'rational' or 'justified' must not be identified with what you ought to do and what is conform to normative reasons.

It is plausible to think that normative reasons and justification thus understood should come apart. Arguably, the fact that a person is rational and has a normative reason does not imply that she is not wrong about this reason. Rationality is not infallible. If this person is wrong about this reason, but not because she is irrational, however, the fact that she is rational may be something to be said in her defence, and this seems sufficient to remove the blame. Thus, it seems that according to Skorupski, justification is a way of removing the blame. A person is justified even if she makes a mistake on her (subjective) reasons to the extent that it is a mistake that any rational subject could have made in the circumstances:

One's belief, action, or feeling is justified if, in the circumstances that prevail, it opens one to no criticism in respect of rationality—whether of competence or care. (2010, 124)

In sum, Skorupski uses 'justified' (or 'rational') in relation to considerations that exculpate by showing the subject in a positive light.³⁷ However, the fact that the terms 'justified' and 'rational' are sometimes used in relation to the fact that a subject is not blamable, or has a certain kind of excuse, and sometimes in relation to what she ought to do (i.e., in relation either to her objective or to her subjective reasons) must not hide the fundamental distinction between normative reasons and considerations that remove the blame.³⁸ Hereafter, I will use 'justified' only in the normative sense.

³⁷For example, Skorupski writes that a belief can be justified because "it rests on excusable errors of reasoning" (2010, 126).

³⁸See also Dancy (2000, 6-7): "Suppose, for instance, I make a not very sensible choice about what arrangements to make about my pension. And suppose that I can later explain the choice I made by pointing out that there were

According to some views of normative reasons, Jane does not do what she ought to do. However, we might think, these views are wrong. Jane is justified (in the normative sense of the term), and not merely excused. What you ought to do is neither determined by your objective reasons nor by your subjective reasons, but by what a rational subject would do in the circumstances. And Jane does what a rational subject would do in these circumstances, and she does it in a way in which a rational subject would do it.

Thus, we must distinguish at least three levels of evaluation of action. First, there is the level of objective reasons. Second, there is the level of subjective (or warranted) reasons. Third, there is the level of what I shall call the expected reasons. The expected reasons are relative to particular circumstances. These are the considerations that we can expect a reasonable subject to treat as reasons in the circumstances in question (that is, these are the motivating reasons a rational subject would have in these circumstances). For example, in FIRE ALARM, that there is a fire is an expected reason. The fact that Jane treats this consideration as a reason to press the button does not show that Jane is not rational, although this consideration is neither an objective nor a subjective reason for Jane.

Given these three levels of evaluation, there are three possible answers to the question what level is normative. Depending on the answer, what shall count as a normative reason or as an excuse will differ.

First, suppose we think that the normative reasons are the objective reasons. Consider MINERS, and suppose that the subject blocks neither shafts. We may say that although the subject does not do what she ought to do, she has an excuse for not doing it, namely, that she does not know what she ought to do (she does not know where the minors are trapped). To the extent that she rationally does the rational thing, she is excused and should not be blamed for her action. On this view, the distinction between objective reasons and subjective reasons aligns with the distinction between normative reasons and excuses.

On this approach, a further distinction is required to distinguish between excuses having to do with subjective reasons and doing the rational thing rationally, and excuses having to do with expected reasons. To illustrate, consider MINERS*. MINERS* is similar to MINERS except that the subject must act instantaneously. She can press the button that blocks shaft A, the button that blocks shaft B, or the button that prevents shaft A and B to be blocked. If no button is pressed instantaneously, all the minors are killed. Suppose that the subject sees instantaneously that if she blocks one shaft there is a chance to save all the minors, and for this reason she presses the button that blocks shaft A. We may think that she has an excuse for not pressing button that prevents shaft A and B to be blocked. After all, it is not instantaneously obvious that the rational thing to do is

some crucial facts that I happened quite reasonably to have got wrong, and in this way, as we might put it, exculpate myself. There is a sense of 'justify' in which I can be said to have justified doing what I did. But this does not show that the balance of reasons was in favour of the action. It wasn't. Indeed, the features in the light of which I made my choice turned out all to be a mistake, and so cannot count even as defeated reasons."

to block neither shafts (assuming that the rational thing to do is to maximize expected utility), and we cannot expect a reasonable subject in this situation to be able to see this instantaneously. On the other hand, it would have been clearly irrational to do nothing. Arguably, although the subject did not act in conformity to her subjective reasons, she should not be blamed, and no criticism can be made in respect of her rationality. Still, on reflection, she would presumably have come to a different decision. FIRE ALARM is a different case illustrating this point.

Second, suppose we think that the normative reasons are the subjective reasons and that you ought to do the rational thing. In MINERS, we will say that the subject knows what she ought to do, namely, to block neither shaft. On this approach, excuses are needed only in cases where the subject does not do the rational thing (or does not do it rationally). Thus, only MINERS* and FIRE ALARM are cases where the subject is excused. On this approach, the distinction between justifications and excuses aligns with the distinction between subjective reasons and expected reasons.

Third, suppose we think that the normative reasons are the expected reasons. On this view, you ought to do what we could expect a rational subject will do in the circumstances. In MINERS, FIRE ALARM and MINERS*, we will say that the subject knows what she ought to do, respectively, to block neither shaft, to press the fire alarm button, and to press a button that blocks one shaft. In each case, the subject is not merely excused for what she does, but justified. On this approach, the distinction between justification and excuse align with the distinction between expected reasons and irresponsibility.

Given that these three approaches give different verdicts as to whether an action is justified (in the normative sense) or merely excused, we may be tempted to use the distinction between justification and excuse to argue in favour or against one or another of these approaches. For example, Baron (2005) argues against the idea that the normative reasons are the objective reasons by saying that, on this view,

it does not matter whether the actor is seen as having erred—erred in a way that deserves to be excused given the weakness of human nature or the particular disabilities of the agent—or is instead seen as having conducted himself admirably, or at least adequately. But to most of us, I think, it does indeed matter whether we are told, “We won’t blame you; we understand that it would be too much to ask you to do better than you did”; or instead, “We won’t blame you; for heaven’s sake, you conducted yourself in a way that was exemplary [or at least: perfectly reasonable].” (2005, 399).

The underlying argument seems to be the following. There are cases where you are excused although you do not do the rational thing (or you do not do what a rational subject would do). In the domain of law, for example, provocation and intoxication may be seen as excuses, and they involve that the subject was not in full control of her action, or not fully responsible for it. We should not expect a

normal subject to resist to (a certain degree of) provocation, nor to foresee all the consequences of her action when intoxicated. Insanity is a factor that removes the blame in a similar way. Now, the first option says that justification is determined by objective reasons and that when one is wrong about objective reasons, rationality is an excuse. But this approach does not distinguish between cases in which the subject is not blamable due to a lack of responsibility, and cases in which the subject is not blamable precisely because she acted as any rational subject would have acted in the circumstances. Still, there is obviously something more to be said in defence of the subject in the latter case than there is in the former. As the former case is a case of excuse and that there is more in the latter case than in the former, it seems to follow that the latter case is not a case of excuse. If so, it must be a case of justification, and hence, the normative reasons are not the objective reasons.

This line of argument supposes that the distinction between justification and excuse is exhaustive. Further, it supposes that intoxication, provocation and insanity should be classified as excuses. These issues are debated.³⁹ According to Gardner, for example, there are four types of cases in which a subject is not blameworthy:

[I]t is best of all if we commit no wrongs. If we cannot but commit wrongs, it is best if we commit them with justification. Failing justification, it is best if we have an excuse. The worst case is the one in which we must cast doubt on our own responsibility. When I say ‘best’ and ‘worst’ here I mean best and worst for us: for the course of our own lives and for our integrity as people. (Gardner, 2007, 88)

The first case is not problematic. It is a case where all the normative reasons speak in favour of the action. In the second case you commit a wrong, but a justified one. The standard example is self-defence. If someone is trying to kill you and the only way to save your life is to kill her, then killing her is justified. Killing someone is generally unlawful, but in situation of self-defence, it is lawful to kill. A wrong is justified if, all things considered, you ought to do it.

Now, a question arises in cases of mistaken self-defence. Suppose you reasonably but falsely think that someone is trying to kill you and that the only way to save your life is to kill her. Suppose you kill her as a result of these reasonable beliefs. There is no denying that, to the extent that your beliefs are reasonable, you are not blameworthy. But is it a case of justified wrong or a case of excuse? According to Gardner’s approach, this is a case of excuse. On this view, you are excused when, although you commit an unjustified wrong, you have justified (false) beliefs (or justified feelings) that rationalize your action. Thus, excuses do not show that you are not responsible for your action, or that your action is not related to reasons. In cases of excusable actions, the subject’s action is typically based on reasons that are not reasons to act, but reasons to believe that the action is right. These reasons does not show that the action is right, but that

³⁹See, e.g., Botterell (2009) for a presentation of the debate.

the subject is not blameworthy.

In contrast, in the last case, the subject is not blameworthy because she is not responsible for her action. Such cases may be seen as cases of exemption rather than cases of excuse. If the agent does not do what she ought to do but she is not responsible for that, she is not blamable. The lack of responsibility involves that the action is (partially) involuntary, which is the case in particular if the rationality of the agent is impaired. These cases can be compared to cases of immunity. A certain status (e.g., being a child, being in a diplomatic position) can make it that you are not under the authority of a certain law. You are not accountable if you act against this law.

Thus, those maintaining that you can be rational but wrong about your normative reasons can distinguish between different ways of avoiding blamability. Further, the fact that we can distinguish three types of cases in which a wrong has been committed without blamability might be used as an argument in favour of a distinction between being rational and doing what one ought to do. Indeed, we might argue, if we identify being rational and doing what one ought to do, we do not account for the fact that there is something more to be said in favour of rational actions that also conform to objective reasons, and that, at least, there is something wrong in actions that do not conform to objective reasons.⁴⁰

The question about what should count as a justification, an excuse, or an exemption, is particularly difficult to answer because it is unclear what criterium should be used to distinguish justification and excuse. We might think that, in contrast to merely excusable actions, justified actions do not require reparations or regrets, even when a wrong is committed. However, this is not the case. In cases of justified wrong, there is still a wrong that is committed. The justification does not cancel it. It just shows that it ought to be done. So, there may still be residual duties, and regrets and reparations may be appropriate.

Sometimes, it is said, it cannot be justified to resist to a justified action, and a third-party is justified in aiding someone to do something this person is justified in doing. For excused actions, things are different. It can be justified to resist to a merely excusable action, and a third-party would not be justified in aiding someone to do something this person is merely excused in doing. These criteria are controversial, though.⁴¹

My aim here is not to settle this issue. However, I have suggested that one way of understanding the notion of epistemic evaluation of action is to relate it to an assessment of subjective reasons. As we can see now, this is not the only option. Indeed, two different levels of evaluation appeal to an epistemic condition. At the level of subjective reasons, we appeal to the considerations that are accessible to a rational subject and would be accessed in ideal circumstances. At the level of expected reasons, we appeal to the considerations that would be accessed by a rational subject and treated as reasons in such or such concrete circumstances. When it is said that the subject's

⁴⁰See Littlejohn (2012, 42-45, 227) for a three-case argument along these lines.

⁴¹See Baron (2005).

epistemic position is good enough (or too bad) for acting, we may refer to either of these levels.

To illustrate, consider FIRE ALARM. In the sense of her subjective reasons, Jane's epistemic position with respect to the consideration that there is no fire is good enough. If she believed this for the right reason and acted on this consideration, she would rationally do the rational thing. However, in the sense of the expected reasons, Jane's epistemic position with respect to this consideration is not good enough. Given Jane's epistemic position with respect to this consideration and the particular circumstances, we should not expect her to treat this consideration as a reason, even if she is rational.

In sum, there are (at least) three levels of evaluation of action, and each level is a candidate for being the normative level. Depending on which level is the normative level, the distinction between justification and excuse will be drawn differently. Further, when evaluating an action from the epistemic point of view, we can focus either on the subjective reasons or on the expected reasons. If so, the notion of epistemic evaluation of action is equivocal.

7.3 The equivocity of the epistemic evaluation of action

7.3.1 Ways in which we can conceive of the epistemic evaluation of action

In the previous section, I have contemplated the notion of evaluation of action in general. I have identified three levels of evaluation and noted that two of them appeal to an epistemic condition, namely, the subject's subjective reasons and the subject's expected reasons for action. Hence, these evaluations involve an epistemic evaluation (i.e., an evaluation of whether this epistemic condition is satisfied). This is one way in which the notion of epistemic evaluation of action is equivocal. In this section, I consider more specifically the notion of epistemic evaluation of action. My aim is to show that there are other ways in which the notion of epistemic evaluation of action is equivocal.

First of all, let me stress that epistemic evaluations of action in this sense differ from purely practical evaluations. Purely practical evaluations assess whether a certain consideration (say, that it is raining) is rightly connected to a certain action (say, taking an umbrella) in order to (potentially) speak in favour of this action, in isolation from the subject's epistemic perspective. Whether there is such a connection between the consideration in question and the action may depend on the subject's desires (e.g., a desire not to be wet), or on practical norms (e.g., a prudential norm). But this practical connection is not supposed to depend on the subject's perspective. In contrast, epistemic evaluations of action are evaluations of whether a certain consideration which is rightly connected to the action (given the subject's desires) is also rightly connected to the subject's perspective for this consideration to potentially justify (or rationalise, or excuse) her action.

To understand the role of the epistemic condition(s) involved in the epistemic assessment(s) of action(s), compare with the (alleged) role of desires in reasons for action, and the related evaluation

of the subject's motivational state when we assess her actions. According to the Humean theory of reasons, for example, p is a reason for S to do A only if, given that it is the case that p , doing A satisfies a desire that S has. (Alternatively, if we want to distinguish rationalising factors and normative reasons, we could propose a Humean theory of rationality, according to which p can rationally motivate S in doing A only if, given that it is the case that p , S satisfies one of her desires by doing A .) Similarly, we may say that p is a reason for S to do A only if S satisfies a certain epistemic condition with respect to p . (If we want to distinguish rationalising factors and normative reasons, we may say that p can rationalise S 's action only if S 's satisfies a certain epistemic condition with respect to p .) In the same way as p can be a reason (or a rationalising factor) of S 's action but not of S^* 's action because S and S^* have different desires, p can be a reason (or a rationalising factor) of S 's action but not of S^* 's action because, even though they have the same desires, S and S^* have different epistemic positions with respect to p .

Of course, it is controversial whether a subject must have desires for there being reasons for her to act. But it is clear that when we evaluate the subject's action, we consider her desires. Likewise, it should be controversial whether S being in a certain epistemic position is required for there being a reason for S to act. But it is clear that when we evaluate the subject's action, we evaluate her epistemic position.

The fact that we appeal to an epistemic condition when we assess a subject's action has led most philosophers to endorse the idea that there is an epistemic norm of action underlying these evaluations. Basically, this norm is supposed to tell us when it is epistemically appropriate (or rational, or excusable) to use a certain consideration as a reason to act, or to use a proposition as a premise in practical reasoning. Philosophers generally suggest that this norm is a *sui generis* norm arising from the concept or nature of the relevant epistemic condition, or from the concept or nature of action or practical deliberation.⁴² For example, philosophers embracing pragmatic encroachment on knowledge often suggest that the (supposed) knowledge norm of practical reasoning and action arises from the concept of knowledge or that of practical deliberation.⁴³ According to those who think that this knowledge norm of action arises from the concept of knowledge, the concept of something known is the concept of something which can properly be treated as a reason (to act). Thus, as a matter of conceptual necessity, if S know that p , it is appropriate for S to treat p as a reason to act (when p is relevant for action). According to those who think that this knowledge

⁴²See Fassio (forthcoming). Thus, Stanley (2005b, 9) writes: "it is immensely plausible to take knowledge to be constitutively connected to action, in the sense that one should act only on what one knows".

⁴³See Stanley (2005b, 89): "my purpose is to establish that knowledge is conceptually connected to practical interests"; Hawthorne and Stanley (2008, 571): "our ordinary folk appraisals of the behavior of others suggest that the concept of knowledge is intimately intertwined with the rationality of action". Fantl and McGrath (2009, 84) write that their principles connecting knowledge and justification of action "articulates broader connections between the concepts of knowledge and justification than those typically recognized" and Hawthorne (2004, 30), that "we operate with a conception of deliberation according to which, if the question whether p is practically relevant, it is acceptable to use the premise that p in one's deliberations if one knows it and (at least in very many cases) unacceptable to use the premise that p in one's practical reasoning if one doesn't know it".

norm arises from the concept of practical deliberation, the concept of practical deliberation is the concept of a deliberation which is appropriate only if the premises are known.

The idea that there is such an epistemic norm of action can be challenged. For example, Fassio (forthcoming) argues that appealing to the existence of an epistemic norm of action in this sense is not the best way of explaining our epistemic evaluations of action. According to Fassio, these epistemic evaluations do not bear on the satisfaction condition of an epistemic norm, but on the regulation conditions of a non-epistemic norm according to which it is appropriate to treat p as a reason if and only if p is a reason (where reasons are conceived as facts or true propositions independent from the subject's epistemic perspective). These epistemic assessments would be instrumental evaluations regarding how the agent can conform to this non-epistemic norm in the actual and relevantly similar circumstances.

Further, if we grant that there is an epistemic norm of action, we should ask whether this norm is objective or subjective. That is, we should ask whether this norm, and the associated reasons, depends on the subject's perspective. Depending on how this question is answered, this can affect what can plausibly be said to be the epistemic norm of action, and even whether it is plausible to say that there is such an epistemic norm in the first place. For example, suppose that we say that knowledge is the epistemic norm of action. Since this norm is supposed to be objective and independent from the subject's perspective, this is supposed to be true even if a particular agent does not know that knowledge is the epistemic norm of action, or even if she does not believe that knowledge is the norm of action. Now, if we grant this, then it is unclear why we do not grant that non-epistemic norms of action, and the associated reasons, are objective too. But if these norms and the associated reasons are objective, then it is unclear why there would be an epistemic norm of action relating the subject's practical reasons to the subject's perspective in the first place. On the contrary, suppose we say that the epistemic norm of action is subjective. For example, suppose we say that it depends on what the subject believes the epistemic norm of action to be. If so, it is unclear why we cannot say that the non-epistemic norms of action and the associated reasons are subjective too. But if they are subjective and depend on the subject's perspective, then it becomes unclear what further role this (subjective) epistemic norm of action is supposed to play.⁴⁴

My aim is to look at the argument from the normative connection between knowledge and action to see whether it supports the rejection of the orthodox account of knowledge. This argument often uses the idea that there is an epistemic norm of action. But this is not required. The argument just needs the claim that there is a unique epistemic condition on appropriate (or rational, or excusable) action, whether or not this epistemic condition is provided by an epistemic norm specific to action and practical reasoning, or merely derived from a non-epistemic norm in some other way. If practical factors are relevant to whether we satisfy this unique epistemic condition, and this epistemic condition is identified with knowledge, there is an argument to the effect that we should

⁴⁴See 9.3.1 for a related discussion.

adopt pragmatic encroachment about knowledge. If this is correct, in order to defend the orthodox account of knowledge, we need not decide whether the epistemic condition which is assessed in epistemic evaluations of actions is issued by an epistemic norm specific to action and practical reasoning.

Thus, for example, Hawthorne and Stanley (2008) claim that there is an epistemic norm of action, because they think that knowledge is a necessary epistemic condition on personal reasons. (What they mean by ‘personal reasons’ is what I have called (rationally) good motivating reasons; these are subjective reasons that motivate you.) They add that it is appropriate for you to treat *p* a reason to act only if *p* is a personal reason to act. (What they mean by ‘appropriate’ is rational permissibility.) They conclude that there is a norm according to which it is appropriate for you to treat *p* as a reason to act only you know that *p*.

On this view, in FIRE ALARM, it is inappropriate for Jane to be motivated by the consideration that there is a fire. Indeed, this is not a personal reason (a consideration that Jane knows). However, this view does not imply that Jane should be blamed or that she does not meet a standard of excusability. Indeed, she treats as a reason something we could expect a rational subject with this epistemic position to treat as a reason in these circumstances. Thus, she inappropriately, but excusably, treat that there is a fire as a reason to act.

However, in their argument in support of pragmatic encroachment, Fantl and McGrath do not directly appeal to an epistemic norm of action connecting an epistemic condition and the appropriateness of treating something as a reason.⁴⁵ They endorse the view that if you know that *p*, then *p* is a reason you have to act (in their terminology, if you know that *p*, *p* is warranted enough to justify you). However, they do not embrace the further claim that it is appropriate for you to treat *p* as a reason to act if and only if *p* is a reason you have. They suggest that if *p* is a reason you have, then you are (normatively) justified in treating it as such. But they deny that if you are (normatively) justified in treating *p* as a reason to act, then *p* is a reason you have to act. (Indeed, against Hawthorne and Stanley’s claim that it is appropriate to treat something as a reason only if it is a (personal) reason, and assuming that (personal) reasons are considerations that the subject knows, they argue that someone could have a justified but false belief that *p*, and hence, could be (normatively) justified in treating *p* as a reason even if *p* is not a reason.) Still, their argument in favour of pragmatic encroachment does not require the existence of such an epistemic norm on action and practical reasoning. It just requires that knowledge is sufficient to turn a reason there is into a reason one has.

In sum, it is controversial whether (and in what sense) there is an epistemic norm of action. But since this issue does not really affect the argument from the normative connection between knowledge and action advanced against the orthodox account of knowledge – which, on this score, merely requires the uncontentious claim that there are epistemic evaluations of actions pointing to

⁴⁵See Fantl and McGrath (2009, 124-125).

necessary and/or sufficient epistemic condition(s) for justified (or rational, or excused) action – I set it aside. For sake of simplicity, however, and given that the great majority of epistemologists in this debate grant that there is an epistemic norm of action, I shall call the relevant epistemic condition ‘the epistemic norm of action’.

I have argued above that epistemic evaluations of action are equivocal because they can invoke an epistemic condition necessary and/or sufficient for either

- (1) subjective reasons or expected reasons.

I will now show that these epistemic evaluations can also be equivocal between the evaluations of

- (2) unreflective rationality and reflective rationality,

and between the evaluations of

- (3) possibilities decisive for action, subjective reasons to act, and considerations warranted enough to justify the action.

Drawing these distinctions is crucial to understand and assess the different competing accounts of the epistemic evaluation of action, and their potential in compromising the orthodox account of knowledge.

To illustrate these distinctions, consider the following case:

BAD SURGEON John is a surgeon. He wants to operate a patient and he needs a scalpel. There is a scalpel on the table, and it is safe. John has no reason to think that it is safe. However, John takes it and starts operating.

Intuitively, there is something wrong in what John is doing, even if the scalpel is in fact safe. The problem is that, in some sense or other, John is not epistemically well enough positioned with respect to the fact that the scalpel is safe.

There are different ways of formulating what is wrong, though. We may want to say that the fact that the scalpel is safe is not a reason John has (that the scalpel is safe is not a subjective reason for John); that a rational subject in these circumstances will not use the scalpel without checking whether it is safe (that the scalpel is safe is not an expected reason for John); that if John is motivated by the unconscious thought that the scalpel is safe, this motivating reason is (rationally) bad (John is unreflectively irrational); that if John deliberates and concludes that he has a sufficient reason to take this scalpel – namely, that the scalpel is safe –, his deliberation is not rational (John is reflectively irrational); that John must not act as if the scalpel is safe; and that the fact that the scalpel is safe cannot justify John’s action.

A basic problem is that although all these claims are true in this case, they do not come to exactly the same thing. I have already explained the distinction between subjective and expected reasons presented in (1). I shall now look at the distinctions presented in (2) and (3).

7.3.2 Unreflective and reflective rationality

Some epistemic evaluations are focused on the subject's mental attitudes, or her mental actions. The expression that philosophers use to capture the mental attitude (or mental action) that is supposed to be epistemically evaluated diverge. Some philosophers use the expression 'treat *p* as a reason for acting' (where '*p*' stands for the relevant consideration).⁴⁶ On this formulation, for example, the norm says that it is not appropriate for John to treat the proposition that the scalpel is safe as a reason unless John satisfies the relevant epistemic condition with respect to this consideration. The following expressions are also used: 'act for the reason that *p*' or 'act on *p*' , 'use *p* as a premise in practical reasoning' , 'rely on *p* in practical reasoning and action' , and 'use (the belief that) *p* as a premise in practical reasoning or (the belief that) *p* as a reason for acting'.⁴⁷

It not always clear whether and when these differences in formulation matter. However, in general, it is important to distinguish at least two different types of attitudes (or mental actions) that we might want to evaluate (from the epistemic point of view). On the one hand, there are active, explicit, controlled and deliberate reasonings and responses to reasons. On the other hand, there are passive, implicit and involuntary directed reasonings and responses to reasons. This distinction is a familiar one.⁴⁸ To respond to reasons, or to reason, is to form, revise or retain an attitude (belief, intention, emotion) by being guided by a certain content. This content inclines the rational subject toward or against the relevant attitude. However, the subject may or may not express (to herself) the attitude she has toward that content (and may or may not cause herself to acquire, retain or lose the relevant attitude). Or she may or may not be reflectively aware of the attitudes and the reasons for them. In sum, the response or the reasoning may be either active and reflective or passive and unreflective.

There are different ways of understanding this distinction. Neta (2009), for example, distinguishes, on the one hand, 'act for the reason that *p*' , and on the other hand, 'use *p* as a premise in practical reasoning' and 'treat *p* as a reason' . As he notes, while a piping plover may engage in broken wing display behavior "for the reason that" a predator is approaching her young, it may lack the "powers of reasoning" required to "use this as a premise in a practical reasoning" (Neta, 2009, 684-685). Presumably, Neta has active and reflective practical reasoning in mind. Further, a piping plover does not possess the concept of reason. Neta takes it that 'treating something as a reason' involves to exercise powers of metacognitive representation and the deployment of the concept of reason (Neta, 2009, 685-686). Thus, a piping plover cannot treat something as a reason, in this sense.

In a similar way, Skorupski (2010) distinguishes between 'act for the reason that *p*' and 'act

⁴⁶See Hawthorne and Stanley (2008).

⁴⁷See Hawthorne and Stanley (2008), Hawthorne (2004), Brown (2008), ?.

⁴⁸See, e.g., Owens (2000), Kolodny (2005, 520), Frankish (2004, 24-33), Skorupski (2010, 57-59, 462-464), Arpaly and Schroeder (2012).

from / out of the reason that p . Skorupski uses the latter formulation to express the idea that the subject has “the capacity for full-blooded reason-responsiveness – the capacity to reflectively identify and audit reasons, at a level that does require possession of the concept of reason, of itself, and of the present.” (Skorupski, 2010, 59)⁴⁹

Some philosophers do not seem sensitive to the difference between passive and active responses to reasons. Perhaps they think that being motivated by p or acting for the reason that p does not require to judge that p is a (good) reason to act. Hawthorne and Stanley (2008), for example, identify the expressions ‘treat p as a reason’ and ‘act for the reason that p ’, and they do not seem to think that acting for the reason that p requires to judge that p is a reason.⁵⁰

The idea that acting for a reason does not require to judge that p is a reason is in line with the traditional psychological approach of rational motivation that goes back to Davidson (1963). According to this approach, to say that a subject acts for a reason amounts to saying that she has a belief-desire pair that causes the action in a non-deviant way. Roughly, the desire is a pro-attitude toward a certain kind of action having a certain property, and the content of the belief is that a particular action (seen under a certain description) is a way of doing the kind of action that is desired. For example, when we say that the piping plover engages in broken wing display behavior for the reason that a predator is approaching her young, we say that she has a desire, say, to protect her young, and a certain belief, say, that engaging in broken wing display behaviour is way of protecting her young, and these attitudes cause the action in the right way.

This standard version of the psychological model of motivation is controversial, and my aim is not to defend it here.⁵¹ What I want to suggest, rather, is that it is clear that sometimes we respond directly or spontaneously to our reasons (like the piping plover), and sometimes the response is deliberate (like when we consciously weigh the pro’s and con’s before acting). Thus, even if we grant that treating p as a reason to act does not require deliberation or reflection on whether p is a (good) reason, we still need to account for cases in which the response is deliberate.

We might reject the distinction between active and passive reasonings and responses to reasons on the basis of the claim that to reason and to respond to reasons always require some sort of deliberative process. Thus, on one prominent view of rational motivation, to be motivated by p , or to act for the reason that p , is to judge that p is a reason. There are two versions of this deliberative or reflective model of motivation. The most common is the normative conception of motivation, according to which for S to be motivated by p is for S to judge that p is a normative reason for S

⁴⁹See also Skorupski (2010, 21, 57-59).

⁵⁰See ?, 535n7.

⁵¹One problem that has been raised is that this model does not capture the idea that when you act intentionally (i.e., for a reason) you know (without observation) what you are doing (see Anscombe (2000, 13-15)). Indeed, a belief-desire pair can cause an action (in a non-deviant way) without the subject knowing that she is doing it (see Velleman (2000, 2-3); Setiya (2007)). Following Davidson, we may reply that the idea that acting intentionally requires to know what one is doing is too strong: “A man may be making ten carbon copies as he writes, and this may be intentional; yet he may not know that he is; all he knows is that he is trying” (1980, 50). See Setiya (2007) for further discussion.

to φ .⁵² This view amounts to the claim that acting for a reason is to act under the guise of the good. Those who reject this claim may endorse the second version, according to which for S to be motivated by p is for S to judge that p is S's motivating reason to φ . If we do not always take our motivating reasons to be normative reasons (i.e., considerations that show that our action is right), this latter view allows the possibility of actions done for reasons taken to be bad by the subject.⁵³

Alternatively, we might think that there is a distinction to be made between active and passive responses and reasonings, but that it is merely a difference of degree. Thus, it is often objected to deliberative or reflective accounts of motivation that animals can act for a reason although they do not possess the concept of reason. Still, if we grant that we can possess concepts by degree, we might think that animals can act for reasons to some degree, to the extent that they possess the concept of reason, self and action, to a certain degree.⁵⁴

However, there are reasons to think that the distinction between passive and active reasonings and responses to reason is genuine, and that it is not merely a difference of degree. I will put forth three arguments.

First, there are cases where a subject can be seen as acting for a reason although she does not judge that she has a reason. As Arpaly notes,

[i]t is not a provocative view that an accomplished tennis player, for example, does not have time to deliberate on all her moves during a fast-paced game. Not only that, but given the complex factors to which she responds, she is unlikely to be able to reconstruct all her reasons for action after the game. However, even after the ball is served, we can legitimately judge her moves as rational ("That was brilliant!") or criticize her for irrationality ("What on earth were you thinking there?"). (Arpaly, 2003, 51-52).

Further, there are cases where a subject seems to act for a (good) reason although on reflection she judges that she does not have such a (good) reason:

STUDENT. I am convinced that I ought to be a lawyer, but only because I was pressured into it by my parents. As I go through law school, I truly believe that I am suited to this kind of work. I do not respond to the clues that indicate otherwise: the fact that I spend much less time working than my peers; that I often feel tired and lethargic; that I never get very good grades. I would never act on these reasons, as grounds on which to quit. Still I might decide to quit, and be moved to do so, unconsciously, by beliefs that correspond to these facts—finding my own decision both capricious and hard to explain. (Setiya, 2007, 35)⁵⁵

⁵²See, for example, Scanlon (1998); Velleman (2000); Korsgaard (2008); Alvarez (2010).

⁵³See Setiya (2007).

⁵⁴See Skorupski (2010).

⁵⁵See also Arpaly (2003) for cases of inverse akrasia, in which the subject do the right thing for the right reason against her best judgement.

Cases in which a subject can be seen as acting for a reason while not judging that she has a reason, or while judging that she has no (good) reason, are controversial. Some philosophers think that these cases feature non-intentional actions or mere activity.⁵⁶ Let me just note that such cases are cases in which the subject's action is guided by a certain content on the basis of which we can provide a rationalising explanation of the action. In this respect, they are very different from cases like moving one's leg to and fro. Thus, it is at least *prima facie* plausible that these cases are genuine cases of acting for a reason, and not cases of mere activity. If so, as they clearly are not cases of acting out of a reason (or from a reason), they strongly suggest that there is a genuine distinction between two ways of being motivated: either by a conscious and deliberate judgement on one's reasons or directly by grasping a certain content.

Second, as Arpaly and Schroeder (2012) argue, the distinction between acting for a reason in the non-deliberative and non-voluntary sense and acting for a reason "par excellence" (i.e., in the deliberative sense) is required if we want to make sense of the idea that deliberations (theoretical and practical) can be rational or irrational, and hence, are performed for reasons. A deliberation is a kind of voluntary mental action the constitutive aim of which is to determine what to believe (theoretical deliberation) or what to do (practical deliberation). Sometimes it is rational to deliberate, and sometimes it is not rational to deliberate. For example, arguably, it would not be rational for Jane to deliberate in FIRE ALARM. Now, suppose that acting for a reason always requires some kind of deliberation on the reason for which one acts. Thus, when you deliberate for a reason, this requires a further deliberation on your reason to deliberate. Call this further act of deliberation the second-order deliberation. The second-order deliberation can also be rational or irrational. If so, there is a reason for which the second-order deliberation is undertaken. This requires a further deliberation. And so on. If we assume that acting for a reason requires an act of deliberation on the reason for which one acts, a vicious regress looms. Only an infinite mind could perform all the required deliberations. And as deliberations are actions, they can occur only if there is a reason for which they are made.

Third, the distinction between passive and active responses and reasonings is basic in many approaches in psychology, philosophy of mind and epistemology. In psychology, dual-process theories distinguish, on the one hand, a fast and automatic reasoning process based on heuristics, and, on the other hand, a conscious reasoning process operating slowly and rule-based.⁵⁷ Those who reject dual-process approaches often maintain that there is still a distinction to be drawn.⁵⁸ In philosophy of mind, it has been argued that we should distinguish two different psychological levels involving two different kinds of acts or attitudes.⁵⁹ In epistemology, some have argued that we

⁵⁶See, e.g., Setiya (2007, 35).

⁵⁷There are various dual-process theories. A plausible criterium for deliberative thinking is the activation of the working memory. See Gerken (forthcoming) for discussion.

⁵⁸See, e.g., Carruthers (2014).

⁵⁹According to Frankish (2004), for example, we must distinguish a basic mind and a supermind. While the basic mind corresponds to unconscious (partial) beliefs and (partial) desires responsive to Bayesian norms of rationality,

should distinguish two different levels of responses to reasons.⁶⁰

In sum, it is very plausible that there is a genuine distinction between, on the one hand, passive responses to reasons and reasonings, and on the other hand, active responses to reasons and reasonings. The reflective or deliberative model of motivation is adequate for some cases, but it is too strong if we want to account for all the cases of motivation by reasons.

Once we grant the distinction between acting for a reason and acting from a reason, a natural question arises. Is there only one epistemic norm? While some philosophers are explicitly interested in formulating a norm regarding active and reflective responses to reasons (e.g., Neta 2009), other philosophers seem to be implicitly more interested in formulating a norm regarding passive and non-reflective responses and reasonings. If there are two different epistemic norms, these philosophers may talk past each others.⁶¹

7.3.3 Possibilities decisive for action, subjective reasons to act, and considerations warranted enough to justify the action

Other ways of understanding the notion of epistemic evaluation of action focus on the action viewed in isolation from the agent's attitudes. For example, Fantl and McGrath (2002; 2007) use the expression 'act as if p'. On this formulation, we may say that John the surgeon should not act as if the scalpel is safe. By 'acting as if p', Fantl and McGrath means doing the thing that it is rational to do given p.⁶² On this approach, the epistemic evaluation of action has nothing to do with the evaluation of the subject's motivating reasons. When we evaluate whether it is rational for S to act as if p, given her epistemic position with respect to p, we do not consider what motivates S in acting. Indeed, a subject S can act as if p, that is, do the rational thing to do to do given p, even if it is not the consideration that p, but the consideration that q, that motivates her.

Further, we should note that, on this approach, the epistemic evaluation of action has nothing to do with the evaluation of the agent's subjective reasons. Recall that subjective reasons are

the supermind corresponds to conscious (flat-out) beliefs and desires responsive to norms of classical reasoning (e.g., practical syllogism).

⁶⁰For instance, Owens (2000) distinguishes between a subterranean and a reflective level. Regarding theoretical rationality, he distinguishes the two levels on the basis of the role played by pragmatic factors. According to him, while awareness of the cost of error can be active at the subterranean level in rationally influencing the formation of the belief, on reflection such a factor cannot move one to form a belief. Owens concludes that we do not have reflective control over our beliefs.

⁶¹Thus, against Hawthorne and Stanley (2008), Neta (2009) argues in favour of a justified-belief-that-one-knows norm of treating something as a reason for action. Gerken (2011, 535n7) objects to Neta that he changes the topic, since Neta's account concerns S's treating p as a reason in the higher-order sense of S's conceiving herself as having p as a reason, while Hawthorne and Stanley are concerned with treating p as a reason in the first-order sense of using or relying on (the belief that) p. This objection from Gerken implicitly assumes that the epistemic norm for treating something as a reason in the first-order sense is different from the epistemic norm for treating something as a reason in the second-order sense.

⁶²Other philosophers (e.g., Ganson (2008)) use the expression 'act as if p' slightly differently, as an equivalent to something like 'doing what one would do if one acted on the assumption that p'. See above, chapter 5 section 4.3 problem 4 footnote 27.

considerations that can rationally motivate a subject. But there are cases where it is true that, given the subject's epistemic position with respect to p , this subject should act as if p , although false that this subject can be rationally motivated by p . To illustrate, suppose that it is slightly probable that it will rain, and if it rains and I have no umbrella, I will be seriously ill. Taking an umbrella is not cumbersome. Arguably, I should act as if it is raining and take my umbrella. But, of course, if I am rational, my reason for taking my umbrella is not that it will rain. That it will rain is only slightly probable. Rather, my reason is that there is a chance that it will rain. Thus, that it will rain is not a subjective reason for me, even if I should act as if it rains.⁶³

Therefore, we must distinguish the epistemic condition that must be satisfied with respect to p for the possibility that p to determine the rational thing to do (acting as if p *versus* acting as if not- p), and the epistemic condition that must be satisfied with respect to p for p to be a subjective reasons. If you satisfy the epistemic condition with respect to p such that the rational thing to do for you is to act as if p , you satisfy the epistemic condition for the possibility that p to decisively weigh in your decision. In such a case, I shall say that p is a decisive possibility for action.

Fantl and McGrath (2009) are more focused on subjective reasons. They use the expression 'p is warranted enough'. With this terminology, we may say that, in BAD SURGEON, the fact or proposition that the scalpel is safe is not warranted enough to justify John's action. According to Fantl and McGrath, if p is warranted enough to be a reason S has to φ (i.e., warranted enough to be a subjective reason), then p is warranted enough to justify S in φ -ing, for any φ (this principle they call 'Safe Reason'). The basic idea is that the epistemic condition on warranted (or subjective) pro tanto reasons is the same as the epistemic condition on warranted (or subjective) sufficient or pro toto reasons. Indeed, it is plausible that whether or not something is a warranted reasons (a subjective reason) does not depend on whether it is a pro tanto or a pro toto reason.

Still, we might think that the way in which they construe the notion of being warranted enough to justify is problematic, if this notion is supposed to capture the notion of subjective reason. According to Fantl and McGrath, the expression 'p is warranted enough to justify S in φ -ing, for any φ ' means that no weaknesses in S 's epistemic position stand in the way of p justifying S in φ -ing, for any φ (see Fantl and McGrath 2009, 66). However, intuitively, there are cases where p is a reason one has (i.e. p is a subjective reason) although a weakness in one's epistemic position with respect to p prevents p from justifying one.

Suppose that I kill you if you sentence Bill when it is not absolutely certain to you that Bill is guilty. Suppose that, in fact, Bill is guilty. You acquire more and more evidence that Bill is guilty. Arguably, at some point before absolute certainty, that Bill is guilty is sufficiently supported by your evidence for it to be a reason you have to φ , for any φ . In particular, at this point, it may be a subjective reason you have to sentence him. Still, a weakness in your epistemic position prevents you from being justified in sentencing Bill, namely, that it is not absolutely certain that

⁶³See also chapter 3 section 4.2.

Bill is guilty (your evidence does not logically entail that Bill is guilty) and if you sentence Bill in this situation, you will be killed. However, the fact that this weakness in your epistemic position prevents the (true) proposition that Bill is guilty from justifying you in sentencing Bill has nothing to do with the question whether this proposition is a reason you have to φ . Rather, it has to do with the relation between this reason you have and what it justifies. To be justified in sentencing Bill, you must have a good reason to do it: that Bill is guilty, and, given the background conditions, that this is absolutely certain. Therefore, there are epistemic weaknesses that do not stand in the way of your having something as a reason, but which may stand in the way of this reason justifying you.

In sum, although it is plausible that the epistemic condition on subjective reasons is not sensitive to whether or not this reason is *pro tanto* or *pro toto*, we should reject the claim that if p is a subjective reason for S , then no weakness in S 's epistemic position stands in the way of p justifying S in φ -ing, for any φ .⁶⁴ We must distinguish the epistemic condition that S must satisfy with respect to p for p to be a subjective reason for S , and the epistemic condition that S must satisfy with respect to p for there being no weakness in S 's epistemic position standing in the way of p justifying S 's φ -ing.

7.4 The main accounts

Everyone agrees that we evaluate actions epistemically. However, since the notion of epistemic evaluation of action is equivocal, the epistemic condition we evaluate may concern different things. This point is crucial when it comes to assessing the argument against the orthodox account of knowledge from a supposed normative connection between knowledge and action. As we shall see, (some of) the distinctions drawn above can be exploited in defence of some putative epistemic norm. In particular, chapter 8 will show how the distinction between justification and excuse can be used to defend the necessity direction of the knowledge norm of action, and chapter 9 will consider defences of the sufficiency direction of the knowledge norm based on the distinction between subjective reasons and considerations warranted enough to justify action.

It does not seem particularly controversial that such distinctions should be drawn, although there is some disagreement on how exactly to draw them. The main point of disagreement rather concerns what the fundamental epistemic condition of our epistemic assessments is, and whether this condition is knowledge, or something else. Other points of disagreement concern the character of the (supposed) norm, its source, whether the epistemic condition of application of the norm is context-sensitive, and, finally, whether we should be monist or pluralist about the epistemic norm of action. I will now look at these points and present the main theoretical options and the accounts which have been defended in the literature so far. The remaining chapters will be devoted to assess

⁶⁴See also Ichikawa (2012, 51) on this point and chapter 9 section 2.2.

these options.

7.4.1 The character and the source of the epistemic norm of action

It is sometimes said that, typically, philosophers think of the epistemic norm of action in terms of an epistemic permission to act.⁶⁵ What does that mean exactly?

First, this claim concerns the character of the supposed norm. The norm is supposed to state an epistemic *permission* in contrast to an epistemic *obligation*. Indeed, many philosophers think that it would be too strong to require of a subject that she treats all the relevant considerations with respect to which she satisfies the supposed epistemic condition as her reasons to act.⁶⁶ However, if this is correct, we could say that you can satisfy the relevant epistemic condition with respect to p, although you may ignore p in your decision without violating the epistemic norm. This is counterintuitive. For example, suppose that a surgeon satisfies the epistemic condition with respect to the consideration that operating quickly is the best thing to do, but she does not do it. She may be criticised on the basis of the fact that she ignores something in her decision that she should not ignore, given her epistemic position. It seems that there are epistemic obligations to treat propositions as reasons to act.

Perhaps the worry according to which it would be too strong to require of a subject that she treats all the relevant considerations as reasons can be attenuated if we think that a subject can treat something as a reason without making an explicit and conscious judgement. Further, we might think that an epistemic obligation is merely *prima facie* and can be cancelled. For example, if the agent do not have the capacities to treat a certain consideration as her reason, the epistemic norm according to which she ought to treat it as a reason may be cancelled.

Second, the claim according to which the epistemic norm action states an epistemic permission also concerns the source of the norm. This claim means that this norm states an *epistemic* permission in contrast to an *overall* (or practical) permission to act, where overall permission requires that there is no decisive interdiction issued by some other (non-epistemic) norm. Thus, you may satisfy the epistemic norm with respect to p for treating p as your reason although other considerations sufficiently and decisively speak against treating p as your reason to act.

For example, suppose that a surgeon satisfies the epistemic condition with respect to the fact that a certain scalpel is safe, but she knows that an evil demon will kill her if she treats the proposition that this scalpel is safe as a reason. It seems that, overall, she ought not treat this proposition as a reason, even if, from an epistemic point of view, this is permitted.

⁶⁵See Benton (2014).

⁶⁶Thus, Hawthorne and Stanley (2008) write: “Our principle concerns what it is *appropriate* to treat as a reason for action, rather than what one *ought* to treat as a reason for action. It would be overly demanding to require someone to treat *all* their relevant knowledge as reasons for each action undertaken. The principle is therefore a claim about what is permissible to treat as reasons for action in a given choice situation”.

This suggests that the supposed epistemic norm is *sui generis* and has its own specific normative source.⁶⁷ As said above, however, it remains controversial whether there is such a *sui generis* epistemic norm governing action and practical reasoning.

My purpose is to investigate the argument from the normative connection between knowledge and action in order to see whether it succeeds in compromising the orthodox account of knowledge. This issue depends neither on whether the norm is formulated in terms of permission or obligation, nor on whether there is such a *sui generis* norm, but on whether there is a unique necessary and/or sufficient epistemic condition underlying our epistemic evaluations of action, and which can be identified with knowledge. If there is such an epistemic condition, that this epistemic condition is also context-sensitive, and that it should be identified with knowledge, then we should abandon the orthodox account of knowledge. As a result, we can leave the questions concerning the character and the source of the norm aside.

7.4.2 The epistemic condition of the norm, context-sensitivity and pluralism.

A more important point of disagreement for our purposes concerns the epistemic condition of the norm. According to a popular view, this condition is knowledge. Bearing in mind all the different ways in which the epistemic norm may be conceived, we can formulate this proposal in the following way:

Knowledge norm: It is appropriate for you to treat p as a reason for acting if and only if you know that p.

Philosophers endorsing the knowledge norm (while sometimes using a different norm-content and character) include Hawthorne (2004), Stanley (2005b), Hawthorne and Stanley (2008), Williamson (2005) and Engel (2009).

An alternative to the knowledge norm of action is the justification to believe norm:

Justification to believe norm: It is appropriate for you to treat p as a reason for acting if and only if you are justified in believing that p.

Some philosophers defend a justification to believe norm with a non-orthodox view of justification. For instance, Littlejohn (2012) argues that justification is factive, Sutton (2007) that justification requires knowledge, and Fantl and McGrath (2009) that justification to believe is determined by practical factors. Related views argue in favour of a justification to believe that one knows norm (Neta, 2009), or a justification to believe that one is in a position to know norm (Smithies, 2012).

The question whether knowledge, or justification to believe, or something else, is the epistemic norm of action is certainly the most debated. Indeed, as said above, the claim that knowledge, or

⁶⁷See Crisp (2005) and Fassio (forthcoming).

justification to believe, is the epistemic norm of action plays a crucial role when it comes to argue against the orthodox account of knowledge.

Indeed, as many have noted, it seems that how good your epistemic position must be with respect to *p* for being rational in treating *p* as your reason for acting depends on your practical situation. If knowledge, or justification to believe, is the epistemic norm of action, then knowledge, or justification to believe, depends on practical factors.

To illustrate, consider the bank cases. In the low stakes case, it is rationally permissible for Keith to come back without checking whether the bank is open, but in the high stakes case, it is not rationally permissible for Keith to come back without checking whether the bank is open. This suggests that Keith meets the epistemic norm of action with respect to the fact that the bank is open in the low stakes case but not in the high stakes case. This may lead to think that the epistemic condition of application of the epistemic norm of action is sensitive to the subject's practical context. If one also thinks that the epistemic norm of action is knowledge, or justification to believe, we are led to pragmatic encroachment on knowledge or justification: whether or not you have a justified belief, or know, varies with your practical situation.⁶⁸ On this view, Keith knows that the bank is open in the low stakes case, but in the high stakes bank case, Keith does not know that the bank is open.

Following Brown (2008, 184), we can state the argument in the following way:

KNOWLEDGE-ACTION ARGUMENT

- 1) Knowledge (/ justification) norm of action: it is appropriate to rely on *p* in practical reasoning iff you know (/ are justified in believing) that *p*.
- 2) In Low but not High, it is appropriate for Keith to rely on the fact that the bank is open on Saturday in practical reasoning.
- 3) In Low but not High, Keith knows (/ is justified in believing) that the bank is open on Saturday.
- 4) Low and High differ only in the stakes.
- 5) Whether *S* knows (/ is justified in believing) that *p* partly depends on the stakes.

There are at least four possible replies to this line of argument against the orthodox account of knowledge.

First, many orthodox philosophers have rejected 1) by denying that there is a normative connection between knowledge and action. In general, these philosophers maintain that the relevant epistemic condition is sensitive to the subject's practical context and they propose an epistemic norm of action in which the subject's practical situation determines, in some way or other, whether

⁶⁸See Fantl and McGrath (2009) for an approach in term of justification and Stanley (2005b) for an approach in term of knowledge.

the warrant she has for believing p is sufficient for acting on p , where this sufficiency is neither necessarily connected to knowledge nor justification to believe:

Warrant account: It is appropriate for you to treat p as a reason for acting if and only if the warrant you have for believing p meets the epistemic standards required by your practical context.

Philosophers defending views along these lines include Gerken (2011) and Brown (2008).⁶⁹

In the same vein, Locke (2015) argues in favour of a norm of practical certainty, where S is practically certain that p if and only if the difference between p 's actual epistemic probability for S and p 's epistemic certainty for S does not make any practical difference for S . On this approach, whether or not the difference between p 's actual epistemic probability for S and p 's epistemic certainty for S makes a practical difference for S is a function of S 's practical situation.

On this type of view, only the epistemic condition that the epistemic norm requires to satisfy varies with the practical context, not the epistemic norm itself. The basic idea is that the epistemic condition of application of the norm include different epistemic standards associated with different practical situations.

We may represent the epistemic condition of application of such a norm as a set of exhaustive and exclusive disjuncts, each one constituted by the conjunction of an epistemic standard and a practical situation. The norm can thus be formulated as follows (where ' s ' is the subject, ' ϕp ' is treating p as a reason, ' PS ' is a practical situation and ' Ep ' is an epistemic position with respect to p):

$Os (\phi p, s \text{ iff } ((E1p, s \ \& \ PS1s) \vee (E2p, s \ \& \ PS2s), \dots, (Enp, s \ \& \ PSn, s))).$

For example, an evidentialist version of this kind of views would have it that the quantity (or quality) of evidence required to satisfy the epistemic condition required by the epistemic norm varies with features of the subject's practical context. On these views, there is a practical element in the scope of the epistemic norm, since the epistemic condition provided by the epistemic norm is determined by practical factors. I shall consider this kind of views in chapter 8.

Second, in reply to KNOWLEDGE-ACTION ARGUMENT, some philosophers grant that there is a variation in the epistemic evaluations of actions which depends on practical factors, but they deny that practical factors affect the epistemic condition required by the epistemic norm. That is, they deny either 2) and 3), or 4). When they deny 2) and 3), they often argue that these shifty epistemic evaluations do not bear on the satisfaction condition of the epistemic norm. According to them, it is epistemically appropriate for high stakes Keith to treat the proposition that the bank

⁶⁹See Gerken (2011, 530): “[I]n the deliberative context, DC , S meets the epistemic conditions on rational use of (her belief that) p as a premise in practical reasoning or of (her belief that) p as a reason for acting (if and) only if S is warranted in believing that p to a degree that is adequate relative to DC ”.

is open as a reason to act, since he knows it, but this is not the only thing we assess when we assess Keith's action.

To give an idea of how this strategy is supposed to work, consider the bank cases. Following Williamson (2005), we may argue that when the cost of error is low regarding the consideration that the bank is open on Saturdays, this consideration may be a (sufficient) reason to act. Therefore, granting that the epistemic norm for treating something as a reason is knowledge, if Keith knows that the bank is open and come back for this reason, she is doing well. However, when the cost of being wrong is higher, it may be that this consideration is not a (sufficient) reason to act, or that it cannot justify this action. Given the high stakes situation, perhaps Keith's reason should be that he knows that the bank is open. If so, Keith does well in the high stakes case only if she acts on the consideration that he knows that the bank is open, and granting that the epistemic norm for treating something as a reason is knowledge, it results that Keith must know that he knows that the bank is open to come back. If Keith does not know that he knows that the bank is open in the high stakes case, this may explain why it is not rational for him to come back without checking whether the bank is open.

More generally, the idea is that in high stakes cases where the subject acts on a proposition p with mere knowledge that p , and without knowing that she knows that p , either she violates a practical norm requiring to act for the sufficient reason *I knows that p*, rather than for the insufficient reason p , or she satisfies the practical norm and acts for the sufficient reason *I know that p*, but she violates the epistemic norm requiring to know the reason for which she acts, which, in this situation, is *I know that p*. Thus, if in the high stakes cases the subject acts with mere knowledge that p and her epistemic position with respect to p is deemed insufficient, the problem need not be that p is inappropriately treated as a reason from an epistemic point of view. After all, the subject knows that p . The problem may be just that p is not a good premise on which to act from a practical point of view.

On this approach, even though the epistemic condition that a subject must satisfy with respect to p may vary with practical factors, this variation does not concern the epistemic norm of action. This variation concerns what considerations are practically, rather than epistemically, appropriate when it comes to justify an action given such or such practical situation. I shall consider versions of this approach in chapter 9.

We might also reject 1) – the claim that knowledge (/ justification) is the norm of action – while maintaining that there is a normative connection between knowledge (/ justification) and action and that the epistemic requirement on rational action varies with practical factors. There are two possibilities here, which give rise to the third and fourth options.

Third, we might adopt a pluralist approach. It is often assumed in the debate that there is only one epistemic norm of action.⁷⁰ Of course, if there is only one epistemic norm of action, if

⁷⁰See, e.g., Littlejohn : "It's clear that there is an epistemic norm that governs practical reason, one that determines

knowledge is the norm of action and the epistemic standards that the subject must meet to satisfy the epistemic norm of action varies with practical factors, then the epistemic standards required by knowledge (and hence, knowledge) varies with practical factors. But one may think that the knowledge norm is merely an epistemic norm of action among others, and that different epistemic norms are in force depending on the situation. On this approach, the variation does not concern the epistemic condition of application that is provided by the epistemic norm. It concerns whether the knowledge norm applies.

For example, one may think that sometimes a knowledge norm is in force, sometimes a justification norm, sometimes a truth norm, etc. In particular, recall that a distinction could be made between active and passive reasonings. It not implausible to think that these reasonings are not governed by the same epistemic norm.

It is important to distinguish this pluralist approach from the claim according to which there is only one epistemic norm of action which is context-sensitive in that it provides an epistemic condition of application sensitive to the practical context. Indeed, according to the pluralist approach, there is not only one epistemic norm, but there are several epistemic norms of action, which are conditioned norms. The epistemic condition of application that each norm provides need not be context-sensitive.

We may represent the proposal in the following way (where ‘s’ is the subject, ‘ ϕp ’ is treating p as a reason and ‘E p ’ is an epistemic position with respect to p):

If C1, then $O_s(\phi p, s \text{ if(f) } E1p, s)$

If C2 then $O_s(\phi p, s \text{ if(f) } E2p, s)$

...

Thus, it might be that there is a knowledge norm of action, but if this norm is conditioned, it is in force only when a certain practical condition obtains. If this condition can vary, then even if there is a knowledge norm of action, sometimes knowledge may be not necessary, or sufficient, for rational action.

This option has been suggested in the literature.⁷¹ It can be defended more or less in the same way as the second option, which proposes a warrant account. Since these approaches differ only in the specific way in which they construe the variability of the epistemic condition – either in terms of a single context-sensitive epistemic norm or in terms of several context-insensitive epistemic norms – I will focus merely on the warrant account and leave this option aside.

whether it's epistemically proper to treat something as a reason to F. It's not clear what this norm is" (2012, 135).

⁷¹See Levin (2008, 366): "What if, instead, it is the norms for assertion and practical reasoning themselves that shift, according to the subjects' interests, values, stakes, or other pragmatic considerations? Then presumably, there wouldn't be onerous practical consequences of our not knowing various "ordinary" propositions about the world, since, at least in some circumstances, we could perfectly well assert, or use as premises in our practical reasoning, propositions that we don't know". See also Levin (2008, 377): "the norms of practical reasoning, like the norms of assertion, seem to be a diverse and complicated lot, involving a variety of practical considerations".

Finally, it pays to note that we can endorse a biconditional version of the (supposed) normative connection between knowledge and action, or a version in which knowledge is merely a necessary epistemic condition on action. Thus, a fourth option is to argue that there is a normative connection between knowledge and action in the sense that knowledge is necessary, but not sufficient, for action. Chapter 8 will argue that knowledge is necessary for action and chapter 9 that knowledge is not sufficient for action. We may then opt either for a shifty view of the norm or for an invariantist view of the norm. In chapter 9, I shall defend the idea that the epistemic norm of action is context-sensitive. Chapter 10 will then defend a context-sensitive certainty account of the epistemic norm of action.

7.5 Conclusion

It is uncontroversial that there are epistemic evaluations of actions. However, the notion of epistemic evaluation of action is equivocal. When evaluating an action epistemically, we may focus on subjective reasons, on expected reasons, on whether the agent appropriately treats a certain consideration as a reason in the non-deliberative sense, on whether she appropriately treats this consideration as a reason in the deliberative sense, on whether a certain possibility determines the rational thing to do, and, finally, on whether no epistemic weaknesses in the agent's epistemic position stand in the way of a certain consideration to justify the action. These are different ways of understanding the so-called epistemic norm of action, and the epistemic condition that must be satisfied by the agent need not be the same in all these approaches. This point is crucial when assessing whether the supposition of a normative connection between knowledge and action compromises the orthodox account of knowledge.

Indeed, it has been argued that if the epistemic norm of action is knowledge (or epistemic justification), and the epistemic standards for being rational to act on p varies with practical factors, then the orthodox accounts of knowledge (and epistemic justification) are in trouble. In reply to this argument, there are four strategies available to the orthodox philosopher. She can try to reject the claim that there is a normative connection between knowledge and action. She can try to show that the variability of the epistemic standards does not concern the epistemic condition provided by the knowledge norm of action correctly understood. She can reject the monist assumption and argue that there are several epistemic norms of actions. Finally, she can argue that knowledge is necessary but not always sufficient for action. In the next chapter, I shall assess the first strategy in considering the arguments in favour of the claim that there is a normative connection between knowledge and action.

Chapter 8

The necessity of knowledge for action, and the warrant account

8.1 Introduction

Is there a normative connection between knowledge and action? A prominent argument against the orthodox account of knowledge is based on the claim that knowledge is the epistemic norm of action. As a result, most orthodox philosophers deny that there is such a connection. Often, these philosophers endorse a warrant account of action instead. In this chapter, my aim is to show that the warrant account is not a fully satisfactory alternative to the knowledge account. In particular, I defend the claim that knowledge is necessary for rational action.

To set up the debate, I start by presenting the main arguments in favour of the claim that knowledge is necessary and sufficient for rational action. These arguments concern in particular linguistic data, bad reasonings, and principled considerations. Then, I consider a general criticism that has been raised against the linguistic argument in favour of the knowledge account of action. I argue that the main attempt to explain the linguistic data in a way which is compatible with the warrant account is not fully satisfying, in particular if one adopts the warrant account. Next, I focus on arguments that cast doubt on the necessity of knowledge for rational action. Against recent criticisms, I defend the standard reply which appeals to the distinction between excuse and justification. In the course of doing it, I also stress that the warrant account does not do justice to our intuitions underlying the principled argument for the necessity of knowledge and the argument from bad reasonings. I conclude that we should adopt the claim that knowledge is always necessary for rational action. Although this does not threaten the orthodox account of knowledge, this leads to reject the warrant account of action.

8.2 Arguments for the knowledge norm of action.

8.2.1 The linguistic argument

Everyone agree that it is most natural to make epistemic assessments of action by using ‘know’ (and its cognates).¹ Consider cases of defence and criticism using ‘know’ and suggesting that knowledge that *p* is necessary for rational action on *p*:

Defence: I did not take my umbrella because I did not know that it was raining.

Criticism: Why did you use this needle? You did not know it was clean.

Consider cases of defence and criticism using ‘know’ and suggesting that knowledge that *p* is sufficient for rational action on *p*:

Defence: The coach know the other team was going to blitz. That’s why it made so much sense for him to hold his tight ends back to block. (Fantl and McGrath, 2009, 60)²

Criticism: The coach shouldn’t have sent his tight ends out on fourthdown. He know the other team was going to blitz. (Fantl and McGrath, 2009, 60)³

These cases using ‘know’ suggest that knowledge is the norm of action, in that knowing *p* would be necessary and sufficient to appropriately act on *p*.

The linguistic argument from the use of ‘know’ to the claim that knowledge is the epistemic norm of action is not merely that there are examples where using ‘know’ is appropriate. More importantly, we must also note that ‘know’ is the *default* word for epistemic assessments of actions. To see this, consider the following case, from Hawthorne and Stanley (2008):

RESTAURANT. Hannah and Sarah are trying to find a restaurant, at which they have time-limited reservations. Instead of asking someone for directions, Hannah goes on her hunch that the restaurant is down a street on the left. After walking for some amount of time, it becomes quite clear that they went down the wrong street.

In this situation, it may be granted that Hannah’s action is wrong because a mere hunch that *p* is insufficient to act on *p*. However, as Hawthorne and Stanley (2008) note,

[a] natural way for Sarah to point out that Hannah made the wrong decision is to say, “You shouldn’t have gone down this street, since you didn’t know that the restaurant was here”.

¹See in particular Stanley (2005b); Hawthorne and Stanley (2008); Fantl and McGrath (2007, 2009); Gerken (2015).

²See also Stanley (2005b, 10).

³See also Fantl and McGrath (2007, 561).

Why is it natural to use ‘know’? If knowledge is not necessary or sufficient for rational action, this is a surprising fact.

We do not merely use ‘know’ as a default term. We *prominently* use this word in epistemic assessments.⁴ These two claims are different. There is no *a priori* reason to think that most of our epistemic assessments of action are made by default – in the absence of the relevant information concerning the specificity of the case. The claims that ‘know’ is used as a default term *and* prominently provide a stronger argument in favour of the knowledge norm than either of these claims alone.

Finally, to augment the linguistic argument, we should note that knowledge questions are always relevant when it comes to make a practical decision. As Reed (2010, 232) writes:

[Y]ou are trying to decide whether to check if the train stops in Foxboro because it is extremely important that you get there as quickly as possible. You have not yet decided whether it is rational to check if the train makes that stop. One of the relevant factors in your decision would presumably be an answer to the question, do you know the train will stop there?

If knowledge is not always relevant for rational action, it is unclear why an answer to the question “Do you know that p?” is always relevant to your practical decision. In addition, a positive answer to this question is always decisive, even when the stakes are high.⁵ This suggests that knowledge that p is always a necessary and sufficient epistemic condition for rationally acting on p.

In sum, the facts that we use ‘know’ by default and prominently to defend and criticize actions, and that ‘knowledge’ ascriptions are always relevant and decisive when it comes to make a practical decision, provide a strong argument in favour of the thesis that knowledge is the norm of action.

8.2.2 The bad reasoning argument

A second type of consideration suggests that knowledge is necessary. Consider the following case:⁶

BAD REASONING. A subject owns one ticket in a lottery with 1000 tickets, and there is only one winning ticket. The winner earns 1000 dollars. As a matter of fact, her ticket is a loser ticket. The subject has no information about the lottery result, and she reasons in the following way:

1. The ticket is a loser ticket.

⁴See, e.g., Gerken (2015). To my knowledge, there is no systematic study showing that ‘know’ is prominent in epistemic appraisal of action, but ‘know’ is one of the ten words most used (see Gerken (forthcoming)). Gerken agrees that the use of ‘know’ is prominent, but he is not committed to the claim that ‘know’ is used as the default term, which seems to be a stronger claim.

⁵See Fantl and McGrath (2007, 562).

⁶See Hawthorne (2004, 29, 85, 174-175) and Hawthorne and Stanley (2008).

2. So if I keep the ticket, I will get nothing.
3. But if I sell the ticket, I will get one dollar.
4. So I should sell the ticket.

This is a bad reasoning. Many have the intuition that this reasoning is bad because the first premise is not known.⁷ If there is no connection between knowledge and appropriately acting on a proposition, why do we have the intuition that the reasoning is wrong because the first premise is not known? This intuition strongly suggests that knowledge that p is necessary for using p as a premise in a practical reasoning, and hence, for appropriately acting on p .

8.2.3 The principled arguments

The principled argument for the sufficiency of knowledge

In favour of the claim that knowledge is always sufficient for rational action, it may be argued that there is a principle to the effect that it is required not to ignore what one knows when it comes to acting and to act accordingly. Four considerations support this claim.

First, we often reason in conformity to the contrapositive of the principle that if a subject knows that p , then it is rational for her to act on p . Consider the following case from Fantl and McGrath (2007, 570):

DICTIONARY. After an argument over the spelling of a word, in which you claimed you knew how it was spelled, I might catch you looking the word up in a dictionary: “Well, if you knew how it was spelled, you wouldn’t have had to check, would you?! You don’t know!”

The fact that it seems appropriate to reason in this way is easily explained on the basis of the principle that if you know that p , it is rational for you to act on p (provided that the choice is p -dependent). If knowledge that p is always sufficient for rational action on p , you need not check whether p if you know that p . Therefore, as you check the spelling of this word, on the assumption that this principle is true, you must not know the spelling of this word.

Second, the following sentences sound incoherent:⁸

- (1) I know that the best thing to do is to take an umbrella, however, it is not rational for me to take an umbrella.
- (2) John knows that the best thing to do is to take his umbrella, but it is not rational for John to take it.

⁷Note that it is not irrational to sell the ticket on the basis of the expected utility of this action.

⁸See Fantl and McGrath (2007).

The fact that such sentences sound incoherent shows the difficulty of explaining the relevance of ‘knowledge’ ascriptions to rational action in terms of Gricean implicatures, since these sentences are instances of attempted cancellation of the supposed implicatures.⁹ In the absence of a different pragmatic explanation, this apparent incoherence is best explained on the basis of the principle according to which knowledge is always sufficient for rational action.

Third, if we assume that knowledge is not always sufficient for practical deliberation, but that knowledge is always sufficient for theoretical deliberation, it has to be possible for situations to arise where rationality would command to be akratic. If it cannot be rational to be akratic, the epistemic norm of practical rationality cannot be stronger than the epistemic norm of theoretical rationality.¹⁰

To see this, consider a case in which you know that p , and you know that if p , then A is the best thing you can do.¹¹ Arguably, from your knowledge that p and that if p then A is the best thing you can do, you can come to know that you should do A . Suppose that you know that you should do A . Thus, you have a justified belief that you should do A . Now suppose that there are cases where the epistemic norm of rational action is stronger than knowledge. In such cases, although you know that p , it may be that you cannot rationally act on p , nor treat p as a reason to act, nor use p as premise in a practical reasoning concluding in an intention to do A . If so, you may have a rational belief that you ought to do A (indeed, you know you should do A), but at the same time, it would be irrational for you to intend to do A (indeed, you do not satisfy the epistemic norm of action with respect to a relevant consideration). In sum, it would be rational for you to believe that you ought to φ while not intending to φ . This entails that rationality would sometimes require to be akratic. But *enkrasia* – the principle of rationality according to which you ought, if you believe that you ought to φ , intend to φ – is perhaps one of the most basic principle of rationality.¹²

Fourth, if the epistemic norm of practical rationality were (sometimes) stronger than the epistemic norm of theoretical rationality, we should and would (sometimes) segregate between reasons available for action and reasons available for belief. However, such a segregation seems “barmy”.¹³ For example, suppose that you consider whether it is safe to cross an ice pond. It would be strange if you took the fact that the park authorities told you that the ice is thick enough to cross as a good reason to believe that it is safe to cross while criticising those who cross on this basis. If you think one should not cross on the basis of what the park authorities have told, then presumably, you think one should not believe that it is safe to cross on this basis. In addition, such a segregation does not seem to correspond to the way in which we in fact deliberate. This suggests that a common epistemic norm governs theoretical and practical rationality.¹⁴ It is disputed whether

⁹See Fantl and McGrath (2007) and chapter 4.

¹⁰See Fantl and McGrath (2002, 2007, 2009); Littlejohn (2012, 2014).

¹¹See Fantl and McGrath (2002, 2007, 2009, 71-72).

¹²See Broome (2013, 22-25).

¹³Fantl and McGrath (2009, 74)

¹⁴See also Littlejohn (2012, 202-209).

knowledge rather than justification is the epistemic norm of theoretical deliberation. However, it seems undeniable that if you know that p , then p can be treated as a reason to believe that q (provided that p obviously supports q). As Fantl and McGrath (2009, 73-76) suggest, if there is a common epistemic norm for theoretical and practical rationality, and knowledge is sufficient for theoretical rationality, it follows that knowledge is sufficient for practical rationality.

In sum, four considerations speak in favour of the idea that it is irrational not to act on something one knows (when this thing is relevant for action). This suggests that knowledge is always sufficient for rational action.

The principled argument for the necessity of knowledge

On the basis of principled considerations, we can also argue that knowledge is always necessary for rationally treating p as a reason to act. When we assess actions, we sometimes focus on the agent and whether she is following the norm, rather than on whether she merely conforms to it. Someone can do the right thing for the wrong reason, or for no reason. In such cases, we may say that the subject does something right, but wrongly. If the subject has no excuse, she can be criticised on this account. These assessments concern whether the subject intends to conform to the norm when she acts, and they are sensitive to what we take to be the subject's (rational) beliefs about what she is doing.

Suppose that S (rationally) believes that in doing A , she conforms to N . For example, suppose that John (rationally) believes that by stopping at a red light, he conforms to traffic rules. If S does A for the reason that A is a way of conforming to N , typically, we may say that S follows N in doing A . If John stops at a red light for the reason that stopping at a red light is a way of conforming to traffic rules, typically, we may say that John follows the traffic rules in stopping at a red light. Suppose that N^* is a non-normative rule for S . That is, suppose that the rule N^* does not apply to S . Suppose that S rationally and correctly believes that N^* is not a normative rule applying to him. We should see no problem if S does A while believing that in doing A she does not meet N^* . For example, if Paul is not playing chess, and he moves a pawn in violation of the rules of chess, while knowing this, he should not be criticised. In contrast, suppose that N is a normative standard that applies to S . Suppose that S rationally and correctly believes that N is a normative standard. Clearly, if S does A^* while believing that in doing A^* she does not meet N , she can be criticised. S is acting against (what she correctly sees as) her reasons. If John does not stop to a red light while believing that he violates the traffic rules (and that these rules apply to him), he can be criticised.

Now, suppose that knowledge is not the norm of action in the sense that knowledge is not necessary for rational action. Then, there should be no problem if a subject (rationally) believes that she does not know that p , while at the same time acting for the reason that p (or while treating p as her reason for action). However, it seems that a subject who knows, or can tell, that she does

not know that *p* while, at the same time, treating *p* as her reason for acting, is irrational.¹⁵ Consider the following dialogue:

A - Why do you take your umbrella?

B - Because it is raining.

A - Do you know that it is raining ?

B - No, I do not. So what? I do not know whether it is raining, but I take my umbrella for the reason that it is raining.

B's last claim is clearly odd and sounds incoherent. This is surprising if knowledge is not necessary for treating something as a reason for action. Of course, it can be rational to take an umbrella even if one does not know that it is raining, but what we treat as a reason in such a case should not be that it is raining. It should be that it is probable that it raining, or that there is a chance that it is raining. But if knowledge is not necessary for rationally treating something as a reason to act, it is unclear why a subject could not coherently take her umbrella for the reason that it is raining while believing that she does not know that it is raining. That this seems incoherent is an argument in favour of the necessity of knowledge.

8.3 The linguistic argument and the warrant account

In the previous section, I have presented the main arguments that can be proposed in favour of the claim that knowledge is the epistemic norm of action. However, it is often thought that if we endorse the claim that knowledge is the epistemic norm of action, we must reject the orthodox account of knowledge. Therefore, many orthodox philosophers have tried to show that these arguments are not decisive, and they have put forth other considerations speaking against the claim that knowledge is the norm of action. The main alternative to the knowledge account of action that has been developed on these bases is the warrant account. In this section, I shall consider criticisms regarding the linguistic argument. However, I shall argue that the main proposal to date aiming at explaining these linguistic and conversational considerations without appealing to a normative connection between knowledge and action is not fully satisfactory and in tension with the warrant account itself.

8.3.1 The significance of the linguistic argument

Against the linguistic argument in favour of the knowledge norm of action, philosophers have noted that the fact that we appropriately use 'know' in some cases does not show that we may

¹⁵See DeRose (2009, 94) for a similar argument in favour of the necessity of knowledge for proper assertion: "What would show that the knowledge account of assertion is too strong are cases in which it's apparent that a speaker properly asserts what she doesn't even reasonably take herself to know."

appropriately use this term in all cases, and, as a matter of fact, we do not always use ‘know’ to make epistemic assessments of actions.¹⁶ Sometimes, we use stronger or weaker epistemic terms, like ‘certain’, ‘think’ or ‘believe’. This may cast some doubt on the idea that there is a good argument from the use of ‘know’ to the claim that knowledge is the norm of action. Of course, if knowledge is necessary for rational action, it may be appropriate to criticise a subject acting on a certain proposition, and to defend a subject refraining from acting on a proposition, by using epistemic terms weaker than knowledge. Indeed, we can make such criticisms and defences on the basis of the fact that the subject does not meet a necessary condition for knowledge. Likewise, if knowledge is sufficient for rational action, it may be appropriate to defend a subject acting on a certain proposition, and to criticise a subject refraining from acting on a proposition, by using epistemic terms stronger than knowledge. Indeed, we can make such defences and criticisms on the basis of the fact that the subject meets a sufficient condition for knowledge.

However, there are cases where we use epistemic terms weaker than knowledge to defend a subject acting on a certain proposition, and to criticise a subject refraining from acting on a proposition:

Defence for acting on p (weaker): I am late, but I had every reason to think that the train left at 12.01.

Criticism for not acting on p (weaker): You should have got potatoes. You thought we didn’t have any. (Brown, 2008, 171).

In addition, there are cases where we use epistemic terms stronger than knowledge to defend a subject refraining from acting on a certain proposition, and to criticise a subject acting on a proposition:

Defence for not acting on p (stronger): I checked whether the bank was open on Saturdays because I was not sure.

Criticism for acting on p (stronger): a mother might criticise her teenager’s late departure from a party by saying ‘You shouldn’t have left so late. You weren’t certain there’d be a bus home that late’ or ‘You shouldn’t have left so late. You didn’t know for sure that there’d be a bus at that time’. (Brown, 2008, 172).

The fact that we use epistemic terms stronger and weaker than ‘know’ in these ways may elicit different reactions regarding the linguistic argument in favour of the knowledge norm of action.

It would be overreacting to reject this argument for the reason that it would prove too much, though. Consider the fact that we sometimes use ‘certain’. Gerken (forthcoming, 6.3.a) is right that “[i]t would be misguided to jump from the fact that [“you weren’t certain”] is natural to the conclusion that certainty is the epistemic norm of rational action”, and that, likewise, it would be

¹⁶See Brown (2008, 170), Gerken (forthcoming)).

misguided to jump from the fact that ‘know’ is natural to the conclusion that knowledge is the epistemic norm of action.¹⁷ However, that does not mean that these linguistic facts should not be explained or are insignificant. If they are best explained under the assumption of a certain epistemic norm of action, this constitutes an argument in favour of this epistemic norm.

In addition, as said above, the linguistic argument is not a mere jump from the fact that ‘know’ is natural, in the sense of “sometimes used”, to the claim that knowledge is the norm of action. Rather, the complete argument is that ‘know’ is the default term, that this term is used prominently, and that ‘knowledge’ ascriptions are always relevant and decisive for action. Even if it is in some sense “natural” to use the word ‘certain’, it is clear that this term is not the default and prominent term of epistemic assessments of actions. On the basis of these linguistic considerations, it makes much more sense to defend a knowledge norm than a certainty norm (on the assumption that knowledge does not entail certainty).

Let us grant that our ordinary practice of assessing action is significant, and that it shows a certain variation in the epistemic terms employed. We might think that this “suggests that the standard for practical reasoning varies with context: sometimes the standard is knowledge, sometimes it is less than knowledge, and sometimes it is more than knowledge” (Brown, 2008, 171). Yet, endorsing a norm with context-sensitive epistemic standards is not enough to do justice to the aforementioned linguistic considerations. Again, ‘know’ is the default and prominent term of epistemic assessments of actions, and ‘knowledge’ ascriptions are always relevant and decisive for action. If the norm provides context-sensitive epistemic standards, these facts should be explained.

We are left in the following situation. Philosophers endorsing the knowledge account of rational action face the challenge of explaining why we sometimes use a certain epistemic vocabulary in ways apparently incompatible with the thesis that knowledge is the norm of action. Orthodox philosophers rejecting the knowledge account of rational action face the challenge of explaining why ‘know’ is the default and prominent term of epistemic assessment, and why ‘knowledge’ ascriptions are always relevant and decisive when it comes to make a practical decision. If the latter philosophers can provide a good explanation of these linguistic considerations and the advocates of the knowledge account cannot give a good explanation of the variations in our epistemic vocabulary, then, overall, the linguistic considerations do not support the knowledge account of action.

However, I shall show that the main attempt to date to explain these linguistic data without appealing to a normative connection between knowledge and action is not fully satisfactory, in particular on the assumption that the warrant account is true. Further, I shall argue in section 4 that the fact that we use epistemic terms weaker than ‘know’ is fully compatible with the idea that knowledge is necessary for rational action. I shall consider whether knowledge is always sufficient in the next chapter.

¹⁷Gerken thinks that it is obviously misguided to jump from the fact that it is natural to use ‘certain’ to the conclusion that certainty is the epistemic norm of action because he thinks it is obvious that certainty is not the epistemic norm of action. To my mind, this is open to debate. See also ?, 532.

8.3.2 The linguistic considerations and the warrant account

As said above, the main competitor of the knowledge account of rational action is the warrant account. It has been developed by Gerken.¹⁸ According to Gerken's account, the epistemic norm of action provides different epistemic standards associated with different deliberative contexts, where a deliberative context is determined by a set of practical factors. More precisely, according to the warrant account:

Warrant Account. In the deliberative context, DC, S meets the epistemic conditions on rational use of (her belief that) p as a premise in practical reasoning or of (her belief that) p as a reason for acting (if and) only if S is warranted in believing that p to a degree that is adequate relative to DC. (Gerken, 2011, 530).

The practical factors that determine the degree of warrant that is adequate include in particular the urgency of action, the availability of evidence and other options, the stakes, and the social role of the agent. However, is this account compatible with the aforementioned linguistic considerations?

In order to explain the prominence of knowledge in epistemic assessments of action, Gerken appeals to the idea that in normal cases of epistemic assessments of action, knowledge-level warrant is (very) frequently necessary and sufficient for rational action:

Normal Coincidence. In normal cases of epistemic assessment, the degree of warrant necessary for S's knowing that p is frequently necessary and very frequently sufficient for the epistemic permissibility of S's acting on (the belief that) p. (Gerken 2015; forthcoming, 6.3c)

On the basis of Normal Coincidence, Gerken maintains, we can explain why 'know' is prominent in epistemic assessments of actions while upholding the warrant account.

However, the problem is that Normal Coincidence seems *ad hoc* in this framework. Why is it that in most normal cases of epistemic assessments, knowledge-level warrant is (very) frequently necessary and sufficient? We might think that this normative fact, if it is a fact, explains the descriptive fact that we prominently use 'know' in (normal) epistemic assessments of actions. Still, this normative fact is itself surprising if we endorse Gerken's warrant account. According to this account, there are many parameters that can determine the deliberative context. They can combine in many different ways, and, presumably, each one can affect the epistemic warrant more or less strongly. Thus, it is very surprising that given all the ways in which these parameters can combine and the different degrees to which they can affect the epistemic warrant, they provide in (very) frequent normal cases exactly the same epistemic condition, namely, knowledge-level warrant. In sum, if the warrant account is true, it is very unlikely that Normal Coincidence is true, and hence,

¹⁸See Gerken (2011, 2015); Gerken (forthcoming).

it is particularly *ad hoc* to assume it without further explanation.¹⁹

Perhaps appealing to the idea that it is only in *normal* cases of epistemic assessments that knowledge-level warrant is frequently required can help to show that Normal Coincidence is not at odd with the warrant account. After all, the advocate of the warrant account may insist that practical situations in which more or less than knowledge-level warrant is required are (very) frequently abnormal. Even if practical factors can combine in many different ways and affect the warrant to different degrees, this is mostly in *abnormal* cases that they affect the warrant in such a way that more or less that knowledge-level warrant is required.

However, this appeal to the distinction between normal and abnormal practical contexts is problematic. First, the claim that ‘know’ is prominent in epistemic assessments of action is a statistical claim regarding our actual epistemic assessments. In contrast, Normal Coincidence is a claim about a statistical fact given normal (epistemic and practical) circumstances. But the relation between frequency and normality is unclear in Gerken’s account. If one thinks that normal practical situations are not necessarily the most frequent in the actual circumstances, then it is difficult to see how Normal Coincidence could help in explaining the fact that ‘know’ is prominent in actual circumstances. On Gerken’s approach, the notion of abnormal practical circumstances seems to be sometimes determined in terms of frequency, and sometimes not.

Thus, Gerken considers cases similar to the high stakes bank case, such as the following:

SURGEON. A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetised on the operating table. The operation hasn’t started as the surgeon is consulting the patient’s notes. The student is puzzled and asks one of the nurses what’s going on:

Student: I don’t understand. Why is she looking at the patient’s records?
She was in clinic with the patient this morning. Doesn’t she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn’t operate before checking the patient’s records. (Brown, 2008, 176).

¹⁹It is worth noting that Douven (2006, 469-470) uses a similar strategy regarding the norm of assertion. He argues for a rational credibility norm of assertion and then argues that, to the extent that, in most cases, we believe what it is rational to believe and our rational beliefs are true and not true by mere luck, it is plausible that in most cases what is rationally credible is known. This may explain why ‘know’ is so prominent as a term of epistemic assessment. However, in Gerken’s approach, the epistemic requirement for action is determined by the practical context. There is no reason to think that practical contexts are uniform regarding the different parameters that they involve. Thus, the idea that there is a coincidence between the warrant required by most (normal) practical context and the warrant required by knowledge is much more surprising.

Then, he writes:

For example, it should be clear enough that the practical contexts in [such cases] are abnormal due to the abnormally high stakes. Again, a broad externalist consideration motivates this assumption. It is infrequent that a single action determines life or death.

Here, it seems that the fact that a kind of action is not frequent makes it abnormal. However, Gerken continues:

Of course, surgeons face life-or-death scenarios more frequently. But this may be acknowledged by saying that surgeons are frequently in contexts with abnormally high stakes.

Here, it seems that the fact that the action is frequent does nothing to show that it is normal. As a result, it is very unclear how frequency and normality are related, and hence, it is difficult to see whether this distinction can help to show that there is no tension between the warrant account and Normal Coincidence.

In addition, note that it is problematic to classify cases with high stakes (like the surgeon's case) as abnormal on the basis of the fact that they feature high stakes. This would manifest a tendency to categorize most of the practical situations in which the practical factors are such that they require a warrant stronger (or weaker, if the stakes are very low) than knowledge as abnormal. But then, Normal Coincidence would tend to be trivially true, for given the very notion of a normal practical situation used, most situations in which more or less than knowledge-level warrant is required would count as abnormal situations.

As a result, it is unclear how appealing to a distinction between normal and abnormal practical situations could help to avoid the tension between the warrant account and the actual prominence of 'know' in epistemic assessments of actions. Either normality is not defined in terms of frequency, and hence Normal Coincidence does not explain the prominence of 'know' in actual epistemic assessments of actions. Or normality is defined in terms of frequency, and then, given all the different ways in which the practical factors can combine and affect the warrant, it is quite odd that most actual practical situations require exactly a knowledge-level warrant.

Further, let me note that on Gerken's view, the notion of knowledge-level warrant is not factive. Still, 'know' is a factive term. If the warrant required for action is never factive, it is odd that we use a factive term prominently to make epistemic assessments of actions.²⁰

In light of these considerations, it does not seem that appealing to Normal Coincidence is sufficient to explain why we naturally and prominently use 'know' in epistemic assessment of actions, in particular if we assume the warrant account.

Finally, Normal Coincidence cannot help to understand why 'knowledge' ascriptions are always decisive and relevant when it comes to make a practical decision. 'Knowledge' ascriptions are always

²⁰Gerken (2015) acknowledges this point.

relevant and decisive, at least in their first-person version, in the sense that saying ‘I know that p’ settle the question whether to act on p, and when the issue is whether to act on p, such assertions are never challenged as inappropriate (e.g., a reply such as “The question is not whether you know that p” would be very odd).²¹ Indeed, if the warrant account is true, and, more generally, if a context-sensitive account of the epistemic norm of action and the orthodox view of knowledge are true, there are situations where more or less than a knowledge-level warrant is required. In both of these cases, it would not be fully relevant to consider what one knows before acting. But then, why is it that ‘knowledge’ ascriptions (or denials) are relevant or decisive even in these cases?

On this score, Gerken (forthcoming, ch. 8) appeals to the idea that ‘knowledge’ ascriptions may have the function of recommending for and against various courses of action. The basic idea is that an assertion may have different illocutionary effects, and hence, fulfil different communicative functions. Thus, for example, by asserting “The window is still open” you may make the indirect speech act of commanding to close the window. In the same way, ‘knowledge’ ascriptions may serve different communicative functions. In particular, it has been argued that they can provide assurance (Lawlor, 2013), identify reliable informants (Craig, 2010), etc. Gerken adds that they may fulfil the function of directing action by generating a directive implicature, or that they have a directive force.

Like declarative implicatures, directive implicatures are conversational implicatures. They are contextually variable and can be cancelled. We must distinguish particularised and generalised implicatures. Particularised implicatures heavily rely on the particularity of the context and the background assumptions ascribed to the interlocutors, while a generalised implicature often constitutes the default interpretation of the speech act. For example, “Can you pass the salt?” generally implicates a request for the salt. This is the default interpretation of this speech act.

If ‘knowledge’ ascriptions are typically associated with a generalised directive implicature, we can explain why they can be most of the time relevant and decisive when it comes to make a practical decision. Even if in the particular context, more or less than knowledge with respect to p is required to do A, and the question is whether to do A, saying “I know that p” may generate a directive to do A. This directive is epistemically proper only if the ‘knowledge’ ascriber is, given the context, in the relevant epistemic position with respect to p (or with respect to a relevant proposition q that is implicated by asserting p), so that p (or q) can properly figure in a line of reasoning concluding in a directive to do A. For example, suppose that the question is whether we should take an umbrella, and this depends on whether it will rain. By saying “I know that it is very cloudy”, I may implicate that it is very likely that it will rain, and thereby, I may direct the action of taking an umbrella.

Still, if this is the right explanation, the directive implicature generated by a ‘knowledge’ ascription must be cancellable. However, on this score, there are disanalogies between, on the one

²¹See chapter 4 section 8.

hand, this supposed directive implicature, and, on the other hand, other generalised implicatures of ‘knowledge’ ascriptions and directive implicatures of other assertions.

To see this, consider Brown’s surgeon case, and suppose that the surgeon says “Of course, I know that it is the left kidney that is diseased. But that does not mean that I should not check the patient’s record”. This is an instance of a supposed cancellation of the supposed directive implicature. Now, suppose that the student challenges the surgeon by saying: “But why do you need to check if you know?”. This challenge does not seem obviously inappropriate. Of course, the surgeon might insist and say something like “The stakes are high, I should double-check”. But the student also might insist: “I am aware that the stakes are high. But you know that you are right. So, why checking?” It may seem that in saying that he knows that it is the left kidney that is diseased, the surgeon is somewhat committed to the idea that he need not check, and this commitment is quite difficult, and perhaps impossible, to cancel. At least, assuming that there is such a commitment easily explains why the student’s challenge is not obviously inappropriate.

Contrast with a case in which the implicature of the ‘knowledge’ ascription concerns whether the subject is a reliable informant. Suppose that I say to Robert “John knows that p, but I do not mean to suggest that John is a reliable informant with regard to p or p-related matters in general.” This is an instance of cancellation of the implicature of flagging John as a reliable informant. Now, suppose that Robert challenges what I just said by saying “But if John knows that p, why is John not a reliable informant with regard to p or p-related matters?”. It is really unclear how this challenge is appropriate. When I say that John knows that p, I am in no way committed to the idea that John is a good informant with regard to p or p-related matters. It seems clear that there is no necessary connection between knowing p and being a reliable informant with regard to p or p-related matters.

The fact that the student’s challenge seems much more appropriate than Robert’s challenge is a first reason to doubt that we should explain the relevance of ‘knowledge’ ascriptions to action in terms of directive implicatures. Indeed, it does not make sense to challenge the (supposed) cancellations of implicatures of ‘knowledge’ ascriptions concerning the reliability of the informant, whereas it makes sense to challenge the (supposed) cancellations of implicatures of ‘knowledge’ ascriptions concerning directives to act on the known proposition. This suggests that we should not account for the directive force of a ‘knowledge’ ascription in the same way as we account for the flagging implicature of these ascriptions.

Second, note that, for some reason or another, we may pretend not to understand an obvious directive implicature typically associated with the speech act under consideration. For example, if asked “Can you pass me the salt?”, I may just reply “Yes I can!”. This may happen, for example, if I want to make a joke, or if I want to signify that I am not disposed to help. However, with ‘knowledge’ ascriptions, things are different. For example, in situations where more or less than knowledge is required, it is very odd to reply to someone saying “I know that p” (if p is relevant for

action) “The question is not whether or not you know that p”, or “So what?”. But if the approach in term of directive implicature were right, this reply could just amount to a pretention not to understand the directive implicature, and could be explained, e.g., as a joke. This asymmetry between cases like “Can you pass me the salt?” and cases of ‘knowledge’ ascription also casts doubt on the fact that ‘knowledge’ ascriptions are always relevant and decisive for action because they are associated with directive implicatures directing a certain course of action.

Third, if ‘knowledge’ ascriptions could be appropriate just in virtue of this supposed directive force, there would be cases where a subject who does not know the relevant proposition could assert that she knows it just to direct the relevant course of action. However, consider a low stakes case where Keith does not know whether the bank is open but the cost of error is so low that coming back on a mere hunch that the bank is open is rationally permissible. It is unclear that Keith can appropriately say “I know that the bank is open” to direct going back. And, indeed, suppose that his wife notes that he has no evidence for this fact. The following defence form Keith would be difficult to accept: “I did not mean that I had any evidence whatsoever for the fact that the bank is open, but only that we should come back”. Also, it is far from clear that, in this situation, Keith could appropriately assert “I do not know that the bank is open, but we must go back tomorrow for the reason that it is open”. (Of course, Keith could say that they should come back because it is likely enough that the bank is open, but this is different from acting for the reason that the bank is open or treating this proposition as a reason.) Still, if the approach in terms of directive force were right, this would just be an instance of cancellation of the supposed directive force of ‘knowledge’ denials, namely, a directive not treat a certain proposition as a reason, and if we are in a context where knowledge (or knowledge-level warrant) is not required for treating something as a reason, this claim should sound perfectly acceptable.

Finally, it may seem that this approach faces the charge of *ad hocness* in the same way as the approach in terms of declarative implicatures developed by Rysiew or Brown.²² One cannot just assume that it is in virtue of directive implicature associated with ‘knowledge’ ascriptions – rather than in virtue of the semantic meaning of ‘S knows that p’ – that ‘knowledge’ ascriptions fulfil their communicative function of directing action. We must explain why and how such a pragmatic meaning is associated with ‘knowledge’ ascriptions. Arguably, Gerken’s claim that the concept of knowledge and the word ‘know’ can be used as good heuristics to make epistemic assessments of action may help to explain why this approach is not *ad hoc*. However, the claim that it is a good heuristic to use the concept of knowledge and the word ‘know’ to make epistemic assessments of action is plausible only if we assume Normal Coincidence (and the idea that the normal practical contexts are the most frequent). If Normal Coincidence is in tension with the warrant account, as I have tried to show, then it is unclear how the approach in terms of directive implicatures is in a better position than the approach in terms of declarative implicatures in respect of *ad hocness*.

²²See chapter 4.

In summary, it is difficult to see how the warrant account can deal with all the linguistic data. If it appeals to Normal Coincidence, then the explanation seems *ad hoc* and, anyway, it is insufficient. If it also appeal to directive implicatures of ‘knowledge’ ascriptions, it faces, in particular, the problem that the supposed implicatures do not seem always cancellable.

Thus, if the advocate of the knowledge norm can explain why we sometimes use epistemic terms stronger and weaker than ‘know’, then the knowledge account does better than the warrant account, overall. In the remaining of this chapter, I shall show that the advocate of the knowledge norm may explain why we sometimes use terms weaker than ‘know’ to defend a subject acting on a certain proposition. Basically, the explanation is that these defences do not constitute justifications, but excuses. To develop this point, I shall focus on the main types of cases that have been put forth against the necessity of knowledge for rational action. On the way, I shall explain why the fact that we sometimes use weaker terms to criticise a subject for not acting on a proposition do not compromise the knowledge account of reasons. I shall turn to cases where more than knowledge seems required in the next chapter.

8.4 Arguments against the necessity of knowledge and the excuse manoeuvre

8.4.1 The epistemic cases and the excuse manoeuvre

The epistemic cases

According to the thesis that knowledge is the epistemic norm of action, knowledge is necessary and sufficient for rational action. Against the claim that knowledge is always necessary, some philosophers have put forth two types of epistemic cases which consist of cases in which the subject does not know the target proposition but may seem to appropriately act on the target proposition given her epistemic position with respect to it. The first type of case features a justified false belief, the second type of case features a Gettiered true belief. Consider a case of a justified but false belief:

RESTAURANT*. Hannah and Sarah have a time-limited reservation for a restaurant. Hannah is extremely well warranted in believing that the restaurant is located down the street. She has checked the restaurant’s website, a reliable guidebook and re-checked with the concierge at their nearby hotel. She decides to go there with Sarah. As they realize that the restaurant has recently moved, Sarah complains that Hannah should not have decided to go down the street because she did not know that the restaurant was there. (Gerken, 2011, 532-533)²³

²³See also Brown (2008); Neta (2009); Fantl and McGrath (2009).

Hannah does not know that the restaurant is down the street since this is false. Still, it seems that Sarah's complain is unreasonable and that Hannah's decision was reasonable.

Consider a case in which a subject does not know the target proposition because she is in a Gettier-like situation:

ERNEST. As Ernest is walking across the plains of Masai Mara, he sees a rock that looks just like a big lion lying on a nearby hill. Consequently, he forms the belief that there is a lion on the hill. After brisk practical deliberation, he loads his rifle and slowly backs away whilst keeping an eye on the lion look-alike rock. As it happens, a real lion lies hidden in the grass on the hill invisible to Ernest. (Gerken, 2011, 535).²⁴

Ernest does not know that there is a lion, because his justified belief is true by luck. Still, it is unclear why he could be criticised for acting in this way. It seems that he did something reasonable. Further, in contrast to RESTAURANT*, Ernest's action and belief are objectively correct. Thus, it is particularly difficult to see in what sense his action could be wrong. The anti-Gettier condition does not seem to have any normative or rational significance.²⁵

With the help of these two cases, we might try to run an incoherence argument against the necessity of knowledge for rational action. Consider the following propositions:

1. If a subject S violates a norm (of rationality), then S is blamable.
2. In RESTAURANT* and ERNEST, the subject is not blamable.
3. In RESTAURANT* and ERNEST, the subject does not know the relevant propositions.
4. Knowledge is necessary for rational action.

1, 2, 3 and 4 are inconsistent. If 4 is true, then either 1, 2, or 3, is false. Given that 2 and 3 seem undeniable, the advocate of a knowledge norm have mainly rejected 1.²⁶

Consider why we should reject 1. As explained in chapter 5, there are various ways in which we can blamelessly violate a norm. In particular, we may have a justification, an excuse, or an exemption.

²⁴See also Brown (2008); Neta (2009); Fantl and McGrath (2009); Littlejohn (2009, 2012).

²⁵See Littlejohn (2012, 2009)

²⁶It worth noting that some philosophers have also argued that 3 is false (see Lord, 2010). Indeed, we might think that in ERNEST and RESTAURANT* the subjects know propositions that make their action justified. For example, although Ernest does not know that there is a lion, he knows that there is something that looks much like a lion, and this is a sufficient reason for Ernest to loads his rifle and slowly backs away whilst keeping an eye on the lion look-alike rock. Likewise, although Hannah does not know that the restaurant is open, she knows that the restaurant's website, the concierge and the reliable guidebook indicate that the restaurant is down the street. These things are sufficient reason for Hannah's going down the street. The problem of this approach, though, is that it is focused on the assessment of the action in isolation from the reasons for which the agent acted. Arguably, the reason why Ernest loads his rifle is not that it looks as if there is a lion, but that there is a lion. Hannah's reason is not that the concierge and the reliable guidebook indicate that the restaurant is down the street, but that the restaurant is open.

It does not seem that the subject in RESTAURANT* or ERNEST has a justification for violating the supposed knowledge norm of practical rationality. Indeed, compare these cases with a case in which if you treat a consideration as a reason, an evil demon kill you. Even if you satisfy the epistemic norm of rational action with respect to this consideration, it is overall practically wrong to treat it as a reason. In such a case, you have a decisive practical reason to violate the epistemic norm of practical rationality with respect to this consideration. But Hannah and Ernest have no such practical reason to violate the supposed epistemic norm of practical rationality in this particular case. So, RESTAURANT* and ERNEST do not feature cases of justified wrong.

A more natural way of thinking about these cases is the following. Although Hannah and Ernest do not meet the epistemic norm of action with respect to the relevant proposition – so that it is not, according to the knowledge account, a reason they have to act – they have an excuse which explains why they are not blamable.

Gerken's objection

Against such an excuse manoeuvre, Gerken (2011, 537-543) has raised an objection that can be stated as follows:

1. An excuse manoeuvre must not be *ad hoc*.
2. To avoid *ad hocness*, a principled account of excuses is required.
3. In the cases under consideration, if there is an excuse, it must be epistemic.
4. Therefore, a principled account of epistemic excuses is required.
5. There is no satisfying principled account of epistemic excuses currently available.
6. Therefore the excuse manoeuvre is *ad hoc*.

Gerken's efforts are mainly directed at showing that 5 is true. What could be the principled account of epistemic excuses? Taking a clue from the debate concerning the norm of assertion, Gerken considers DeRose's distinction between primary and secondary appropriateness.²⁷ Thus, concerning the knowledge norm of assertion, DeRose writes:

As happens with other rules, a kind of secondary propriety/impropriety will arise with respect to this one. While those who assert appropriately (with respect to this rule) in a primary sense will be those who actually obey it, a speaker who broke this rule in a blameless fashion, by asserting something she didn't know but reasonably thought she did know, would in some secondary sense be asserting properly, and a speaker who asserted something she reasonably thought she did not know, but in fact did know (if this is possible), would be asserting improperly in that secondary sense. An example

²⁷See also Unger (1975); Williamson (2000); Hawthorne and Stanley (2008, 286).

of someone reasonably but falsely thinking they know something is provided by Ginet's famous 'barn case' ..., which can be (and in fact most naturally is) imagined so that the character is reasonable in thinking that he knows that what he's looking at is a barn: Henry quite reasonably thinks that he knows that he's seeing a barn, because he reasonably thinks he is in a very normal situation; unbeknownst to him, he's in a region teeming with fake barns and in fact doesn't know, though in fact he happens to be seeing a real barn right now. In many other cases, a speaker reasonably thinks she knows, though in fact she doesn't, for what she believes is not even true. The knowledge account of assertion would lead us to expect that though such speakers are breaking the rule for assertion, they are blameless for doing so and their assertions are warranted in a secondary way, since they reasonably take themselves to know what they assert. Thus, our sense that such speakers are at least in some way asserting appropriately and are not to be blamed for their assertions does not falsify the knowledge account of assertion, which would lead us to expect just that. (DeRose, 2009, 93-94)

Likewise, we might say that in RESTAURANT* and ERNEST, Hannah and Ernest reasonably believe that they know the proposition on which they act, and hence, are in a secondary sense acting properly, although they violate the primary epistemic norm of action. Satisfying the secondary norm provides Hannah and Ernest with an epistemic excuse.

Against this approach, Gerken raises what I shall call the 'no higher-order belief objection'. Suppose that Hannah and Ernest form no belief as to whether they know the target proposition. Alternatively, suppose that we are considering subjects who do not have the concept of knowledge, and hence, cannot form such beliefs. If the knowledge norm applies in these cases, we cannot explain why these subjects can blamelessly act on a proposition that they do not know by appealing to the fact that they have a rational belief that they know the proposition on which they act. Therefore, this principled account of epistemic excuses is not sufficient to account for all the possible counterexamples to the sufficiency of knowledge.

Next, to try to deal with these further putative counterexamples, Gerken considers the possibility that what excuses the subject has to do with the fact that this subject has propositional justification for the proposition that she knows the target proposition, rather than with the fact that she is doxastically justified in believing the proposition that she knows the target proposition. However, we may think that there is a psychological constraint on propositional justification. For example, we may think that p is justified for S only if S has the psychological capacity to form the belief that p . If so, the problem remains for cases in which the subject lacks the concept of knowledge. On the other hand, if one thinks that there is no psychological constraint on propositional justification, it is unclear how the fact that a proposition is justified for a subject could excuse her. Indeed, it may seem that if believing this proposition is beyond the subject's ability, this proposition cannot excuse her. Thus, Gerken writes:

[I]f propositional warrant is entirely detached from the agent's cognitive life, it hard to see how it excuses *the agent* from breaking the (alleged) knowledge-norm of action... [I]t is not clear whether "impersonal" propositional warrant [i.e., propositional justification without psychological constraint] can excuse the agent in any circumstances. (Gerken, 2011, 541)

If so, the only remaining option for a principled account of epistemic excuses seems to be that of appealing to the idea that in RESTAURANT* and ERNEST, the subject is excused because she rationally (although falsely) believes the proposition that she treats as a reason. However, according to Gerken:

While promising, this approach is little but a collapse of the knowledge account into a warrant account, such as (WA). It is compatible with a warrant approach to hold that the degree of warrant normally required for knowledge is what is normally required for action/practical reasoning. But a central reason for distinguishing between "warrant normally required for knowledge" and "knowledge itself" is this: In certain epistemically abnormal cases, the degree of warrant that would have met the warrant condition on knowledge in normal circumstances is insufficient for knowledge. In some such cases, the warrant may nevertheless be sufficient for rational action/practical reasoning. As mentioned, cases of warranted false belief and Gettier-style cases are paradigms of such epistemically abnormal cases. So, to say that knowledge is normally required for action/practical reasoning but that warranted belief will do in abnormal cases amounts to accepting a version of a warrant account compatible with (WA). (Gerken, 2011, 542)

This passage is a bit unclear. Indeed, it seems clear that according to the knowledge account, a warranted belief will not do in abnormal cases to the extent that it is not a piece of knowledge. So, it is unclear why adopting a warrant account of excuses would lead the knowledge account of reasons to collapse into a warrant account of reasons. For sure, according to the knowledge account and the excuse manoeuvre under consideration, it seems that a warranted belief will do in abnormal cases to excuse the subject. But it will not do to provide the subject with a reason (i.e., to justify her action).

Presumably, Gerken's underlying argument is the following. Given the view of epistemic excuses under consideration, the advocate of the knowledge norm and its opponent should agree that a warrant to believe the target proposition can do a sufficient defensive work to remove the blame in the abnormal cases. But if a warrant to believe can to a sufficient defensive work in the abnormal case, it can do a sufficient defensive work in the normal cases too. Hence, thinking that a sufficient defence in the normal cases require more than a warrant, namely, knowledge, seems spurious.

Reply to the no higher-order belief objection

Although these remarks may cast some doubt on the excuse manoeuvre, they are far from decisive. First, consider the no higher-order belief objection. This objection puts forth that there can be no secondary appropriateness, in DeRose's sense, for subjects who do not possess the concept of knowledge, or who do not entertain a higher-order belief about their knowledge. However, we can weaken in some way our understanding of this notion of secondary appropriateness. Instead of saying that the subject is excused because she rationally believes that she knows the target proposition, we may just say that the subject is excused because she rationally assumes that things are normal, or that she does not know that things are abnormal. (DeRose's quotation above makes it clear that he is focused on whether the subject is rational in assuming that things are normal). Presumably, rationally assuming that things are normal does not require a higher-order belief that things are normal, nor the deployment or the possession of the concept of normality.²⁸ Thus, by weakening our interpretation of what is required by DeRose's secondary appropriateness, we can save this approach from the no higher-order belief objection.

Second, consider Gerken's worry that, to the extent that it is not related in some way to the subject's doxastic state, propositional justification for a certain relevant consideration cannot excuse. The idea underlying this worry seems to be that in order for a consideration to be able to excuse an agent, it must be in some sense psychologically accessible to her, or related to her "cognitive life" in some way. However, why should we think this?

Beyond the distinction between the notions of epistemic and non-epistemic excuses (where the first one has to do with the subject's epistemic position and the second one has to do with something else), we may distinguish between excuses having to do with something peculiar to the subject's environment, and excuses related to something peculiar to the subject's internal condition. As Baron (2005, 390-391) notes:

Excuses thus divide into two rather different groups, the first of which is much more akin to justifications than is the second:

- (1) those (such as duress) that come into play because the situation was such that it was extremely difficult for that actor and would be extremely difficult for most actors to avoid acting wrongly or unlawfully; and
- (2) those (insanity being the paradigmatic instance) that come into play because of some peculiarity about the actor that makes it very difficult for him or her to act as the law requires.

²⁸It seems to me that Gerken should grant this point. Indeed, arguably, S knows that p only if S rationally assumes that things are (epistemically) normal. As animals and infants can have knowledge that p, they can rationally assume that things are epistemically normal. Further a subject can know that p although she has no higher order belief that things are epistemically normal. Note that Gerken uses the term 'presuppose' rather than 'assume'. See Gerken (2013, ch. 2).

Similarly, Botterell (2009, 174-175) distinguishes, on the one hand, duress and necessity, whereby the accused can claim that “she ‘couldn’t have done otherwise’ due to external circumstances beyond her control” (where these external circumstances respectively refer to the action of someone else or to a naturally occurring event), and, on the other hand, insanity, where the defence has to do with conditions that are internal to the subject.²⁹ Let me label the first kind of excuses ‘environmental excuses’ and the second kind of excuses ‘personal excuses’.

The distinctions between, on the one hand, epistemic and non-epistemic excuses, and, on the other hand, environmental and personal excuses, are orthogonal. If an excuse is epistemic, it has to do with a certain relation between a subject and the world. The excuse may be due either to a peculiar feature of the subject, in which case it is a personal epistemic excuse, or to a peculiar feature of the world, in which case it is an environmental epistemic excuse. In both of these cases, there is no reason to think that the exculpating feature should be psychologically accessible to the subject in order to excuse her action. For example, there is no reason to think that if a subject is excused because she is insane, then she has psychological access to the fact that she is insane. Likewise, there is no reason to think that if a subject is excused because things are such that he would, under certain psychological conditions, be justified in believing that he knows the proposition on which she acts, then this excuse is psychologically accessible to her or related to his cognitive life in some way. In particular, if the epistemic excuse is environmental, it need not be in any way related to the subject’s cognitive life.

Briefly stated, the idea that this is *the agent* who is excused, and that she is excused in virtue of her epistemic position, suggests in no way that the excuse is personal, rather than environmental. If the excuse is environmental, it need not be related to the subject’s psychological or cognitive life. Even if the excuse were personal, this would suggest in no way that the agent should be in a position to tell (or think) what it is. Therefore, it is unclear why the fact that impersonal propositional justification is not related to the subject’s cognitive life should cast any doubt on the fact that a subject can be excused because she has impersonal propositional justification with regard to the fact that she knows that p.

Reply to the collapse objection 1: Unequal Defences

Further, it is unclear why a principled excuse manoeuvre appealing to the mere fact that a subject can be defended because she has a warranted belief in the target proposition should collapse into a kind of warrant account of reasons. There is a risk of collapse only if the way in which the subject’s action is defensible in the abnormal case (in which the subject does not know the target proposition) is in no way different from the way in which the subject’s action is defensible in the

²⁹On the basis of such a distinction, Botterell (2009, 174-175) notes that the distinction between justification, excuse and exemption is not always clear. Indeed, if self-defence is a justification (a justified wrong), then it is unclear why necessity is a mere excuse. Further, we might think that insanity is more like an exemption than an excuse. I set this debate aside.

normal case (in which the subject knows the target proposition). In particular, if the defence is stronger in the normal case than it is in the abnormal case, knowledge provides something more, in terms of defence, than what a mere warrant to believe provides. If so, it might be argued on this basis that while knowledge justifies, warrant to believe merely excuses. There would be no collapse of the knowledge account into a warrant account.

There are two reasons to think that there is a normative asymmetry between defences based on pieces of knowledge and defences based on warranted false beliefs. The first one is that while knowledge may be sufficient to defend a certain action in a certain context, a mere warrant to believe could be insufficient to defend this action in this context. To see this, consider the following case (after Hawthorne and Stanley (2008)):

REMARRIAGE. Hannah's husband Mordechai has gone off to war, and goes missing for many years. It would not be reasonable to believe that Mordechai is still alive. Hannah remarries after waiting five years, believing her husband to be dead. As a matter of fact, Mordechai is dead.

In this case, it seems that Hannah knows that Mordechai is dead, and that her remarriage cannot be reasonably criticised. Now, contrast with the following case:

REMARRIAGE*. Hannah's husband Mordechai has gone off to war, and goes missing for many years. It would not be reasonable to believe that Mordechai is still alive. Hannah remarries after waiting five years, believing her husband to be dead. As a matter of fact, however, Mordechai was in captivity. After reemerging from captivity, Mordechai complains to Hannah that she shouldn't have remarried without knowing that he had died.

It seems that, in this case, Mordechai's complain is not unreasonable. And, following Hawthorne and Stanley (2008), we may think that "it is reasonable for Mordechai not to be satisfied with the excuse that Hannah had a justified belief that he was dead." A warrant to believe is typically, or in most cases, a good epistemic defence. However, it is not clear that it is sufficient to provide an excuse in all cases, in particular when being right was particularly important.

The difference between REMARRIAGE and REMARRIAGE* suggests that when an action is defended on the basis of what the subject knows, the defensive work is not done merely by the warrant the subject has to believe the target proposition, but by the consideration that the subject is right, and not right by mere luck. Merely citing the warrant could not have been sufficient to defend the action.

More generally, it seems that there are cases – in particular when the stakes are high – where subtracting the truth while preserving the warrant to believe will affect whether the subject's action is criticisable. If so, that means that a warrant to believe does not always do a sufficient defensive

work. If knowledge can do a sufficient defensive work in these cases but a mere warrant cannot, this suggests that knowledge and mere warranted belief are not on a par, normatively speaking.³⁰

The second argument in favour of the asymmetry between defences based on knowledge and defences based on mere warrant to believe is that there is still something that seems wrong in the abnormal case in which the subject does not know the target proposition, like in ERNEST or RESTAURANT*, while nothing seems wrong in the normal version of the case. Let us return to the debate on the knowledge norm of assertion and consider what DeRose writes on this issue:

It would also be trouble for the knowledge account if we had a clear sense that there is nothing at all amiss in a speaker's asserting something she does not know in the cases where she reasonably takes herself to know what she asserts. But I certainly lack any such clear sense; in fact, I instead have a fairly clear sense that something is amiss in such cases – though, of course, the exact nature of what has gone wrong is open to a lot of negotiation. (DeRose, 2009, 94).

In brief, the idea is that when you assert something that you do not know, there is still something that seems wrong even if we grant that you reasonably believe that nothing is wrong. For example, suppose that you realise that what you said is, in fact, false. It would not be unnatural for you to retract your assertion by saying “I was wrong, I should not have said this”. This retraction would be difficult to explain if a mere warrant to believe were enough to justify your assertion.

Likewise, the advocate of the knowledge norm of action may insist that there is a normative asymmetry between the abnormal case and the normal case because there is still something that seems wrong in the former while nothing seems wrong in the latter. One way of showing this normative inequality consists in focusing on the different normative consequences of the actions performed in the two cases. In particular, wrongs blamelessly committed on the basis of reasonable but false beliefs may call for stronger duties of reparation than wrongs blamelessly committed on the basis of knowledge. If reasonable but false beliefs defend wrongs in the same way as pieces of knowledge, however, it is unclear why wrongs blamelessly committed on the basis of the former could call for stronger duties of reparation than wrongs blamelessly committed on the basis of the latter.

To illustrate, consider the following case, from Littlejohn (2012, 210-212):

LOAN SHARK. Harry is behind in his payments to Bobby, the loan shark. Bobby gave Harry a severe beating last week and a warning that if he missed another payment, Bobby would kill him. The payment is due today and Harry does not have the money. He borrows a revolver and hangs out at Audrey's restaurant hoping that Bobby will

³⁰Fantl and McGrath (2009, 96-100) use a subtraction argument to show that knowledge is not necessary for rational action and that knowledge-level warrant is sufficient. However, if there are cases where merely subtracting truth changes whether the subject's action is criticisable then, given that justification entails lack of criticisability, subtracting truth can affect whether the subject's action is justified.

leave him alone in public. He is shocked when he sees a man with a menacing look he takes to be Bobby come in and walk straight towards him. Harry says “I won’t let you get me, Bobby!” and he pulls out his revolver and takes aim.

Audrey has a pipe. She knows that the man Harry takes to be Bobby is really Bobby’s twin brother. She knows that while he might look dangerous, he is a threat to no one. She knows that to stop Harry from firing at Bobby’s twin, she will need to club him with her pipe. She does so, intervening on behalf of Bobby’s brother.

Leo sees Audrey swinging her pipe at Harry. Leo believes that Bobby’s twin is Bobby and believes that Audrey must be helping the loan shark kill Harry. Leo grabs a pipe and swings it at Audrey on the reasonable but mistaken belief that he is helping to defend Harry from a loan shark and an accomplice. Leo connects and knocks Audrey unconscious, but only after Audrey connects and knocks Harry unconscious.

At the end of the story we have two unconscious subjects. Harry was knocked unconscious by Audrey because she knew that Harry was going to shoot an innocent person if she did not intervene. Leo knocked Audrey unconscious just as she was striking Harry.

After the police come, Leo sees Bobby and Audrey begin to stir. There is enough pain reliever to help one of these subjects but not enough to help both.

As Littlejohn (2012, 212) notes, it seems that “[g]iven just enough pain reliever to help one, [we] ought to assist Audrey rather than Harry.” This is strange if Leo’s action on Audrey is as justified as Audrey’s action on Harry. Indeed, if Leo’s action and Audrey’s action are normatively on a par, it is unclear why our duty to repair the (supposedly justified) wrong committed by Leo on Audrey is stronger than our duty to repair the (uncontroversially justified) wrong committed by Audrey on Harry. In contrast, on the assumption that the wrong resulting from Leo’s action is not justified but merely excusable, while the wrong resulting from Audrey’s action is not merely excusable but justified, we can do justice to our normative intuitions.

As a result, it is unclear why accepting the idea that a warrant to believe that *p* can excuse treating *p* as a reason would put oneself at risk of collapsing the knowledge account into a kind of warrant account. If there were no normative difference between a defence in terms of mere warranted beliefs and a defence in terms of knowledge, then, surely, such a collapse would seem almost unavoidable. However, it is clear that there are ways of showing that there is a normative difference.

Reply to the collapse objection 2: Rationality-based account of reasons, warrant account and knowledge account.

Of course, there is something in common between the subject who acts on the epistemically reasonable but false beliefs in the abnormal situation, and the subject who acts on what she knows in

the normal situation. They are both, in some sense of the term, ‘rational’. In particular, we may say that they are both rationally coherent. But does that mean that they are normatively on a par? And can the advocate of the warrant account plausibly say that this is this shared property of rational coherence that can make both of these subject justified? Whether reasonable but false beliefs provide justifications or mere excuses raises the fundamental question whether reasons should be grounded in rational coherence.

According to the rationality-based account of reasons:

RBR. p is a reason for S to φ if and only if rationality requires (/it is coherent for) S to treat p as a reason to φ ³¹

This means that if p is a reason for S but S does not treat p as a reason, then S is irrational or incoherent.

If we assume a rationality-based account of reasons, it makes sense to think that when you reasonably believe that p is the case, or when you reasonably believe that you should treat p as a reason to act, then p can justify you. Indeed, one might think that if you reasonably believe that p is the case, or if you reasonably believe that you should treat p as a reason to act, rationality requires you to treat p as a reason to act. According to the rationality-based view of reasons, if rationality requires you to treat p as a reason to act, then p is a reason to act. As a result, since reasons can justify, if you reasonably believe that p or that you should treat p as a reason to act, p can justify you.

However, RBR is controversial. As noted in chapter 7, it is unclear that rational requirements provide us with reasons. First, they have wide scope. Second, the claim that reasons are grounded on the subject’s attitudes faces a bootstrapping objection. It is problematic to think that a subject may have a reason in virtue of attitudes that she has if she has no independent reasons to have these attitudes in the first place.³²

In addition, it does not seem that an advocate of the warrant account can adopt RBR. Indeed, it is compatible with the warrant account that you may have a warranted but false belief that p is warranted enough for you to act on p . In this situation, it is rationally coherent for you to treat p as your reasons to act. However, as your belief that p is warranted enough for you to act on p is false (although warranted), that does not make it the case that p is warranted enough for you to act on p . Hence, that does not make it the case that p is a reason that can justify you.

Presumably, the advocates of the knowledge norm will also reject the rationality-based view of reasons. When they say that knowledge is the norm of practical rationality, they do not have *structural* rationality in mind, but *substantive* rationality, where substantive rationality is understood in terms of appropriate responses to reasons, and not in terms of rational coherence among attitudes.

³¹See Scanlon (2014, 7-14).

³²See, e.g., Scanlon (2014, 98-100).

If we reject RBR, we may say that it is substantively irrational to act on p when one does not know that p , although, to the extent that one reasonably believes that p , it may be coherently rational to act on p . We may then say that the first notion of rationality has to do with justification and reasons, and the second notion of rationality has to do with excuses and what we could expect a rational subject would do in the circumstances.

Thus, the advocate of the knowledge norm may grant that it would be rationally incoherent for Ernest not to treat the proposition that there is a lion as a reason, and for Hannah not to treat the proposition that the restaurant is down the street as a reason to go there. However, that does not mean that by treating these things as reasons, Ernest and Hannah respond to reasons. After all, Ernest is not sensitive to what he takes to be his reason, namely, that there is a lion. Were there no lion, he would still believe that there is a lion. Likewise, Hannah is not sensitive to what she takes to be her reason, namely, that the restaurant is down the street. Given this absence of sensitivity, it is difficult to see them as *responding* to reasons, and hence, it is difficult to see them as substantively rational.

In brief, the criticism of the excuse manoeuvre appealing to a collapse seems to rely on the fact that in the normal and abnormal cases, the subjects are on a par with respect to rationality. However, the notion of rationality that is appealed to is that of rational coherence. If so, the criticism of the excuse manoeuvre appealing to a collapse is compelling only if some sort of rationality-based account of reasons is assumed. Still, we have seen that both the advocates of the knowledge norm of action and the advocate of the warrant account are likely to reject RBR.

Finally, if the criticism of the excuse manoeuvre appealing to a collapse could be put to work against the knowledge norm, it could be put to work against the warrant account as well. For, again, given the warrant account, it has to be possible for abnormal cases to arise in which the subject is warranted in believing that p is warranted enough for her to act on p , while this normative belief is false. If this is possible, the advocate of the warrant account will have to appeal to an excuse manoeuvre to explain why the subject in such an abnormal situation is not blamable when she treats p as her reason to act while she is not warranted in doing so. But if appealing to an excuse manoeuvre in the cases of false warranted normative beliefs is acceptable, then it is unclear why appealing to an excuse manoeuvre in the cases of false warranted factual beliefs would not be acceptable.

8.4.2 The practical cases and the excuse manoeuvre

So far, I have considered cases in which less than knowledge seems intuitively sufficient to defend the subject's action because the subject is warranted in believing the target proposition (and may rationally assume that things are normal). I have defended the idea according to which such cases do not compromise the knowledge account of practical reasons because we may say that such defences

do not constitute justifications, but epistemic excuses.

There are also cases where the subject acts on less than knowledge although she is not even warranted in believing the target proposition. Thus,

[c]onsider someone who has a .4 credence that it will rain, and on those grounds brings an umbrella on their walk. According to the Reasons-Knowledge Principle, they are subject to criticism if they fail to know whether there is a chance that it will rain. But surely this person is rational, whether or not she acquires enough justification for her belief that there is a chance that it will rain. Perhaps she had to rush out the door, and didn't have time to actually come to know what the epistemic chance of rain was. Surely, such a person is not subject to criticism. (Hawthorne and Stanley, 2008).

In this case, the subject is motivated by (the content of) a mental state that does not amount to knowledge (namely, a .4 credence that it will rain). Still, it seems that, given her practical situation, the subject should not be blamed for this.

Can we say that, in such a case, the epistemic norm of action is overridden by another norm so that the action is still, overall, justified?³³ As Hawthorne and Stanley note, we cannot because the epistemic norm of action “is a norm governing action generally. So it is not an option that it can be trumped by norms governing action in some more general sense”. Indeed, if knowledge that *p* is necessary for *p* to be a reason justifying the action, it cannot be the case that a subject is justified by *p* in doing *A* while not knowing that *p*. Thus, if the subject merely have credence .4 that it will rain but does not know that there is a chance that it will rain, the fact that she has credence .4 that it will rain does nothing to justify her action (although it can explain it). However, again, we may say that the subject is excused for taking her umbrella without (justifying) reasons.³⁴

According to Gerken, however:

the list of practical matters that may render it practically rational to act on less than knowledge, and thus the need for excuses, does not stop at urgency. If the stakes are very low, it may be practically rational to act on less than knowledge that *p*. Likewise, it may be rational to act on less than knowledge that *p* if the putative benefits of acting are far greater than the costs of acting (as may well be the case when bringing an umbrella). So, it appears that there will be an abundance of cases in which someone may be excused from violating the knowledge norm. This suggests that the knowledge account will result

³³Regarding assertion, for example, Williamson (2000, 256) writes: “Sometimes one knows that one does not know *p*, but the urgency of the situation required one to assert *p* anyway. I shout “That is your train”, knowing that I do not know that it is, because it probably is and you have only moments to catch it. Such cases do not show that the knowledge rule is not the rule of assertion (...) they merely show that it can be overridden by other norms not specific to assertion. The other norms do not give me warrant to assert *p*, for to have such warrant is to satisfy the rule of assertion.”

³⁴Compare with Skorupski's FIRE ALARM case in chapter 7 section 2.7.

in a proliferation of excuses. Indeed, it stands in danger of multiplying excuses beyond plausibility. (Gerken, 2011, 544)

For example, Gerken (forthcoming, 6.3.c) considers the following case:

KICKOFF. S believes that the game has started. But the only basis for her belief is that she vaguely remembers a stranger telling her the time of the kickoff in the bar the night before. But both S and the testifier were tipsy, and the fellow didn't seem all that reliable anyhow.

Then, Gerken (forthcoming, 6.3.c) writes:

In KICKOFF, S's testimonial warrant is arguably below the degree required for knowledge. Yet it is epistemically permissible for S to act on her belief that the game has started by turning on the television (...). KICKOFF provides a novel challenge to the idea that knowledge that p is a necessary condition for acting on p. The novelty lies in the structural feature that the costs of acquiring novel evidence are higher than the costs of acting on a false belief.

In reply to this line of criticism, let me note three things. First, it seems that excuses are precisely the kind of things that, we should grant, can proliferate in principle. It is not particularly controversial that there are many kinds of excuses, and that a certain consideration can work as an excuse in some situations but not in others. In contrast to norms, that often have a certain degree of generality, excuses are most of the time very flexible. For example, it is plausible that the fact that you are warranted in believing that p although p is false can excuse you for acting on p in some cases, but not in all cases.. Thus, it is not particularly surprising that excusability is sensitive to the subject's practical situation.

Second, it is unclear why assuming that many different things can excuse will lead to the idea that there will be abundance of cases in which we are excused rather than justified in our actions. After all, we know many things, and in most cases, we need not act on things we do not know. In a case in which Gerken (forthcoming, 8.3.c)) the stakes are very low and one does not know that p, for example, we may grant that it may be practically rational to act as if p, that is, to act on the probability that p. However, such cases do not show that it would be practically rational to treat p – rather than the fact that p is probable – as a reason to act.

Thus, it is not clear that KICKOFF provides a challenge to the necessity of knowledge. Indeed, there are different ways of understanding the expression 'act on her belief that the game as started'. Suppose that S is epistemically rational. Thus, given the evidence she has, she does not have a full belief that the game has started, but merely a certain degree of belief, or a belief that there is a chance that the game has started. If by "acting on her belief" Gerken means "acting on her rational degree of belief" or "acting on the fact that there is a chance that the game has started", the

advocate of the knowledge account can grant that the subject is warranted in doing this, because she knows that there is a chance that the game has started. Now, suppose that S has a full belief that the game has started. That means that S is epistemically irrational, because the warrant she has for believing that the game has started is below the threshold for rational belief. But, in this case, it is unclear how the subject could be epistemically rational in acting on her belief, because this belief is irrational in the first place.

Third, the idea that practical factors can make treating something as a reason epistemically appropriate is in itself problematic. To see the problem, suppose that the stakes are low and Keith does not have enough evidence to know that the bank is open. For example, suppose that he has been told that the bank is open two Saturdays over three. Now, suppose that he decides to come back on the basis of the fact that the bank is open next Saturday. Is he practically rational? It does not seem. The mere fact that the stakes are low does not give him (the fact that) the bank is open next Saturday as a reason to act. To see this, consider the following dialogue:

Keith's wife - Why do you prefer to come back tomorrow?

Keith - Because the bank is open tomorrow.

Keith's wife - But you do not know that the bank is open this Saturday, right?

Keith - I do not know that the bank is open this Saturday. But look, since the stakes are extremely low, that the bank is open this Saturday is a good reason to come back this Saturday.

Keith's wife - ??

It seems that when the evidence is insufficient to rationally believe the target proposition, practical considerations cannot make the evidence good enough for treating it as a reason. (Of course, in this case, it might be rational for Keith to act on the assumption that the bank is open, or to act as if the bank is open, but this is a different matter. That does not mean that Keith can appropriately be motivated by the proposition that the bank is open. If asked why he wants to come back, it would be inappropriate for him to say "Because the bank is open"). This point is related to the principled argument for the necessity of knowledge. If knowledge that p were not always necessary for rationally treating p as a reason, we could easily take ourselves to act for the reason that p while taking ourselves not knowing that p, provided that it is clear that the practical context requires less than knowledge. However, as this dialogue illustrates, it does not seem coherent to act for something we take ourselves not know, even if the stakes are extremely low.

In sum, the warrant account predicts that practical factors can make us justified in treating something we do not know as a reason to act and that they can make us justified in treating something we know we do not know as a reason to act. But if we can be justified in treating p as a reason while taking ourselves as not knowing p, it is unclear why what Keith says seems incoherent.

Finally, note that, as the warrant account has it, if we grant that knowledge is not always necessary, then it becomes difficult to explain why we have the intuition that reasoning on the basis of lottery propositions seems wrong because these premises are not known.

8.4.3 Criticism for not acting on p (weaker)

By distinguishing excuse and justification, the advocate of the knowledge norm of action can explain why we sometimes use weaker terms than knowledge to defend a subject acting on something she does not know, and why subjects who do not act on what they know are not always blamable. In particular, I have appealed to the distinction between substantive and structural rationality, and stressed that the knowledge norm of action should be understood as a norm of substantive rationality, while excuses may arise from the fact that the subject can be seen as structurally rational.

The distinction between substantive and structural rationality can also help to see why we sometimes use terms weaker than ‘know’ to criticise a subject refraining from acting on a certain proposition. Thus, consider the following case (after Brown (2008, 171)):

POTATOES. I leave my partner a note saying we are out of potatoes. My partner returns home from work having stopped in the grocery store to buy beer, but doesn’t get any potatoes. It turns out that we do have potatoes since, by chance, a neighbour happened to bring some over from her garden as a surprise for us. I may criticise his action saying, “You should have got potatoes. You thought we didn’t have any”.

Clearly, it is not rationally coherent for the partner not to buy potatoes if he thinks that he should get some. However, it is unclear that the criticism targets what he should have done in view of the reasons that apply to him. When we say “You should have got potatoes” in this situation, it is disputable that the ‘should’ is normative. And, indeed, suppose that the partner has in fact bought potatoes rather than beer, and realise that the neighbour has bring enough potatoes. It would not be unnatural for him to say: “I should not have got potatoes. I should have bought beer instead”. Therefore, unless we adopt a rationality-based account of reasons, such cases do not threaten the knowledge account of practical reasons.

8.5 Conclusion

In this chapter, I have presented the main arguments in favour of the claim that knowledge is the epistemic norm of action. In particular, I have defended the claim that knowledge is always necessary for action by defending the excuse manoeuvre against recent attacks. I have also argued against the warrant account. First, I have argued that the main attempt to date to explain the

linguistic data without appealing to a normative connection between knowledge and action is not fully satisfying, and, in particular, that it is in tension with the warrant account itself. Second, I have stressed that the warrant account does not do justice to our intuition underlying the wrong practical reasoning argument and the principled argument for the necessity of knowledge. The claim that knowledge is necessary for rational action does not threaten the orthodox account of knowledge but it is incompatible with the warrant account. I conclude that orthodox philosophers should reject the warrant account and adopt the claim that knowledge is necessary for rational action. In the next chapter, I shall consider whether knowledge is sufficient for rational action.

Chapter 9

The context-sensitivity of the norm of action

9.1 Introduction

In the previous chapter, I have argued that knowledge is always necessary to satisfy the epistemic norm of action, and hence, that we should reject the warrant account. The warrant account is a shifty account of the epistemic norm of action, but that was not the problem. In this chapter, I argue that such a shifty approach is right. It follows that an orthodox philosopher – i.e., a philosopher rejecting a shifty view of knowledge and ‘know’ – should reject the claim that knowledge is the epistemic condition that we must satisfy to satisfy the epistemic norm of action. Or, as we can put it, she should reject the claim that knowledge is the epistemic norm of action.

I start by presenting the paradox that the orthodox philosopher faces if she grants that knowledge is necessary and sufficient for action. This paradox relies on the uncontroversial claim that our epistemic evaluations of actions are sensitive to the subject’s practical context. But if these epistemic evaluations are assessments of whether the epistemic norm of action is satisfied, and knowledge is necessary and sufficient for action, then the orthodox view of knowledge is wrong.

However, I note that, in principle, it might be that these shifty epistemic assessments are not evaluations of whether the epistemic norm of action is satisfied. Thus, an available option for the orthodox philosopher is to say that the paradox is merely apparent. On this approach, the sensitivity of our epistemic evaluations to the subject’s practical context would not concern the epistemic norm of action, but something else.

Still, even if there are epistemic evaluations that are not evaluations of whether the epistemic norm of action is satisfied, this does not imply that the epistemic norm of action is insensitive

to practical factors. Thus, a champion of the insensitivity of the epistemic norm of action must provide an explanation why our overall epistemic evaluations of action shift, without compromising her initial assumption.

I consider the main attempt of this kind, which has been proposed by Williamson (2005). A way of presenting this approach is to distinguish an epistemic requirement insensitive to practical factors, which comes from a primary epistemic norm of action, and an epistemic requirement sensitive to practical factors, which comes from a derivative norm. Both of these epistemic requirements are relevant to the overall epistemic evaluation of action. I put forth two arguments against this approach, though. First, I argue that even if we grant the general distinction between primary norms and derivative epistemic norms, using this distinction is ill-motivated when it comes to the epistemic norm (or appropriateness) of action. Second, I argue that whatever we think of this distinction, it does not address the relevant cases properly: Williamson's approach relies on a wrong diagnosis.

Next, I consider whether the distinction between having a reason and having a sufficient reason can be used to defend the insensitivity of the epistemic norm of action. I argue that even if this approach might be useful in other contexts, it cannot account for the standard pairs of cases. In particular, on a natural understanding of what the epistemic norm of action governs, the bank cases show that the epistemic norm of action is sensitive to the stakes.

Finally, I offer a second argument in favour of the sensitivity of the epistemic norm of action. According to this argument, insensitivism about the epistemic norm of action leads to endorse a maximal-standard account of the epistemic norm of action, which is an implausible account.

9.2 The pragmatic encroachment paradox and the context-insensitivity of the epistemic norm of action

9.2.1 The pragmatic encroachment paradox

Consider the three following claims:

PRAGMATIC ENCROACHMENT PARADOX

(Sensitivity Thesis) The epistemic requirement for rationally treating p as a practical reason is sensitive to the subject's practical situation.

(Orthodoxy) Knowledge is insensitive to the subject's practical situation.

(Knowledge Norm of Action) Knowledge is the epistemic norm of action.

These three claims are, it seems, incoherent. If we say that the epistemic requirement necessary and sufficient for rationally treating something as a reason is sensitive to the practical context, and

we say that knowledge is the epistemic condition necessary and sufficient for rationally treating something as a reason (or, in short, that knowledge is the epistemic norm of action), then it follows that knowledge is sensitive to the practical context.

Can we reject the sensitivity thesis? The sensitivity thesis relies on our assessments of particular cases or, as Williamson puts it, on “our ability to discriminate in particular cases between silly and sensible decision-making” (2005, 228). These assessments show that “surely what it is appropriate to use as a premise for practical reasoning at least partly depends on how much is at stake, not just on purely epistemic matters of the traditional kind” (ibid.).

In an orthodox framework, the sensitivity thesis is compatible with different views about the relation between knowledge and the epistemic norm of action. An orthodox philosopher may think that knowledge is merely necessary, or merely sufficient, or neither necessary nor sufficient. Thus, two different types of cases have been proposed in support of the sensitivity thesis. First, some have proposed cases in which a very low epistemic position is sufficient to act on the target proposition. Such cases are used, for example, by the advocates of the warrant account. Recall KICKOFF:

KICKOFF. S believes that the game has started. But the only basis for her belief is that she vaguely remembers a stranger telling her the time of the kickoff in the bar the night before. But both S and the testifier were tipsy, and the fellow didn’t seem all that reliable anyhow. (After Gerken (forthcoming, 6.3.c))

According to those philosophers, in KICKOFF, it is epistemically rational for S to act on the proposition that the game has started (act for this reason, treat this proposition as a reason to act, use it as a premise, etc.) by turning the TV on, for the cost of error is very low.

Now, if we also assume that knowledge is the norm of action, this case is sufficient to threaten the orthodox account of knowledge. Indeed, it seems that we are thereby led to say that sometimes a very weak epistemic position is sufficient to know that p and sometimes, in more normal cases, a stronger epistemic position is necessary to know that p. If so, the epistemic standards for knowledge are sensitive to the the practical context and pragmatic encroachment is true.

However, this way of defending the sensitivity thesis – on the basis of the idea that sometimes satisfying very low epistemic standards is sufficient to act on the target proposition – is more coherent if we want to provide an argument against the knowledge norm of action than if we want to provide an argument in favour of pragmatic encroachment. Indeed, it is implausible that the subject in KICKOFF knows the target proposition. So, thus based, the sensitivity thesis threatens the claim that knowledge is the norm of action rather than the orthodox account of knowledge. If this is right, cases like KICKOFF cannot be used to support the premise of an (alleged) paradox threatening the orthodox account of knowledge.

Further, I have argued in the previous chapter that it is unclear that, in cases like KICKOFF, the subject can rationally act on the proposition (rather than on the probability that the proposition

is true). If so, cases like KICKOFF are insufficient to show that the epistemic requirement for rationally treating p as a practical reason is sensitive to the subject's practical context.

Second, we might want to defend the sensitivity thesis on the basis of cases in which a very strong epistemic position is necessary to act on the target proposition. Consider the high stakes bank case. Suppose that Keith is not troubled by his wife's worries. He decides to come back without checking. Everyone agrees that the fact that the bank is open cannot rationalise his action. In contrast, in the low stakes case, it is clear that this fact can rationalise this action. As Fantl and McGrath write:

When more rides for you in your decision depending on whether p , it is to be expected that you would pay closer attention to possibilities of error you would normally ignore, and rightly so. Few would deny that practical environment with respect to p can affect how strong your epistemic position must be in order to act on p . This is a point long stressed by decision theorists. (Fantl and McGrath, 2009, 32-33)

Thus, philosophers generally argue in favour of the sensitivity thesis by contrasting normal cases and cases in which a very strong epistemic position is necessary – rather than by contrasting normal cases and cases in which a very weak epistemic position is (allegedly) sufficient. In contrast to cases like KICKOFF, in which it is implausible that the subject knows the target proposition, it is not implausible that the subject does not know the target proposition in the high stakes case. If we defend the sensitivity thesis on the basis of cases in which a very strong epistemic position is required, this thesis is liable to raise a tension between the orthodox account of knowledge and the sufficiency of knowledge for rational action. The threat of pragmatic encroachment comes from the sufficiency direction of the knowledge norm of action. We can thus make things more precise in the following way:

PRAGMATIC ENCROACHMENT PARADOX*

(Sensitivity Thesis)* The epistemic requirement for rationally treating p as a practical reason is sensitive to the subject's practical situation so that a particularly strong epistemic position with respect to p is sometimes required.

(Orthodoxy) Knowledge is insensitive to the subject's practical situation.

(Necessity of knowledge) Knowledge is necessary to satisfy the epistemic norm of action

(Sufficiency of knowledge) Knowledge is sufficient to satisfy the epistemic norm of action

By distinguishing clearly the sufficiency direction and the necessity direction of the knowledge norm, we make it clear that there are three ways of rejecting the knowledge norm of action. First, following the advocates of the warrant account, we can reject the knowledge norm of action by rejecting the necessity and the sufficiency of knowledge. This approach is likely to support the sensitivity thesis

with cases in which a very weak epistemic position is (allegedly) sufficient to act on the target proposition (e.g. KICKOFF) and with cases in which a very strong epistemic position is necessary to act on the target proposition (e.g. the high stakes bank case). I have already considered and rejected this option in the previous chapter.

Second, we can choose to reject the knowledge norm of action by rejecting the necessity of knowledge while keeping the sufficiency of knowledge. This approach may support the sensitivity thesis by using cases in which a very weak epistemic position is (allegedly) sufficient to act on the target proposition, but this is not obligatory. We might think, for example, that in the low and high stakes cases, although Keith's epistemic position with respect to the fact that the bank is open is quite strong, it is not sufficient for knowledge. We can then explain why it is epistemically appropriate for Keith to act on the fact that the bank is open in the low stakes case by saying that knowledge is not necessary, and, given the sensitivity thesis, that only in the low stakes case is Keith's epistemic position sufficient.

I have argued in the previous chapter that knowledge is always necessary, so this option is not open to us. But there is another independent reason to reject this option. The main problem, indeed, is that this approach seems to lead to scepticism. If knowledge is always sufficient for rational action and the standards for knowledge do not shift, this means that the epistemic standards for knowledge are extremely high. Thus, the pragmatic encroachment paradox is often constructed with claims similar to the followings:¹

PRAGMATIC ENCROACHMENT PARADOX**

(Sensitivity Thesis)** The epistemic requirement for rationally treating *p* as a practical reason is sensitive to the subject's practical situation so that a (close to) maximally strong epistemic position with respect to *p* is sometimes required.

(Orthodoxy) Knowledge is insensitive to the subject's practical situation.

(Non-Scepticism) Knowledge does not require a (close to) maximally strong epistemic position.

(Sufficiency of knowledge) Knowledge is sufficient to satisfy the epistemic norm of action.

(Sensitivity Thesis)** is based on the idea that the cost of error can be raised so high that it would not be rational for the subject to act on the target proposition unless she has, say, Cartesian certainty about the target proposition. Now, if (Sufficiency of knowledge) is true, this means that either (Orthodoxy) or (Non-Scepticism) is false. The view under consideration grants (Orthodoxy). Hence, if the orthodox philosopher is a non-sceptical philosopher, it seems that she should reject the option according to which knowledge is not necessary but always sufficient for rational action.

¹See in particular Fantl and McGrath (2009). Note that Fantl and McGrath use a claim weaker than (Sufficiency of knowledge), namely, the claim that knowledge-level justification is sufficient.

A third possibility is to reject the claim that knowledge is the epistemic norm of action by rejecting merely the sufficiency of knowledge, while keeping the necessity of knowledge. This approach grants that knowledge does not require to satisfy (close to) maximal epistemic standards – hence, it does not lead to scepticism – but it grants that knowledge is always necessary. I will defend this option in the next chapter.

9.2.2 The equivalence thesis

In the remaining of this chapter, I shall consider a fourth way of responding to these paradoxes. Instead of rejecting the claim that knowledge is the norm of action, this approach has it that the paradoxes are merely apparent. This approach is based on the idea that we should distinguish between different types of epistemic assessments. Indeed, perhaps the claim that knowledge is the epistemic condition that we must satisfy in order to satisfy the epistemic norm of action (or, in short, the claim that knowledge is the norm of action), and the claim that knowledge is insensitive to practical factors, do not together entail that the overall epistemic evaluation of action should not shift with practical factors. Perhaps the overall epistemic requirement, upon which we base our overall epistemic evaluation, is not merely determined by the epistemic norm of action. In other words, the paradoxes rely on the following claim:

(Equivalence Thesis) The epistemic position that is, overall, required with respect to p for being rational to treat p as a practical reason is equivalent to the epistemic position required by the epistemic norm of action.

Suppose that the equivalence thesis is false. Then, even if you satisfy the epistemic norm of action with respect to p , that does not mean that, overall, your epistemic position is good enough to act on p .

It is dubious that the equivalence thesis is true. Consider the following case:²

JUDGE. John is a judge assessing Bill's culpability. An evil demon kills John if John sentences Bill on the basis of non-entailing evidence. In fact, Bill is guilty, and John acquires more and more evidence that Bill is guilty.

Arguably, at some point, that Bill is guilty is sufficiently supported by John's evidence for this proposition (or belief) to be a reason John has to φ , for any φ . Hence, at this point, this proposition (or belief) is a (pro tanto) reason John has to sentence Bill. Still, it would not be rational for John to sentence Bill for the reason that Bill is guilty. Indeed, the evidence that John has is not entailing. If John sentences Bill on the basis of non-entailing evidence, John will be killed. We may say that John does not have the epistemic position that is, overall, required with respect to Bill's culpability

²See also chapter 7 section 4.1.

to treat the proposition (or belief) that Bill is guilty as a sufficient reason to sentence Bill. However, that does not mean that John does not satisfy the epistemic norm of action with respect to Bill's culpability. That Bill is guilty is a (sufficient) reason John has to do other things, for example, to make someone else sentencing Bill. The problem is just that, given John's practical situation, this reason is not sufficient for him to sentence Bill.

This case illustrates a situation in which a subject meets the epistemic norm of action with respect to p although she does not meet the overall epistemic requirement for rationally treating p as her reason for a particular course of action. If we reject the idea that the *overall* epistemic requirement is equivalent to the *specific* epistemic requirement issued by the epistemic norm of action, it is possible that the overall epistemic requirement shifts even if the epistemic requirement issued by the epistemic norm of action do not shift. This may happen if other norms deliver other epistemic requirements.

Rejecting the equivalence thesis is a first step if one wants to maintain that knowledge is the epistemic norm of action in an orthodox framework. Still this is not sufficient. It has also to be shown that, even granted that the equivalence thesis is false, our judgements about silly and sensible decision-making do not compromise the idea that the epistemic norm of action is insensitive to the practical context.

9.3 Insensitivism about the epistemic norm of action

9.3.1 The two norms approach

Primary and secondary appropriateness, primary and derivative norms

Williamson (2005) is the best attempt to date to exploit the falsity of the equivalence thesis in order to defend the knowledge norm of action in an orthodox framework. To present his approach, let me start by some terminological remarks. I characterize Williamson's approach as a 'two norms' approach. However, in this paper, Williamson does not use the word 'norm' but, mainly, the word 'appropriate'. Further, in this paper, he does not use the terms 'primary' and 'secondary' appropriateness or norms. Hence, we might think that my characterization of Williamson's view is wrong, or misleading.

It is clear that Williamson's use of 'appropriate' is normative since, according to him, appropriateness has to do with entitlement and is contrasted with excuses.³ An action or attitude is appropriate when it conforms to a norm. Also, although Williamson does not explicitly distinguish two norms, the point of this paper is to distinguish two questions: that of whether something – a premise – is (normatively) appropriate, and that of whether the subject knows whether this thing is (normatively) appropriate. Thus, according to Williamson, we have two kinds of evaluation:

³See Williamson (2005, 227).

Suppose that an agent relies on the premise q in otherwise impeccable practical reasoning, and (in her context) q is indeed appropriate, but she is in no position to know that q is appropriate. Then our assessment of her reasoning should be ambivalent: it is good that she used an appropriate premise, bad that she was in no position to know that it was appropriate. (Williamson, 2005, 229-230)

This distinction seems at first sight very close to DeRose's distinction between primary and secondary appropriateness presented in chapter 8 (section 4.1.). It is in fact significantly different. DeRose's distinction between primary and secondary appropriateness is closer to the distinction between warrant and reasonableness that Williamson proposes when he discusses the knowledge norm of assertion in *Knowledge and its limits*:

In other cases, one reasonably but falsely believes p , and is in no position to know that one does not know p (...). One cannot discriminate between one's actual circumstances and circumstances in which one would know p . For example, it is winter, and it looks exactly as it would if there were snow outside, but in fact that white stuff is not snow but foam put there by a film crew of whose existence I have no idea. I do not know that there is snow outside, because there is no snow outside, but it is quite reasonable for me to believe not just that there is snow outside but that I know that there is; for me, it is to all appearances a banal case of perceptual knowledge. Surely it is then reasonable for me to assert that there is snow outside. The case is quite consistent with the knowledge account. Indeed, if I am entitled to assume that knowledge warrants assertion, then, since it is reasonable for me to believe that I know that there is snow outside, it is reasonable for me to believe that I have warrant to assert that there is snow outside. If it is reasonable for me to believe that I have warrant to assert that there is snow outside, then, other things being equal, it is reasonable for me to assert that there is snow outside. Thus the knowledge account can explain the reasonableness of the assertion. However, granted that it is reasonable for me to believe that I have warrant to assert p , it does not follow that I do have warrant to assert p . The term 'warrant' has been reserved for the property C in the rule C of assertion. There may be other evidential norms for assertion, if they can be derived from the knowledge rule and considerations not specific to assertion. The reasonableness of asserting p when one reasonably believes that one knows p has just been derived in exactly that way. (Williamson, 2000, 256-257)

In brief, primary appropriateness is a matter of conformity to the norm (having warrant), and secondary appropriateness is a matter of reasonable belief that one conforms to the norm (that one has warrant). If there is secondary appropriateness without primary appropriateness, the subject has an excuse.

It is clear that secondary appropriateness thus understood cannot be identified with knowledge of primary appropriateness, and the absence of secondary appropriateness cannot be understood as the absence of knowledge of primary appropriateness. For secondary appropriateness does not require primary appropriateness, whereas knowing that primary appropriateness obtains requires that primary appropriateness obtains. In other words, the claim that it is bad not to know that primary appropriateness obtains is different from the claim that it is bad not to reasonably believe that primary appropriateness obtains. How can we understand the difference?

I suggest to understand the difference as follows. Whereas a subject who knows that primary appropriateness obtains can be seen as following the norm (successfully), a subject who merely reasonably believes that primary appropriateness obtains (and hence, is merely secondarily proper) should be seen as merely *trying* to follow the norm. It is plausible to think that there is something bad if we are not following the norm but just *trying* to follow it (unsuccessfully). Thus, when an action or attitude is *merely* secondarily appropriate, the subject has an excuse, but two things are still bad. First, it is bad that the action or attitude is primarily inappropriate. Second, it is bad that the subject is not in fact following the norm, but is just trying to follow it.

If this correct, we may think that Williamson (2005) is concerned with the distinction between conformity to the norm (primary appropriateness) and non-accidental conformity to the norm (following the norm). He is not focused on the distinction between primary and secondary appropriateness. Thus, even if we think that the distinction between primary appropriateness and secondary appropriateness is not a distinction between two norms, that does not imply that it is wrong to characterize Williamson's approach as a 'two norms' approach. (Also relevant here is the fact that, in any event, Williamson seems to think that the distinction between primary and secondary appropriateness actually refers to two norms, as his evocation of "derivative evidential norm" in the passage quoted above suggests)

Assume that Williamson is concerned with the distinction between conforming to a norm and non-accidentally conforming to a norm (following a norm). It might still be thought that as the question whether we follow a norm is not normative, but descriptive, then there is no norm underlying these evaluations. Thus, regarding the norm of belief, it has been argued that we should distinguish correctness conditions and regulation conditions.⁴ Roughly, a correctness condition says when an attitude, or action, is correct. For example, it is often said that a belief is correct if and only if it is true. This correctness condition is often used to distinguish beliefs from other attitudes (such as imagining, conjecturing, etc.).⁵ From the correctness condition, it is said, we can derive a norm applying to this attitude or action. For example, we may say that you ought to believe that p if and only if p is true. (Of course, this supposes that the notion of correctness is normative, or that we assume that we ought to have correct beliefs.)

⁴See Engel (2013).

⁵See chapter 2 section 2.2.

In contrast, a regulation condition is supposed to tell us which condition(s) must be satisfied for the attitude to be susceptible of being made correct, or for the subject to be considered as following a certain norm (i.e., as making an attitude or action correct with respect to this norm). That is, a regulation condition answers the (non-normative) question: Which conditions must be satisfied for the attitude or action to be liable to be (made) correct? This can take the form of an instrumental assessment regarding the means necessary to satisfy the end of having a correct attitude, or performing the correct action. For example, a belief that p is correct if and only if p is true. But to make your belief as to whether p correct, for example, you have to see whether p is true. A regulation condition for this norm is thus *seeing* that p . If there is no way of securing a correct belief as to whether p , for example because there is no evidence as to whether p or not p , then no regulation condition can be satisfied.

There can be tensions between a putative norm and possible regulation conditions. Suppose we say that you ought to believe that p if p is true. Now suppose that the proposition that it is raining and nobody believes that it is raining is true. Since this proposition is true, I ought to believe it, according to the supposed norm of belief. But if I believe this proposition, it is false. Therefore I ought not to believe it, according to the supposed norm. But this compromises the alleged norm. Arguably, a norm cannot be such that there are cases in which there is no logically possible way of following it. Here, the problem is that there are cases in which it is impossible to follow the norm according to which if p is true then you ought to believe that p . In other words, for some cases, there are no regulation conditions that it is logically possible to satisfy to follow this alleged norm: propositions that cannot logically be truly believed ('blindspot propositions') constitute an objection to such a norm. Although such propositions can be true, there is no means to have a true belief that they are true.

Let us grant the distinction between correctness condition and regulation condition. We may want to evaluate an attitude or action with respect to its correctness condition or with respect to its regulation conditions. The two assessments are not equivalent. You can have a correct belief that p without satisfying any regulation condition. Suppose that a correct belief is a true belief. If you happen to believe p and it happens that p is true, your belief is correct. If the norm of belief is truth, then your belief is conform to the norm. Still, we might think that your belief is, in some other sense, bad. You did not know that the belief was true, or correct; you took no means to secure a correct belief. Your belief is bad with respect to the regulation condition.

However, according to this train of thought, regulation conditions are not norms. In chapter 7 (section 3.1) I have considered this point. The question was whether we needed to assume an epistemic norm of action to account for our epistemic assessments of action. I considered in particular Fassio's view, according to which epistemic assessments of action bear on the regulation conditions of a non-epistemic norm (a norm according to which you ought to treat p as a reason if and only if it is a reason, where reasons are facts or true propositions), rather than on a satisfaction

(/ conformity) condition of an epistemic norm (Fassio, forthcoming). So, if Williamson is concerned with the distinction between primary appropriateness and regulation conditions, and if assessments of regulation conditions do not bear on norms, then Williamson's approach is not a 'two norms' approach.

There are two replies here. Consider the question whether there really is an epistemic norm of action. Remember I said that answering it was not essential for my purposes, for it is not essential to my opponent's view that such a norm is in force. What is essential to my opponent's view is that there is an epistemic standard of evaluation associated with actions, that this standard is variable, and that it is knowledge. Still, following most philosophers, and for the sake of simplicity, I decided to keep the formulation 'the epistemic norm of action' to designate the epistemic standard upon which we base our epistemic evaluations of action, whether or not we should conceive of this standard in terms of a genuine norm. Similar considerations apply here. It is essential to Williamson's approach that there are two epistemic conditions that we (can) use to evaluate actions. One is given by the (alleged) epistemic norm of action. The other has to do with the satisfaction of a regulation condition (if my interpretation of Williamson is correct). For the sake of simplicity, we can call these two conditions 'norms'.

There is a second, and perhaps more important, reason to use this terminology here. It is that this terminology fits better with what Williamson says. Firstly, recall that Williamson says that reasonableness (DeRose's secondary appropriateness) relies on a derivative (evidential) norm. So there is no reason not to consider regulation conditions as arising from a derivative norm too.

Secondly, Williamson suggests in several places that the reason why it is bad not to know that a premise is appropriate has to do with prudential considerations.⁶ But there are norms underlying our prudential evaluations.

Thirdly, in a more recent writing, Williamson distinguishes between primary norms from secondary and tertiary norms:

Let N be a norm for agents in situations. In a given situation, an agent may comply with N , or fail to do so. A given norm typically generates various derivative norms. In particular, there is the secondary norm DN of having a general disposition to comply with N , of being the sort of person who complies with N . Whether an agent A complies with DN in a situation S depends not just on whether A complies with N in S but also on whether A complies with N in possible situations beyond S .

(...) There is also a tertiary norm ODN of doing what someone who complied with DN would do in the situation at issue (' O ' for 'occurrent'). For example, I comply with ODP in a situation S if and only if I do in S what a promise-keeper would do in S . The status of ODN is derivative from that of DN , and thus doubly derivative from that of

⁶See Williamson (2005, 233; 235).

N. (Williamson, forthcoming)

Here too, the motivation for such distinctions has to do with different kinds of evaluations of what is somewhat “bad” or “good”:

Despite such differences between the three norms, if we endorse a primary norm N, we shall probably be sympathetic to the corresponding secondary norm DN and the tertiary norm ODN too. If it is bad not to keep a promise, then there is also something bad about not being generally disposed to keep one’s promises, and something bad about not doing what someone so disposed would do. (Ibid.)

The fact that the normative force of DN (or ODN) is partially derivative from the normative force of N (or N and DN) does not imply that DN (or ODN) is not a norm:

Usually, if you ought to comply with N, then you ought to comply with DN too, even though the underpinning of the latter ‘ought’ is derivative from, and more complicated than, the underpinning of the former ‘ought’. We can therefore call DN a norm too. (Ibid.)

Similarly, it seems that if a norm N is in force for you then, usually, you ought to follow this norm (although this is not required by N itself; N only requires conformity). Suppose that knowing N and knowing that doing A is a way of conforming to N in the circumstances are necessary regulation conditions. That is, suppose these are conditions we must satisfy in order to follow N. Then, we can say, you (derivately) ought to know N and that doing A is a way of conforming to N.

This strongly suggests that Williamson has in mind a normative reading of ‘bad’ when he says that it is bad not to know that a premise is appropriate. If so, according to him, there is a genuine norm underwriting this evaluation. Therefore, I think that it is correct to describe Williamson’s view as a ‘two norms’ approach.

Williamson’s approach: the derivative epistemic norm

With these preliminary remarks in mind, let us now turn to the discussion of Williamson’s approach. Williamson agrees that the orthodox view of knowledge and the claim that knowledge is the norm of action *seem* to have an implausible consequence, namely, that the standards for epistemic appropriateness of practical reasoning do not shift with the practical context. However, he argues, drawing this consequence is to quick. Indeed, as anticipated, according to him, we may distinguish between different levels of evaluation:

Suppose that an agent relies on the premise q in otherwise impeccable practical reasoning, and (in her context) q is indeed appropriate, but she is in no position to know that q is appropriate. Then our assessment of her reasoning should be ambivalent: it is good

that she used an appropriate premise, bad that she was in no position to know that it was appropriate. No simple verdict captures the complexity of the position. (2005, 230).

Similarly, Hawthorne and Stanley write:

intuitions go a little hazy in any situation that some candidate normative theory says is sufficient to make it that one ought to F but where, in the case described, one does not know that situation obtain. (Hawthorne and Stanley, 2008).

I shall call the evaluation appealing to knowledge of appropriateness ‘derivative epistemic appropriateness’ to distinguish it from primary appropriateness and from reasonableness (or DeRose’s secondary appropriateness). Williamson suggests two ways in which (what I call) derivative epistemic appropriateness can be affected by practical factors. First, practical factors can affect the importance of being derivatively epistemically proper. That is, practical factors can affect the normative force of the derivative epistemic norm:

How harshly should we judge practical reasoning in which the agent relies on an appropriate premise without being in a position to know that it is appropriate? A natural answer is: it depends on how much is at stake. If not much, then it seems unreasonably pedantic to condemn the reasoning. But if matters of life and death are at stake, the charge that the agent was not in a position to know that the premise was appropriate becomes more serious. (2005, 230).

The use of the term “pedantic” suggests that, according to Williamson, derivative epistemic appropriateness is always required, although in cases in which the stakes are very low this requirement is very weak. Intuitively, the idea is that when the stakes are low it is not that important to be primarily appropriate, and hence, not that important to follow the norm. Here, practical factors are not supposed to affect whether there is a requirement of derivative epistemic appropriateness, but merely the force of this requirement. To use Dancy’s terminology (2004, ch. 3) we may say that practical factors can be *intensifiers* of the derivative epistemic norm.

The second way in which Williamson suggests that practical factors can affect the derivative epistemic norm is by affecting this derivative epistemic norm:

How many iterations [of knowledge] are relevant depends on how much is at stake. (2005, 231).

Thus, what is supposed to be possibly affected by practical factors is not merely the strength of the derivative epistemic norm, but also the epistemic position that this norm requires. On some practical contexts, the derivative epistemic norm requires knowledge of primary appropriateness. On others, the derivative epistemic norm requires knowledge of knowledge of primary appropriateness.

And so on. Using Dancy's terminology, we may say that practical factors can be *enablers* of a derivative epistemic norm with such or such epistemic condition.

Armed with the distinction between primary appropriateness and derivative epistemic appropriateness, we may understand Williamson as claiming that the epistemic norm of action is a primary norm, and that overall epistemic appropriateness is a function of both primary (epistemic) and derivative epistemic appropriateness. While primary (epistemic) appropriateness is insensitive to practical factors, overall (epistemic) appropriateness is sensitive to practical factors, due to the fact that derivative epistemic appropriateness is sensitive to practical factors.

Now, consider the pairs of low and high stakes cases that are typically put forth in favour of the sensitivity of the epistemic norm of action (e.g., the bank cases). According to Williamson, the subjects' epistemic positions in the low and high stakes cases are typically such that these subjects are in a position to know the target proposition, but not to know that they know it.⁷ Suppose the subjects' epistemic positions in the low and high stakes cases should be seen in this way. With the help of the distinction between primary and derivative epistemic appropriateness, and the knowledge norm of action, we can explain why the low stakes subject and the high stakes subject differ in overall epistemic appropriateness. Neither of these subjects meet the derivative epistemic norm. In other words, even if these subjects satisfy the conformity condition of the epistemic norm of action, they are not in a position to follow this norm. Although this does not really matter for the low stakes subject, this matters a lot for the high stakes subject. Thus, the difference that appears at the level of the overall epistemic evaluation ultimately comes from a difference occurring at the level of derivative epistemic appropriateness, not at the level of primary epistemic appropriateness. In sum, the subjects are on a par as far as primary epistemic appropriateness is concerned, but not on a par as far as derivative epistemic appropriateness is concerned. Hence, they are not on a par as far as overall epistemic appropriateness is concerned. While it is roughly overall (epistemically) appropriate for the low stakes subject to act on the target proposition (it would be "pedantic" to criticise her action on the basis of the fact that it is secondarily inappropriate), it is clearly not overall (epistemically) appropriate for the high stakes subject to act on the target proposition.⁸

At this stage, we might worry that the two norms approach is forced to say that in such cases, the two norms conflict. That is, the primary epistemic norm requires to treat the target proposition as a reason – since the subject knows it – and the derivative epistemic norm requires not to treat it as a reason – since the subject does not know that she knows it. Since the subject knows that the bank is open, then she should be inclined to come back (she knows the situation, and what she should do if this situation is the case). Still, since the subject does not know that she knows that

⁷In Williamson's terms, this means that the subjects' beliefs in the target proposition are safe (there are no close possible worlds in which the subjects have a false belief based in the same way), but a belief in the fact that they know the target proposition would not be safe (there are close possible worlds in which the higher-order beliefs based in the same way are false).

⁸See Williamson (2005, 232-233).

the bank is open, she should be inclined to refrain from coming back, since she is not in a good enough epistemic position to follow the epistemic norm. But if the agent is in this conflictual state, how can she act?

Williamson agrees that the agent is in an ambivalent normative situation. He suggests, however, that the fact that the subject does not satisfy the derivative epistemic norm with respect to the primary epistemic norm destroys the motivational force associated with the satisfaction of the primary epistemic norm. Thus, he writes:

If stakes are high enough, prudent human agents will engage in third-order reasoning about whether to trust their second-order reasoning about whether to trust their first-order reasoning, and so on. Consider this dialogue (with oneself or another):

Q1. Is *q* the case?

A. Yes.

Q2. Did you have warrant for your answer to Q1?

A. Yes.

Q4. Did you have warrant for your answer to Q2?

A. I don't know.

At any point in such an interrogation, anything less than a positive answer seems to destabilize the previous positive answer, and therefore all the earlier positive answers in a domino effect. For what use is an assertion if the speaker is unwilling to stand over it at the next level up? Yet each further question in effect demands a further iteration of knowledge. (Williamson, 2005, 233)

In light of these considerations, we can suggest that Williamson's answer to the present worry would be that in such cases there is no motivational conflict in the agent's state. Given that the agent does not satisfy the derivative epistemic norm of action, the satisfaction of the primary epistemic norm loses its motivational force on the subject.

Brown's objection

A different worry has been raised against this approach by Brown (2008, 180-181). As Brown notes, it is unclear why the low and high stakes subjects should be seen as not in a position to know that they know the target propositions. And if they are in a position to know that they know, this explanation cannot work.

However, as said above, according to Williamson, there are two ways in which practical factors can affect derivative epistemic appropriateness. In some cases the practical factors affect what the derivative epistemic norm requires:

This extension of the point to higher iterations of knowledge enables the insensitive invariantist to handle cases that [may be offered] in which by ordinary standards subjects not only know but know that they know, being far from cases in which they falsely believe that they know; (...) although they know that they know, they lack some higher iteration of knowledge. (2005, 234)

Thus, even if we assume that both the low and high stakes subjects know that they know the target proposition, that does not mean that they both satisfy the derivative epistemic norm of action. It may be that the high stakes subject, but not the low stakes subject, should know that she knows that she knows the target proposition to satisfy the derivative epistemic norm. Hence, it may be that only the low stakes subject is overall (epistemically) proper.

Further, there is another reply available to Williamson. This reply is based on the idea that while the low stakes subject can be in a position to know that she knows the target proposition, the high stakes subject cannot be in a position to know that she knows. Indeed, typically, there is a possibility of error that is salient to the high stakes subject (e.g., in the high stakes bank case, that the bank has changed its hours). But, as Williamson notes, there is a bias typically associated with the salience of possibilities of error. We have a tendency to take salient possibilities of error as more probable than what they are in fact:

the high standards case is constructed in just such a way as to focus the ascriber's attention on considerations that tell against the ascription of knowledge, more specifically, on possibilities of error. They may be psychologically salient because the practical costs of error are high for the subject or the ascriber, or simply because they have been evoked in vivid and convincing detail. No wonder the ascriber gives them increased weight. One effect of fictional violence on television is to make viewers overestimate their chances of being the victims of violent crime: they suffer an illusion of danger. Might not an illusion of epistemic danger result from exposure to lurid stories about brains in vats, evil demons, painted mules, or gamblers who bet the farm? (2005, 226)

Given this bias, a judgement about whether she knows is less reliable for the high stakes subject than it is for the low stakes subject:

bias in our beliefs about which first-order beliefs constitute knowledge makes those second-order beliefs less reliable, thereby making it harder for them to constitute second-order knowledge. (2005, 235)

Thus, Williamson can argue that the asymmetry between the low stakes subject and the high stakes subject regarding the derivative epistemic norm has to do with the fact that the low stakes subject can satisfy this norm (she can know that she knows the target proposition because her higher-order judgement would be reliable enough given that it is not affected by a bias) while the high stakes

subject cannot (she cannot know that she knows the target proposition because her higher-order judgement would not be reliable enough given that it is affected by a bias).

The epistemic norm of action and the derivative epistemic norm

Although Williamson's approach has sufficient resources to respond to Brown's worry, it faces two more serious objections. The first one concerns whether the distinction between primary appropriateness and derivative epistemic appropriateness can apply to the epistemic norm itself. The second one concerns the validity of Williamson's diagnosis of the cases.

We may object to the way in which Williamson uses the distinction between primary appropriateness and derivative epistemic appropriateness. Williamson writes:

It is tempting to object that such a position [in which the premise of the practical reasoning is primarily proper but the subject does not know it] cannot arise, because it is the nature of appropriateness as a premise for practical reasoning to be epistemically accessible to the agent. For she must reason from what is accessible to her and cannot be condemned for not taking account of anything else. On the objector's view, whenever q is appropriate, the agent is in a position to know that q is appropriate: in the terminology of my *Knowledge and its Limits*, that q is an appropriate premise for practical reasoning is a luminous condition. But it is argued there that only trivial conditions are luminous

In this passage, Williamson seems to consider a view according to which the subject's rational beliefs about appropriateness are infallible. Presumably, this kind of view is based on the idea that we should define appropriateness in terms of the subject's rational beliefs about what is appropriate (or, at least, in terms of the normative beliefs favoured by the subject's epistemic perspective). On this kind of view, it is rational for S to believe that S ought to ϕ if and only if S ought to ϕ .⁹ If this view is correct, there is no primary appropriateness without being in a position to know that primary appropriateness obtains. Indeed, suppose that the premise is primarily appropriate for S . In virtue of what it takes to be an appropriate premise for S , S has a rational belief that the premise is appropriate, or at least, it is rational for S to believe that the premise is appropriate. Hence, if S has a rational belief that the premise is appropriate, this belief logically cannot be false, and thereby, that the premise is appropriate is known.

Against this view, Williamson appeals to his antiluminosity argument. The antiluminosity argument has some flaw.¹⁰ Be it as it may, I am inclined to grant the general distinction between primary appropriateness and something like a derivative epistemic appropriateness related to regulation conditions. Few would deny that if the law sanctions crossing at a green light, then if you

⁹See Gibbons (2013) for this kind of view.

¹⁰It relies on a safety condition on knowledge according to which if S knows that p at t , then p is true at $t+1$. Still, the usual safety condition on knowledge requires something weaker, namely, that if S knows that p at t , then S does not falsely believe p at $t+1$ on the same basis.

cross at a green light your action is primarily (legally) proper. However, it seems that if you did not know (or had not reason to think) that the light was green, something was wrong in your action. It was not proper in a different sense.

Still, one may grant that there is a general distinction between primary appropriateness and derivative epistemic appropriateness, while objecting to Williamson's approach that this distinction does not make sense when it comes to *epistemic* appropriateness. Contrary to what Williamson suggests, the objection need not be based on the idea that whenever q is (primarily) appropriate for action, the agent is in a position to know that q is (primarily) appropriate for action. It may just be based on the weaker claim that whenever q is *epistemically* (primarily) appropriate for action, then q is derivatively epistemically appropriate. This does not mean that whenever q is (primarily) epistemically appropriate for action, then S knows (that she knows...) that q is epistemically appropriate for action. The thought is just that the notion of (primary) epistemic appropriateness just *is* the notion which is also captured by the notion of derivative epistemic norm.

It is important to stress that this view does not entail that being epistemically appropriate is a luminous condition, or that if you are epistemically proper, then you know that you are epistemically proper. It merely entails that, when the premise is epistemically appropriate, the question whether or not you know that it is epistemically appropriate does not matter. That is, there is no sense in which it is required that you know that you are epistemically proper when you are epistemically proper. If it matters whether you know that you satisfy some epistemic condition with respect to p before acting on p , then this epistemic condition is insufficient for epistemic appropriateness. In brief, one may argue that it is not always true, contrary to what Williamson (2005, 230) claims, that when "an agent relies on the premise q in otherwise impeccable practical reasoning, and (in her context) q is indeed appropriate, but she is in no position to know that q is appropriate", then "it is bad that she was in no position to know that it was appropriate". In particular, when the appropriateness in question is epistemic, then, arguably, there is nothing "bad" in the fact that the agent is not in a position to know this epistemic appropriateness.

In sum, we can agree with Williamson that having the authority and knowing (or being in a position to know) that one has the authority are two different notions. But the question is whether knowing (or being in a position to know) that one has the authority always matters, that is, if there is always a sense in which it is "bad" not to know (or not to be in a position to know) that one has the authority even when one has it. According to the view under consideration, knowing (or being in a position to know) whether one has the epistemic authority never matters when one has the epistemic authority. That is, there is no (secondary) sense in which it is "bad" not to (be in a position to) know that q is epistemically appropriate if q is epistemically appropriate. If the question whether one knows that one has the epistemic authority really matters, this means that we do not have the epistemic authority in the first place if one does not know that one has it.

Two considerations militate in favour of the identification of the epistemic norm of action with

the derivative epistemic norm. First, the epistemic norm of action, like the derivative epistemic norm, is supposed to apply to all kinds of reasons for actions (legal, moral, etc.). It is not just a norm besides the others. It plays the special role of saying when the reasons for action are sufficiently available, or accessible, to the subject.

Second, given that the derivative epistemic norm delivers a general epistemic condition, the idea that there is a further general epistemic condition required for practical reasoning delivered by a primary epistemic norm is ill-motivated. In contrast to primary *practical* appropriateness for action, it is unclear what specific role the alleged primary *epistemic* appropriateness for action is supposed to play. When we consider the difference between primary practical appropriateness and derivative epistemic appropriateness, we have a clear sense that we are assessing something different. In general, primary appropriateness is a matter of the action being conform to the norms (legal, moral, etc.), and derivative epistemic appropriateness is a matter of being a good agent with respect to these norms. Being a good agent involves in particular caring about these norms, and hence, being interested in whether or not one conforms to them. Presumably, this is why we may think that there is still something bad with someone conforming to a norm while not knowing that she conforms to it.

However, when it comes to the epistemic norm of action, it is unclear why we should think that someone who already meets this norm – and hence, sufficiently cares about what there is reason for her to do – should also care about whether she sufficiently cares about what she has reason to do. If she sufficiently care about what she ought to do, then she should not care more. If she should care about whether she sufficiently care about what she ought to do, and she does not, then she does not sufficiently care about what she ought do. Invoking a distinction between primary epistemic appropriateness and derivative epistemic appropriateness with regard to primary epistemic appropriateness seems spurious.

In other words, the epistemic appropriateness related to the epistemic norm of action is the appropriateness of the epistemic access to the premise. So, it is in the nature of this epistemic appropriateness that, if it obtains, nothing more in terms of access is required. But a requirement that one knows that one has this access seems to be legitimate only if a requirement to have a better access to the premise is legitimate. And this latter requirement is legitimate only if the access to the premise was epistemically not good enough in the first place.

Perhaps Williamson could reply that a distinction between primary epistemic appropriateness and derivative epistemic appropriateness is required because two different norms are involved. Thus, it seems that for Williamson, derivative epistemic appropriateness arises from a prudential norm related to “habits of good decision making” (2005, 235).¹¹ However, even if we assume that the source of the derivative epistemic norm is prudential (in the sense that it is affected by practical factors), it remains unclear why we need a further epistemic norm of action.

¹¹Williamson also refers to “prudent” (2005, 233) and “cautious” (2005, 235) agent and decision-making.

Williamson considers the possibility to reject what I have called his two norms approach, and to propose an epistemic norm of action which is sensitive to the practical context:

The insensitive invariantist [about knowledge] could try to build variation in the required number of iterations of knowledge into appropriateness itself (...): in some cases q would be appropriate iff one knew q , in others iff one knew that one knew q , and so on, depending on the stakes. If such a move is acceptable, it provides the insensitive invariantist with a systematic response to arguments from practical differences to shifting semantic standards for epistemic terms. (2005, 232)

However, he raises two objections against this approach. The first one is the following:

However, the move does not guarantee the epistemic accessibility of appropriateness, for the anti-luminosity argument is quite general. (2005, 232)

While it is true that this way of understanding the epistemic norm of action does not guarantee the epistemic accessibility of (epistemic) appropriateness, it is unclear why this is problematic. This move can guarantee that there is no relevant sense of “bad” in which it is bad that epistemic appropriateness is not known, and this seems sufficient.

The second objection is the following:

Moreover, [this move] unperspicuously mixes together considerations at different levels. In knowing that one knows q , one knows a truth about one’s own epistemic state; in knowing q itself, one may know a truth simply about the external world. (2005, 232).

In reply to this objection, we may note that it might be wrong, in the first place, to understand the shift in epistemic requirement manifested by the pairs of cases in terms of iterations of knowledge, or higher-order epistemic positions. Perhaps we should understand the shift primarily in terms of strength of epistemic position with respect to the target proposition. I shall now show that Williamson’s diagnosis in this respect is, indeed, wrong.

The problem with Williamson’s diagnosis

Whether or not we think that two different norms can be involved (a primary epistemic norm and a derivative norm), Williamson’s approach faces a second objection, namely, that it relies on a wrong diagnosis of the standard pairs of cases. The diagnosis is that the shift concerns secondary appropriateness, understood in terms of higher-order epistemic positions. If this diagnosis is wrong and the shift concerns primary appropriateness – a first-order epistemic position with respect to the target proposition – then whatever we think of the distinction between primary and secondary norms, this distinction is irrelevant to understand the standard cases. Indeed, whether or not the

subject is secondarily proper, it might still be argued on the basis of these cases that the (primary) epistemic norm is sensitive to the practical context.

Three considerations can be used to show that Williamson's diagnosis is wrong. Firstly, the discrepancy between primary and secondary appropriateness is supposed to generate hazy or ambivalent assessments. However, our judgements regarding the standard pairs of cases are in no way ambivalent or hazy. It is clear, for example, that low stakes Keith should come back without checking whether the bank is open and that the high stakes Keith should not come back without checking whether the bank is open. If high stakes Keith comes back without checking, there is no sense that, at least, his action was primarily epistemically proper. Thus, the shift in epistemic requirements seems to concern a primary (epistemic) norm.

Secondly, if Williamson's diagnosis is right, then the high stakes subject should be interested in whether she is secondarily proper, for this is where the problem is supposed to lie. Hence, she should be interested in whether she knows that she knows (that she knows...) the target proposition. Thus Williamson writes:

Hi [the high stakes subject] has far more reason than Lo [the low stakes subject] has to check on such practical reasoning, to engage in second-order practical reasoning about whether to trust the first-order practical reasoning. (2005, 232-232)

However, the intuition is not that the high stakes subject should check her first-order reasoning, but rather, that she should check whether the target proposition is true. For example, in the high stakes bank case, intuitively, Keith should check whether the bank is open, not whether she knows (that she knows...) that the bank is open. It would be a level confusion to think that acquiring more evidence for p is to acquire more evidence for the proposition that one knows that p .

Thirdly, note that whether or not the subject's epistemic position is good enough to act on the target proposition does not seem to depend on the question whether the subject has reflective capacities. For example, it is plausible to think that animals and small infants cannot audit their own reasons or engage in higher-order reasonings about their first-order reasonings. Yet, it still seems that they can (and should) adjust their behaviour and their epistemic position in function of practical considerations, such as stakes. This suggests that the primary epistemic norm of action itself is sensitive to the stakes.

In sum, even if we grant that there is an interesting distinction between a primary *epistemic* norm and a secondary norm concerning our epistemic position with respect to whether this primary epistemic norm is satisfied, this distinction is irrelevant to understand the standard pairs of cases. Thus, even if we grant this distinction, we may still think that these cases support the claim that the epistemic norm of action is sensitive to practical factors.

9.3.2 The insufficient reason approach

The epistemic norm of action and the favouring relation

In the previous section, I have argued against Williamson's attempt to exploit the falsity of the equivalence thesis – the thesis according to which the overall epistemic requirement is fully determined by the epistemic norm of action. I have argued, in particular, against the relevance of the distinction between primary and derivative (epistemic) appropriateness to the cases under consideration. Still, perhaps there is another way of exploiting the falsity of the equivalence thesis.

In general, rejecting the equivalence thesis amounts to granting that there are cases where your epistemic position with respect to p is insufficient for you to do A , even if p favours doing A and p is a reason you have to do A because you satisfy the epistemic norm of action with respect to p . Indeed, we must distinguish two questions. One is whether a reason is warranted for a subject. Another is what action this reason favours (and how strongly) for this subject.¹² There is no reason to deny that the considerations which figure as reasons in the favouring relation can include an epistemic condition. Recall JUDGE. That Bill is guilty may be a reason John has to sentence Bill even if, due to the threat of the evil demon, John should not sentence Bill for the reason that Bill is guilty, and the reason for this is that John's epistemic position with respect to Bill's culpability is too weak. To defend the insensitivity of the epistemic norm of action, we may thus appeal to the general idea that practical factors can affect the overall epistemic requirement with respect to p by affecting an epistemic consideration relevant to whether p is a *sufficient* reason for the action.

Let me stress, however, that rejecting the equivalence thesis on this basis has no direct consequence on whether or not the epistemic norm of action is sensitive to practical factors. It merely opens the *possibility* that practical factors affect the overall epistemic requirement merely by affecting the favouring relation, and without affecting the epistemic requirement issued by the epistemic norm of action. Thus, this approach still needs an explanation of the standard pairs of cases typically offered in favour of the sensitivity of the epistemic norm of action (e.g., the bank cases) in terms of sufficient and insufficient reasons. I shall argue, however, that an explanation of this sort cannot do justice both to what the epistemic norm of action, naturally understood, is supposed to govern and to what shifts in the relevant pairs of cases.

The insensitivity of the epistemic norm of action defended by the insufficient / sufficient reason distinction

Can we use the distinction between insufficient and sufficient reason to defend the insensitivity of the epistemic norm of action? This distinction has been used in a close but slightly different context to defend, against cases purporting to show the contrary, that knowledge (or justification to believe that one knows) is always sufficient for action. I shall start by considering this debate, and next,

¹²On the favoring relation, see Dancy (2004, ch. 3) and Skorupski (2010, ch. 2).

whether the same kind of defence can be used by the advocate of the insensitivity of the epistemic norm.

Some philosophers have put forth cases where the subject is supposed to know (or to have justification to believe that she knows) although, intuitively, she should not act on the target proposition because her epistemic position is too weak. Philosophers advocating for a knowledge (or justification to believe that one knows) norm have defended their account by using the distinction between having a reason and having a sufficient reason. Roughly, these philosophers claim that the cases in which more than knowledge is required with respect to the target proposition do not compromise the sufficiency of knowledge (or justification to believe that one knows) for action. Indeed, according to them, in these cases a stronger epistemic position is required with respect to the proposition for having a sufficient reason to perform the action in question, but not to satisfy the epistemic norm of action with respect to this proposition.

The most discussed case which is supposed to speak against the sufficiency of knowledge (or justification to believe that one knows) for action is a case from Brown (2008, 176):

SURGEON. A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetised on the operating table. The operation hasn't started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:

Student: I don't understand. Why is she looking at the patient's records?
She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records.

Neta (2009) defends a justification to believe that one knows norm, and he grants that the surgeon knows the target proposition and is justified in believing that he knows the target proposition. According to him, however, this is a case in which "it is rationally permissible to treat a known proposition as a reason for acting, but it is also rationally permissible—in fact, it is rationally required—[to] double-check that p." (Neta, 2009, 697). Neta reconstructs the surgeon's reasoning in the following way:

SURGEON'S REASONING

Premise 1: The patient's left kidney is diseased.

Premise 2: Diseased kidneys must be transplanted.

Lemma: Therefore, transplant the patient's left kidney.

Premise 3: Before any kidney transplant, it is absolutely imperative to double-check which kidney needs to be transplanted.

Conclusion: Therefore, double-check the patient's records.

As we can see, the surgeon is supposed to (appropriately) use the proposition that the patient's left kidney is diseased as a reason, but this does not imply that he should not double-check. According to Neta "Brown's point is that it is rationally required to double-check the truth of a proposition before proceeding to take steps on the basis of its truth. But this point, while true, is perfectly compatible with the claim that permissibly treating a proposition as a reason for acting requires knowing (or justifiably believing oneself to know) that the proposition is true." (Neta, 2009, 697).

Ichikawa (2012) makes a similar claim:

Brown's argument is valid only on the assumption that that the disease is in the left kidney would be a sufficient reason for operating without checking the charts. But Brown has given us no reason to think that is so. It's entirely open to the defender of the knowledge norm to argue that knowledge of *p* is sufficient for *p* to be a reason, but that in this case, *p* isn't a good enough reason for action. Perhaps, for instance, under the circumstances, the proposition that the chart indicates that the disease is in the left kidney would need to be held as a reason in order to justify operating—this would straightforwardly explain the need to check the chart. Or perhaps a higher epistemic status than knowledge is needed vis-à-vis the proposition that the disease is in the left kidney—this too is consistent (...) on the supposition that only a proposition like the surgeon knows that the disease is in the left kidney would be a good enough reason for such a decisive action as removing a kidney, and Brown's surgeon does not, until she double-checks, know that she knows. Other candidate reasons are also available. (Ichikawa, 2012, 51)

Ichikawa concludes that "claims about who knows what, and about what actions are appropriate or inappropriate, are in general insufficient to provide test cases for the knowledge norm of practical reasoning. They are certainly not themselves sufficient to refute it. It is utterly unsurprising—and fully consistent with the knowledge norm—that often, propositions are genuinely held as reasons, but are nevertheless not sufficient reasons to perform particular actions. (Ichikawa, 2012, 50)

I will argue in the next chapter that Neta and Ichikawa's strategy is unsuccessful to account for Brown's case correctly understood.¹³ Here, I am merely interested in whether a similar strategy can be used to defend the insensitivity of the epistemic norm of action in general.

¹³In particular, Brown's point is not merely that the surgeon should double-check, but that she should double-check *because* the possibility of error is significant. Thus, Brown's claim is not merely a claim about what action is appropriate or inappropriate, but a claim about *why* a certain action is appropriate or inappropriate.

To be put to work, this strategy must show that stakes do not affect the epistemic norm of action but merely whether the target proposition is a sufficient reason for action. Consider the bank cases, and, following this suggestion, let us reconstruct what is supposed to be Keith's reasoning in the high stakes case:

HIGH STAKES REASONING

Premise 1*: The bank is open tomorrow

Premise 2*: If the bank is open tomorrow, I should come back tomorrow

Lemma*: Therefore, let's come back tomorrow

Premise 3*: Given that the stakes are high, it is absolutely imperative to double-check whether the bank is open

Conclusion*: Therefore, double-check whether the bank is open

Let us suppose that this is how Keith should reason in the high stakes case. There is still an important difference between SURGEON'S REASONING and HIGH STAKES REASONING. Premise 3 is supposed to be true in virtue of rules of surgery, for example, or in virtue of the fact that surgeons must operate on what the chart indicates. Whether or not premise 3 is appropriately used is independent from whether or not premise 1 is appropriately used. But in virtue of what is premise 3* true? Is there a general norm saying that when the stakes are high on whether p, we should double-check whether p? Even when the stakes are high on whether p, it makes sense to check whether p only if there is a risk that not-p. But if so, whether premise 3* is true depends on whether premise 1* is appropriately used. Indeed, as Locke writes:

Roughly, to premise that p in a particular piece of deliberation is to deliberate in a way that ignores, at least for the purposes of that deliberation, the possibility that not-p. (Locke, 2015, 87).

If so, if premise 1* is epistemically appropriate, then it is appropriate to ignore the possibility that the bank is closed in HIGH STAKES REASONING, at least epistemically speaking, that is, when the question is whether p is true (there may be reasons to consider the possibility that not-p that have nothing to do with the question whether p is true, for example, if it is required to double-check whatever the evidence we already have). But if it is epistemically appropriate to ignore the possibility that the bank is closed in this deliberation, then it is unclear why premise 3* is appropriately used. Reciprocally, if premise 3* is appropriately used, then it seems that premise 1* is not epistemically appropriate.

More formally, the line of reasoning is the following:

1. If S satisfies the epistemic norm of action with respect to p, then p is an epistemically appropriate premise for S (assumption)

2. A premise p is epistemically appropriate for S if and only if it is epistemically appropriate for S to ignore the risk that not- p in her deliberation (assumption)
3. Therefore, if S satisfies the epistemic norm of action with respect to p , then it is epistemically appropriate for S to ignore the risk that not- p in her deliberation (Safe Warranted Reason) (by 1. and 2.)
4. If high stakes make it that S should double-check whether p , then there is a risk that not- p that S cannot appropriately ignore in her deliberation, epistemically speaking (assumption).
5. If high stakes make it that S should double-check whether p , then S does not satisfy the epistemic norm of action with respect to p (by 3 and 4)

Note that (Safe Warranted Reason) is fully compatible with the falsity of the equivalence thesis. Indeed, that p is a warranted reason in this sense does not imply that if p is a warranted reason for S , then S 's epistemic position with respect to p is, overall, sufficient for S to act on p . In SURGEON'S REASONING, for example, it might be granted that premise 1 is an epistemically appropriate premise even if the surgeon should double-check. But this requires to grant that this is not because there is a risk for the surgeon that premise 1 is false that double-checking is required of him. Thus, it is not sufficient to reject (Equivalence thesis) to reject (Safe Warranted Reason), but (Safe Warranted Reason) is all that is required to defend the sensitivity of the epistemic norm of action.¹⁴

Indeed, (Safe Warranted Reason) implies that if the *mere* risk that not- p is the reason why one should not do A , then p is not a reason one has to do A . And in the standard high stakes cases, like the high stakes bank case, it is clear that the subject should double-check because there is a possibility that the target proposition is false. Therefore, the advocate of the insensitivity of the epistemic norm of action must either reject (Safe Warranted Reason) or the intuition that what shifts in the standard pairs of cases is the appropriateness of rejecting the possibility that the target proposition is false.

Suppose we want to reject (Safe Warranted Reason). Then, we are led to reject the idea that to treat p as a reason (or to premise that p) is to ignore the possibility that not- p in the deliberation. But if so, we could in the same reasoning allow the premise that p and the premise that not- p is

¹⁴Fantl and McGrath (2009) rely on something that may seem stronger than (Safe Warranted Reason). They rely on (Safe Reason): if S knows that p , then p is warranted enough to justify S in ϕ -ing, for any ϕ , where "in saying that p is warranted enough to justify you, we are saying that no weaknesses in your epistemic position with respect to p —in your position along any of the truth-relevant dimensions—stand in the way of p justifying you. (Fantl and McGrath, 2009, 66). I agree with Ichikawa (2012) that this may seem too strong. In particular, this seems to lead to (Equivalence Thesis). In contrast, (Safe Warranted Reason) is compatible with the falsity of (Equivalence Thesis), because it merely relies on the idea that if p is warranted enough to justify you, the *mere* risk that not- p do not stand in the way of p justifying you. Now, if we understand Fantl and McGrath as saying that if p is warranted enough to justify you, then no *mere* weaknesses in your epistemic position with respect to p stand in the way of p justifying you, then (Safe Reason) and (Safe Warranted Reason) are equivalent.

possible. But we never reason in this way. As Fantl and McGrath note:

We do not find ourselves weighing p – a reason to ϕ – against *there is a serious risk that not- p* – a reason not to ϕ . (Fantl and McGrath, 2009, 80)

For example, we do not find that Keith should reason by premising, on the one hand, that the bank is open, and, on the other hand, that there is a risk that the bank is closed. Therefore, it is difficult to reject (Safe Warranted Reason). On the other hand, it does not seem particularly controversial that what shifts with the stakes is the significance of the possibility of error. That is, it seems clear that in the low stakes case, Keith should premise that the bank is open and should not premise that there is a risk that the bank is closed, but, in the high stakes case, he should premise that there is a risk that the bank is closed and he should not premise that the bank is open. If so, the stakes affect the epistemic norm of action (what premises are epistemically appropriate), and not whether a consideration is a sufficient reason (whether a premise can justify the action).

In sum, it may well be that the overall epistemic requirement of action is not always fully determined by the epistemic norm of action. And it may well be that, sometimes, you satisfy the epistemic norm of action with respect to p even if p is not a sufficient reason you have to do A because your epistemic position with respect to p is not good enough. However, this is not sufficient to defend the insensitivity of the epistemic norm of action. Indeed, on a natural understanding, the epistemic norm of action tells us when it is appropriate to take for granted a proposition in our deliberation, that is, to ignore the possibility that this proposition is false. But the standard pairs of low and high stakes cases show that whether you can appropriately ignore a risk of error precisely depends on the stakes. The distinction between sufficient and insufficient reason helps us to show that (Equivalence thesis) is false, but it does not help to show that (Safe Warranted Reason) is false, or to show that what shifts with the stakes is not the significance of the risk of error.

9.4 The non-maximalist argument in favour of the sensitivity of the epistemic norm of action

So far, I have argued that appealing to a distinction between primary and derivative epistemic appropriateness, or to a distinction between having a reason and having a sufficient reason, is insufficient to defend insensitivism about the epistemic norm of action. The standard pairs of low and high stakes cases strongly suggest that stakes affect the epistemic norm of action by affecting whether you can legitimately ignore the risk of error.

I shall now offer another argument in favour of the sensitivity of the epistemic norm of action. As, to my knowledge, there is no independent argument in favour of the insensitivity of the epistemic norm of action, I shall conclude that the epistemic norm of action is sensitive to the practical context.

In favour of the claim that the epistemic norm of action is sensitive to the practical context, the following argument can be proposed:

1. If the epistemic norm of action is insensitive to the practical context, then it is a (non-sensitive) maximal-standard norm (assumption).
2. If the epistemic norm of action is a (non-sensitive) maximal-standard norm, in most cases ordinary subjects cannot satisfy this norm (assumption).
3. In most cases, ordinary subjects can satisfy the epistemic norm of action (assumption).
4. Therefore, the epistemic norm of action is not insensitive to the practical context (by 1, 2, 3)
5. Either the epistemic norm of action is sensitive to the practical context or it is insensitive to the practical context (assumption)
6. Therefore, the epistemic norm of action is sensitive to the practical context (by 4 and 5)

It seems to me that the only controversial assumption in this argument is 1. Why should we think that this assumption is true? It can be established by the following reductio:

1. S can satisfy the epistemic norm of action with respect to p even if there is a risk, call it 'R', that not-p (non-maximalist assumption)
2. If the risk R that not-p can be properly ignored in a situation S1, it can be properly ignored in any situation S (insensitivist assumption)
3. In any situation S, the risk R that not-p is properly ignored (by 1 and 2)
4. For any risk that not-p, we can design a high stakes situation, S*, in which it is not appropriate for the subject to ignore the risk that not-p. (assumption)
5. Therefore, the risk R that not-p is properly ignored in any situation S but there is a situation S* in which it is not properly ignored. (by 3 and 4)

Of course, 5 is a contradiction. So, if we assume that the norm is insensitive, we must reject either 1 or 4. It is difficult to reject 4. Think in particular about high stakes cases in which more evidence is easily available. It would be silly not to consider this evidence before acting. If so, the advocate of an insensitive norm must reject 1 and endorse a maximal-standard norm.

9.5 Conclusion

In this chapter, I have argued that although the overall epistemic requirement to act on a proposition is not always fully determined by the epistemic norm of action, we should still grant that the

epistemic norm of action is sensitive to practical factors. As a result, an orthodox philosopher must reject the knowledge norm of action. In addition, if knowledge is always sufficient for rational action, then either we should be sceptic or we should endorse pragmatic encroachment. Hence, an orthodox philosopher must also reject the sufficiency of knowledge. In the next chapter, I shall show that we have other independent reasons to think that knowledge is not always sufficient for rational action and I shall then defend a certainty norm of action which is sensitive to the practical context.

Chapter 10

The insufficiency of knowledge for action, and the certainty account

10.1 Introduction

In the previous chapter, I have argued that we must reject the insensitivity of the epistemic norm of action and that the only remaining option to avoid scepticism and pragmatic encroachment is to deny that knowledge is always sufficient to meet the epistemic norm of action. In this last chapter, my aim is to argue, first, against the sufficiency of knowledge on grounds independent from the denial of pragmatic encroachment. Second, it is to defend a certainty norm for action.

To argue against the sufficiency of knowledge, I start by considering cases in which more than knowledge is intuitively required. I argue that, to explain these cases, an advocate of the sufficiency of knowledge cannot appeal to the distinction between primary and derivative norm, or to the distinction between having a reason and having a sufficient reason. Next, I consider whether the claim that knowledge is sufficient for action is plausible given what knowledge is supposed to imply. In particular, I consider epistemic modal infallibilism about knowledge – the view according to which if a subject *S* knows that *p*, then there is no epistemic possibility for *S* that not-*p*. If epistemic modal infallibilism is true (and compatible with non-scepticism), then there is a principled argument in favour of the sufficiency of knowledge for epistemically rational action. However, I reject the main arguments in favour of epistemic modal infallibilism, and, finally, the main argument in favour of the sufficiency of knowledge.

To defend the certainty norm for action, I argue that this account satisfies all the desiderata that an account of the epistemic norm of action should satisfy. I emphasize that a certainty norm of action does not threaten the orthodox view of knowledge, which thereby undermines the third

argument against the orthodox view of knowledge.

10.2 Against the sufficiency of knowledge

10.2.1 Cases in which knowledge is insufficient

In chapter 8, I have noted that we sometimes use terms stronger than ‘know’ to criticize and defend actions in a way that seems incompatible with the idea that knowledge is always sufficient for rational action. More precisely, there are cases where we use epistemic terms stronger than ‘know’ – such as ‘certain’, ‘absolutely certain’, or ‘know for sure’ – to defend a subject refraining from acting on a certain proposition, and to criticise a subject acting on a proposition:

Defence for not acting on p (stronger): I checked whether the bank was open on Saturdays because I was not sure (∧ it was not certain) that the bank was open.

Criticism for acting on p (stronger): A mother might criticise her teenager’s late departure from a party by saying “You shouldn’t have left so late. You weren’t certain there’d be a bus home that late” or “You shouldn’t have left so late. You didn’t know for sure that there’d be a bus at that time”. (Brown, 2008, 172).

If knowledge is not always sufficient, we have a straightforward explanation why we sometimes use terms stronger than ‘know’, like ‘certain’, to criticise a subject acting on a proposition and to defend a subject refraining from acting on a proposition. We use these terms in situations where knowledge is insufficient for action.

Further, some philosophers have proposed cases where the subject is supposed to know a proposition although, intuitively, her epistemic position with respect to the target proposition is not good enough to act on it. We can classify these cases in two groups. First, there are what we may call ‘non-reflective cases’, in which the subject has knowledge without certainty. Second, there are what we may call ‘reflective cases’ in which the subject has knowledge without the required higher-order epistemic position. Although philosophers have mainly focused on the first type of cases, both are important to argue in favour or against one or another epistemic norm of action.

The main unreflective case which has been discussed in the literature is the surgeon case proposed by Brown (2008, 176):¹

SURGEON. A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetised on the operating table. The operation hasn’t

¹See also Reed (2010, 228-229; 232) for similar cases.

started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records.

In cases in which you know that p but it is not certain for you that p , it seems that the fact that it is not certain for you that p can be relevant to whether or not you should act on the target proposition. In particular, if the stakes are high, certainty may be required of you. Of course, this sort of cases is possible only if knowledge does not entail certainty. But the idea that knowledge does not entail certainty is shared by virtually all epistemologists nowadays.²

A reflective case is proposed by Neta (2009, 688):

HISTORY EXAM. You are taking an oral history exam, and you come across the question "in what year was Abraham Lincoln assassinated?" You know that the answer is 1865. But you are momentarily struck by a neurotic diffidence about your memory of this historical fact. After you hear the question, the first thought that goes through your mind is: I believe that the answer is 1865, but of course I don't know that it is! Although the latter conjunct is false—you do know that the answer is 1865—it's also true that you believe that you don't know that the answer is 1865, and it's furthermore true that you don't believe that you do know that the answer is 1865. A fortiori, you do not justifiably believe that you know that the answer is 1865. In the end, you decide to answer the question by saying "1865", on the basis of the proposition that the answer is 1865 as a reason for so answering the question.

In cases in which you (justifiably) believe that you do not know that p , it seems that it is inappropriate for you to treat p – rather than the fact that p is probable – as a reason to act, even if you know that p .³

In chapter 9, I have considered two ways in which the champions of the insensitivity of the epistemic norm of action try to explain our shifting judgements about silly and sensible decision-making when a particularly strong epistemic position is required with respect to the target proposition. Some appeal to the distinction between primary and secondary appropriateness, and they

²See also chapter 6.

³Neta uses HISTORY EXAM to argue in favour of a justification to believe that one knows norm of action. SURGEON puts pressure on Neta's account, though, because the surgeon can be seen as justifiably believing that he knows.

argue that while primary appropriateness is insensitive to the practical context, secondary appropriateness may require a particularly strong epistemic position. Others appeal to the distinction between having a reason and having a sufficient reason, and they argue that while having a reason is insensitive to the practical context, having a sufficient reason may require a particularly strong epistemic position.

When defending the sensitivity of the epistemic norm of action in the previous chapter, I have set aside the question whether knowledge is always sufficient. I have merely argued that these distinctions cannot be used to reject the claim that the epistemic norm of action is sensitive to practical factors, and that, given that the epistemic norm of action is sensitive to practical factors, the following conditional is true: if knowledge is always sufficient to satisfy the epistemic norm of action, then either we should be sceptic or we should embrace pragmatic encroachment. I have then concluded that the only option for the non-sceptical orthodox philosopher is to reject the claim that knowledge is always sufficient. However, typically, these distinctions have been used to defend the sufficiency of knowledge. I will now argue that neither of these distinctions can be used to reject the cases speaking against the claim that knowledge is always sufficient.

Consider Williamson's defence of the sufficiency of knowledge. It appeals to the distinction, explained in the previous chapter, between primary and derivative norms and the idea that the knowledge norm is a primary norm.⁴ Suppose we argue that the surgeon and the student are primarily proper but not derivatively proper. That is, suppose we say that they satisfy the epistemic norm of action – indeed, they know the target proposition – but that, given the practical context, their higher-order epistemic positions with respect to their first-order knowledge are insufficient. And suppose we say that this is what explains why we judge that, overall, they should not treat the target proposition as a reason. However, what is so bad with the surgeon's higher-order epistemic position? There is no reason to think that the surgeon does not know that he knows (that he knows...) the target proposition. Further, the problem with the surgeon's epistemic position does not seem to be that he does not have the required higher-order position to know, but merely that he does not know with certainty the target proposition. If he acquires first-order epistemic certainty with respect to the target proposition, this seems sufficient for him to appropriately act on the target proposition. If he acquires the (supposedly) necessary higher-order position to know, that does not guarantee that he will improve his epistemic position in the right way, since, in the end, the target proposition could still be known to be known (to be known...) without certainty. What is so bad with the student's higher-order epistemic position? We may think that even if the student does not know that he knows, and is thereby derivatively improper, given that the stakes are not high, it would be "pedantic" – to use Williamson's word – to condemn his action on the basis of the fact that he is secondarily improper. Presumably, the cost of error is low and the reward in accuracy is high. And, after all, according to Williamson, this is what happens with low stakes

⁴See Williamson (2005).

Keith in the bank cases.

Consider Neta and Ichikawa's appeal to the distinction between having a reason and having a sufficient reason.⁵ According to Neta,

Brown's point is that it is rationally required to double-check the truth of a proposition before proceeding to take steps on the basis of its truth. But this point, while true, is perfectly compatible with the claim that permissibly treating a proposition as a reason for acting requires knowing (or justifiably believing oneself to know) that the proposition is true. (Neta, 2009, 697).

Ichikawa agrees:

claims about who knows what, and about what actions are appropriate or inappropriate, are in general insufficient to provide test cases for the knowledge norm of practical reasoning. They are certainly not themselves sufficient to refute it. It is utterly unsurprising—and fully consistent with the knowledge norm—that often, propositions are genuinely held as reasons, but are nevertheless not sufficient reasons to perform particular actions. (Ichikawa, 2012, 50)

According to him, what we have in SURGEON is merely:

a pair of intuitive verdicts: one attributing knowledge, and another denying appropriateness of action:

1. The surgeon knows (before reading the chart) that the disease is in the left kidney.
2. It would be inappropriate for the surgeon to remove the left kidney without first collecting more evidence.

If these are the only verdicts we have, then it seems true that these cases do not show that knowledge is sometimes insufficient to satisfy the epistemic norm of action. Perhaps, in SURGEON, *the disease kidney is the is the left kidney* is not a sufficient reason to remove the left kidney, and a sufficient reason would be *the chart indicates that the disease kidney is the left kidney*. And perhaps, in HISTORY EXAM, *the answer is 1865* is not a sufficient reason, and a sufficient reason would be *I know that the answer is 1865*.

However, these are not the only verdicts that we have. There is the further verdict that it would be inappropriate for the surgeon to remove the left kidney without first collecting more evidence *because there is a possibility of error*. Brown's point is not merely that the surgeon should double-check, or that it is appropriate to collect more evidence before removing the kidney. It is that the

⁵Neta (2009) uses the distinction between having a reason and having a sufficient reason to resist cases like SURGEON.

surgeon should double-check due to the possibility of error. This is made clear by what the nurse says: “imagine what it would be like if she removed the wrong kidney”.⁶ But I have argued in chapter 9 (section 3.2) that if the reason why a subject should check whether p is that there is a possibility of error that she should not ignore in her deliberation (given that the question is whether p is true), then she does not satisfy the epistemic norm of action. That is, if p is an epistemically appropriate premise in S 's practical reasoning, then *not- p is possible* is not an epistemically appropriate premise for S 's practical reasoning. If p is a reason one has, then it is epistemically appropriate to ignore the possibility that not- p . If so, given that it is appropriate for the surgeon to premise *there is a possibility that the disease kidney is not the left kidney*, even if she knows that the disease kidney is the left kidney, he does not satisfy the epistemic norm with respect to this proposition to treat it as a reason.

In sum, it may be granted that there are cases in which a reason we have is insufficient to act because, although we satisfy the epistemic norm of action, our epistemic position is insufficient for this reason to be a sufficient reason. However, this is not what happens in Brown's surgeon case correctly understood. Similarly, this is not what happens in the high stakes bank case, and it is unclear how this distinction could be used for HISTORY EXAM. The distinction between having a reason and having a sufficient reason does not help the orthodox philosopher advocating for the sufficiency of knowledge.

It is worth noting that SURGEON and HISTORY EXAM also speak against the non-orthodox philosophers advocating for the sufficiency of knowledge. These cases are supposed to be cases in which the subject knows the target proposition while not being in a position to act on it. To defend the sufficiency of knowledge in the non-orthodox framework, we must say that the subject does not know the target proposition in these cases. This move may have some plausibility in SURGEON.⁷ Perhaps in this case the use of ‘know’ is not literal, and the surgeon in fact does not know the target proposition. But it is far less plausible to think that, in HISTORY EXAM, the subject does not (or cannot) know. Indeed, this would lead to say that in the situations in which it is inappropriate for you to act on p because you (justifiably) believe that you do not know that p , then you do not know that p . In these situations, a (justified) belief that you do not know that p would be infallible. The fact that knowledge is sometimes insufficient for action is more strongly established by reflective cases.

10.2.2 Maximal infallibilism and the sufficiency of knowledge

In order to see whether knowledge is always sufficient for action, it is also useful to consider whether this claim can be supported by what we take knowledge to be. Standardly, we can distinguish

⁶See also Brown (2011, 267): “the thought seems to be that knowledge is insufficient because of the chance of error”.

⁷See Fantl and McGrath (2009, 62).

between infallibilist and fallibilist accounts of knowledge. There are different ways of understanding this distinction, but we might think that, on some way of fleshing out infallibilism, necessarily, if S knows that p, then S satisfies the epistemic norm of action with respect to p.

According to Brown (2011, 163), for example, if we flesh out infallibilism about knowledge in terms of maximal epistemic position, we might propose an argument of the following sort in favour of the sufficiency of knowledge:

MAXIMAL INFALLIBILISM ARGUMENT

1. If S knows that p, then S satisfies maximal epistemic standards with respect to p (assumption).
2. If S satisfies maximal epistemic standards with respect to p, then it cannot be coherently required of S that she satisfies stronger epistemic standards (assumption).
3. If S knows that p, then it cannot be coherently required of S that she satisfies stronger epistemic standards (by 1 and 2)
4. If it cannot be coherently required of S that she satisfies stronger epistemic standards with respect to p, then S satisfies the epistemic norm of action with respect to p (assumption)
5. Therefore, if S knows that p, S satisfies the epistemic norm of action with respect to p (by 3 and 4).

Against this line of reasoning, Brown emphasizes that maximal infallibilism has sceptical consequences about knowledge. Indeed, on this view of knowledge, you satisfy maximal epistemic standards only if there is no possible stronger epistemic position with respect to the target proposition. However, it is clear that if knowledge requires a maximal epistemic position in this sense, then we know hardly anything. Since maximal infallibilism has sceptical consequences, it is an implausible account of knowledge. Hence, the argument from maximal infallibilism in favour of the sufficiency of knowledge for rational action fails.

Is there another way of fleshing out infallibilism which avoids scepticism while lending support to the sufficiency of knowledge for rational action? Brown (2011) argues that there is no infallibilist account of knowledge which meets these desiderata. Her line of reasoning is based on the following claims:

AGAINST INFALLIBILISM-BASED SUFFICIENCY

1. For any infallibilist account of knowledge, either it entails maximal infallibilism, or it does not (assumption).
2. Only a maximalist account of knowledge can be used to provide an argument in favour of the sufficiency of knowledge (assumption).

3. Therefore, only an infallibilist account of knowledge entailing maximal infallibilism can provide an argument in favour of the sufficiency of knowledge (by 1 and 2).
4. Maximal infallibilism is an implausible account of knowledge (assumption).
5. Therefore, only an implausible infallibilist account of knowledge can provide an argument in favour of the sufficiency of knowledge (by 3 and 4).

The thought that only maximal infallibilism can provide an argument in favour of sufficiency of knowledge is based on the idea that if you do not satisfy maximal epistemic standards, it is in principle possible that higher epistemic standards are required for action, while if you satisfy maximal epistemic standards, necessarily, you satisfy the epistemic norm of action. In Brown's words:

So long as the epistemic standard necessary for practical reasoning/assertion is not maximal, this leaves it open that the epistemic standard sufficient for epistemically appropriate practical reasoning/assertion is higher than what is necessary. (Brown, 2011, 162)

In contrast,

[a]ssume that the knowledge required for epistemically appropriate assertion/practical reasoning involves having the maximal strength of epistemic position. On that assumption, there is no coherent worry that the epistemic position of someone who meets the knowledge requirement for epistemically appropriate assertion/practical reasoning needs to be strengthened in order to have an epistemic position good enough for assertion/practical reasoning. For their epistemic position is already maximal. (Brown, 2011, 163)

Thus, the general idea is that if there are epistemic standards stronger than those required for knowledge which can be coherently required, the argument in favour of the sufficiency of knowledge for action from what knowledge requires fails. Indeed, it might be that the epistemic norm of action requires, in some cases, to satisfy these stronger standards.

10.2.3 Epistemic modals infallibilism and the sufficiency of knowledge

However, it can be argued that although Brown is right that if maximal infallibilism is true then knowledge is always sufficient for epistemically rational action, she is wrong about why this claim is true. This affects her argument according to which no plausible infallibilist account of knowledge lends support to the sufficiency of knowledge. Indeed, I will show that the reason why sufficiency of knowledge follows from maximal infallibilism might also be used to show that a non-maximal infallibilism can lend support to the sufficiency of knowledge.

According to Brown's proposal, the crucial step in support of the idea that sufficiency of knowledge follows from maximal infallibilism is the following: if S satisfies maximally high standards, it cannot be coherently required of S that she meets higher epistemic standards, and hence, S satisfies the epistemic norm of action (see premise 4 of MAXIMAL INFALLIBILISM ARGUMENT). Yet, from the fact that it cannot be coherently required of S that she satisfies stronger epistemic standards with respect to p, it does not follow that S's epistemic position with respect to p is good enough for S to act on p. This follows only if we also assume that, necessarily, for any proposition p and practical situation PS, there is a sufficient epistemic position with respect to p such that it can be epistemically rational to act on p in PS. But unless we give an argument in favour of this last claim, in principle, it might be that there are cases in which no (logically) possible epistemic position can be good enough for S to act on p. In such cases, even if, provided that S satisfies maximal standards, it is not required of S that she satisfies stronger epistemic standards (indeed, this would be incoherent), it may still be that meeting these maximal epistemic standards are insufficient, and that what is epistemically required of S is just not to act on the target proposition.

In other words, although we can in principle exclude cases in which maximal standards are insufficient to act because satisfying stronger epistemic standards is required, we cannot exclude without argument cases in which maximal standards are insufficient to act because no epistemic standards are sufficient to act on the target proposition. If such cases are possible in principle, MAXIMAL INFALLIBILISM ARGUMENT fails independently of its sceptical consequences.

Now, there is another way of arguing that maximal infallibilism implies that knowledge is sufficient for rational action. This argument does not appeal to the idea that it would be incoherent to require of you to satisfy higher epistemic standards when you satisfy maximum standards (again, this incoherence does not show by itself that you satisfy the epistemic standards for action). Rather, it appeals to the idea that you satisfy the epistemic norm of action with respect to p if and only if you epistemically ought to ignore any possibility in which not-p in your practical deliberation, in the sense that it would be epistemically inappropriate for you to use *not-p is possible* as a premise (see chapter 9). The argument is the following:

MAXIMAL INFALLIBILISM ARGUMENT*

1. If S knows that p, then S satisfies maximal epistemic standards with respect to p (assumption).
2. If S satisfies maximal epistemic standards with respect to p, then S is epistemically justified in excluding any possibility in which not-p (assumption).
3. If S knows that p, then S is epistemically justified in excluding any possibility in which not-p (by 1 and 2).
4. S is epistemically justified in excluding any possibility in which not-p in her practical

deliberation only if S satisfies the epistemic norm of action with respect to p (assumption).

5. If S knows that p, then S satisfies the epistemic norm of action with respect to p (by 3 and 4).

It is important to stress that this argument does not appeal to the contentious idea that if it cannot be coherently required to satisfy higher epistemic standards, then the epistemic norm of action is satisfied. It connects the question whether you satisfy the epistemic norm of action with respect to p with whether you can ignore the possibility that not-p in your practical deliberation. And it is uncontroversial that if you satisfy maximum epistemic standards with respect to p, then there is no not-p possibility that you cannot eliminate. The reason why maximal infallibilism implies that knowledge is always sufficient to meet the epistemic norm of action is not that we cannot meet epistemic standards stronger than maximal standards, but that the epistemic norm of action tells us what possibility of error can be epistemically ignored in our practical deliberation. If you know that p by maximal standards, you can epistemically eliminate all the possibilities of error.

Consider the claim according to which if you satisfy maximal standards then you can epistemically exclude all the possibilities of error. Suppose that this claim should play the key role in the argument from maximal infallibilism to the sufficiency of knowledge. If so, it might be that a version of non-maximal infallibilism could provide an argument in favour of the sufficiency of knowledge. Indeed, all that is required for this account of knowledge is to be such that if you know that p, then there is no epistemic chance, or possibility, for you that not-p. And we may think that satisfying maximal epistemic standards is not necessary to be in a good enough epistemic position to exclude all the possibilities in which not-p.

Following Brown, let us call the view according to which if S knows that p, then there is no possibility for S that not-p, ‘epistemic modals infallibilism’. Brown argues that either epistemic modals infallibilism leads to maximal infallibilism, and hence, it is an implausible account of knowledge, or else epistemic modals infallibilism cannot motivate the sufficiency of knowledge. She writes:

[S]uppose that epistemic modals infallibilism is true: if one knows that p, then it is not possible that not-p. On the assumption that epistemic modals infallibilism does not entail maximal infallibilism, these two suppositions are compatible with the claim that the knowledge required for epistemically appropriate practical reasoning and assertion is not maximal. Yet, if it is not maximal, there is an epistemic position stronger than one’s actual epistemic position. It seems that, in certain circumstances (for example, if stakes are very high), epistemically appropriate practical reasoning or assertion may require this stronger epistemic position. (Brown, 2011, 166)

Still, I have argued that the coherence of requiring higher epistemic standards is not the issue. Again, the mere fact that you cannot coherently satisfy higher epistemic standards does not show

that you satisfy the epistemic norm of action. Rather, the question whether it is epistemically appropriate for S to act on a certain proposition is the question whether there is a chance for S that not-p. Even if S's epistemic position with respect to p is not maximal and there is a coherent worry that a higher epistemic position is required by the epistemic norm of action, if it is assumed that S's epistemic position is such that there is no chance for S that not-p, it is unclear why this coherent worry is relevant to whether the epistemic norm of action is satisfied. Even if the stakes on being right on p are extremely high, if there is no epistemic chance that not-p for S, it is unclear how it could be epistemically irrational for S to act on p. High stakes can play a role only if there is an epistemic chance for the subject that she is wrong about p. If being in a good enough epistemic position to eliminate all the possibilities in which not-p does not require to satisfy maximal epistemic standards, a non-sceptical version of epistemic modals infallibilist might be used to argue in favour of the sufficiency of knowledge for rational action:

EPISTEMIC MODALS INFALLIBILISM ARGUMENT

1. If S knows that p, then it is not epistemically possible for S that not-p (assumption).
2. If it is not epistemically possible for S that not-p, then S satisfies the epistemic norm of action with respect to p (assumption).
3. If S knows that p, then S satisfies the epistemic norm of action with respect to p (by 1 and 2).

Consider the first premise of EPISTEMIC MODALS INFALLIBILISM ARGUMENT. Why thinking that epistemic modals infallibilism is true? This view is mainly motivated by the fact that utterances of the following sentences – often called concessive knowledge attributions (CKAs) – sound odd:⁸

- (1) I know that p, but it may be that not-p.
- (2) I know that p but there is a chance that not-p.

The oddity of these sentences is explained if we adopt what DeRose calls 'Moore's principle':⁹

MP: Whenever a speaker S does or can truly assert, "It's possible that P is false", S does not know that P. (DeRose, 1991, 596)

Moore's principle underlies epistemic modals infallibilism. In order to argue in favour of MP, we have to show that the clashes in (1)-(2) must be explained semantically and not pragmatically. Indeed, Moorean sentences clash, but the clashes are not due to a semantic inconsistency. Consider:

- (3) It is raining but I do not know that it is raining.

⁸See DeRose (1991); Lewis (1996).

⁹See Moore (1959).

It is clear that there is no inconsistency between the proposition that it is raining and the proposition that I do not know that it is raining. And, indeed,

it seems crazy to infer that it is not raining from the fact that I don't know that it is raining. If it is raining and I don't know that it is raining were inconsistent, then each would entail the negation of the other: it is raining would entail I know that it is raining and I don't know that it is raining would entail It is not raining. But, clearly, neither entailment holds. (DeRose, 1991, 597)

This suggests a way of distinguishing pragmatic and semantic clashes:

If a sentence clashes, if it is not clearly implausible to suppose that the entailments hold, and if we can find no good deflationary explanation for the clash, then we have good-reason to suppose that the clash is being produced by a genuine inconsistency in the sentence.

Now, according to DeRose,

when we apply this method to “It is possible that not-P, but I know that P”, I think we find that the sentence does seem to clash or cancel itself, and that, unlike “It is possible that not-P, but P”, it seems reasonable to suppose that the clash one senses may be that of genuine inconsistency, because the required entailments (from It is possible that not-P to I don't know that P and from I know that P to It is not possible that not-P) don't seem crazy. (DeRose, 1991, 599)

On this basis, and in the absence of a good pragmatic explanation of the clashes in (1)-(2), we may thus think that MP, and hence, the first premise of EPISTEMIC MODALS INFALLIBILISM ARGUMENT, is true.

Even if there is a *prima facie* good semantic argument in favour of epistemic modals infallibilism, we might still think that epistemic modals infallibilism about knowledge has sceptical consequences. If epistemic modals infallibilism has sceptical consequences, then, supposing that non-scepticism is true, epistemic modals infallibilism cannot be true. And if epistemic modals infallibilism cannot be true, then there is no available argument in favour of the sufficiency of knowledge for rational action on its basis.

For example, Lewis (1996) writes that fallibilism – as contrasted with sceptical infallibilism – is the view according to which S can know that p even if she has not eliminated every possibility in which not-p. According to Lewis, a possibility is eliminated if it is logically inconsistent with the subject's evidence. If we assume that these uneliminated possibilities are epistemic possibilities, then fallibilism is the view that MP is false: S can know that p even if not-p is possible for S.

Epistemic modals infallibilism would then be incompatible with fallibilism thus understood, and it would lead to a form of sceptical infallibilism.¹⁰

As Huemer (2007) notes:

It is metaphysically possible for someone to have the kind and degree of justification for a proposition that I presently have for the claim that I have two hands, and yet for that proposition to be false. This is amply shown by brain-in-a-vat and similar scenarios. The skeptic could therefore argue as follows:

1. If you know that P, then it is not epistemically possible for you that $\sim P$.
2. It is epistemically possible for you that you do not have two hands.
3. Therefore, you do not know that you have two hands. (From 1, 2.)

The non-sceptical orthodox philosopher advocating for MP must reject the claim that it is epistemically possible for you that you do not have hands. That is, she must say that the possibility of error in question is not epistemic, or that, although consistent with your evidence, this possibility can be *epistemically* eliminated. This requires to distinguish metaphysical possibilities of error and epistemic possibilities of error.

This line of thought is defended in particular by Stanley (2005a). As Stanley notes, “[f]allibilism is the doctrine that someone can know that p, even though their evidence for p is logically consistent with the truth of not-p.” (Stanley, 2005a, 127). This means that S can know that p on the basis of E even if it is logically possible that S has evidence E and not-p is the case. However, according to Stanley, this is compatible with epistemic modals infallibilism, the view according to which if S knows that p, then not-p is epistemically impossible for S. In a nutshell, E and not-p may be logically or metaphysically possible although given that S knows that p (on the basis of E), not-p is epistemically impossible for S. Epistemic possibilities are propositions logically (and obviously) compatible with what the subject knows, and not propositions logically compatible with the subject’s evidence. If so, epistemic modals infallibilism does not lead to scepticism, and given the semantic argument in favour of MP, the first premise of EPISTEMIC MODALS INFALLIBILISM ARGUMENT is plausible.

It has been objected to this approach of epistemic possibilities, in which possibilities of error logically or metaphysically compatible with the subject’s evidence are not necessarily epistemic possibilities, that it contradicts some plausible claims about what epistemic possibilities are, and how they can be eliminated. Thus, Dodd (2010) writes:

According to the received view of epistemic possibility [the one adopted by Stanley (2005a)], CKAs express necessarily false propositions. Is this really inconsistent with

¹⁰See also DeRose (1991, 596). DeRose (1991, 596) and Lewis (1996) endorse MP and epistemic modals infallibilism, but they avoid scepticism by embracing epistemic contextualism about knowledge. What possibilities of error are epistemically relevant depends on the conversational context.

fallibilism? Stanley (...) thinks it's not, though it's unclear why (...). As Lewis says, when S utters a CKA, S is saying (*) S knows p, but can't rule out a possibility in which not-p. Assuming with Lewis that it's in virtue of our evidence that we're able to rule epistemic possibilities out, and that our evidence rules a possibility out iff our evidence isn't consistent with the possibility's obtaining, (*) is equivalent to (**) S knows p, but there is a possibility in which not-p that's consistent with S's evidence. So, given plausible evidentialist assumptions about how we rule possibilities out, CKAs express necessarily false propositions iff (**) is necessarily false. (**) is necessarily false iff [fallibilism] is false. Thus it's plausible that if CKAs express necessarily false propositions, then fallibilism is false. Thus the fact that CKAs sound contradictory does pose a prima facie problem for fallibilists.

However, there are two problems with Dodd's criticism. First, it is clear that it is question-begging. Indeed, an advocate of a knowledge-based account of epistemic chances will just deny that it's in virtue of our evidence that we are able to rule epistemic possibilities out and that our evidence rules a possibility out if and only if our evidence is not consistent with the possibility obtaining. Perhaps it can be said that it's in virtue of our evidence that we are able to rule epistemic possibilities out but that our evidence rules a possibility out if and only if it is sufficient to know that it does not obtain, which, given fallibilism about knowledge, does not require that this possibility of error is metaphysically incompatible with the evidence. More directly, Stanley's claim amounts to the idea that it is in virtue of what we know that we are able to rule epistemic possibilities out and that a possibility of error is ruled out if and only if it is inconsistent with what we know. If so, granting fallibilism about knowledge, Stanley's claim is just that you can (epistemically) rule out a possibility of error even if your evidence is consistent with this possibility obtaining.

Second, contrary to what Dodd suggests, it is not so obvious that possibilities of error that are logically or metaphysically consistent with the subject's evidence are epistemic possibilities for the subject – i.e. possibilities that are not ruled out by the subject's evidence. Indeed, as Huemer notes, this view seems too permissive. Intuitively, not every metaphysical possibility in which the subject has the evidence or justification in question and the target proposition is false should count as an epistemic possibility for the subject.

To see this, consider the two following cases, from Huemer (2007, 121-122):

LOST WALLET. Having lost my wallet, I ask myself: "Where could it be? Could it be at the movie theater?" "No", I reason, "I remember buying gas after that, and I had to have my wallet to do that. Hm, but it might be out in the car..." I then go out to search the car.

PHONE AIRPORT. Mary is picking up Sam from the airport, but she's a little late, so she calls Sam on his cell phone.

Mary: Where are you?

Sam: I'm on the ground; we've just landed.

Mary: Is it possible that you're still in the air?

Sam: No, it isn't. I can see out the window right now, and we're on the ground. I know what the ground looks like! Sheesh.

In LOST WALLET and PHONE AIRPORT, intuitively, the subject can truly say “It is not possible that the wallet is on the theater” or “It is not possible that the plane is still in the air”. This contradicts the claim that not-*p* is an epistemic possibility for *S* if *S*'s evidence for *p* does not entail *p*. Indeed, according to this claim,

the epistemic impossibility claims in Lost Wallet and [Phone] Airport are incorrect – despite my apparent memory of having bought gas after leaving the theater (...) I should still say the wallet may be at the theater. Indeed, the wallet may be on the sun. Similarly, in Airport, rather than observing peevishly that he knows what the ground looks like, Sam might have more appropriately responded to Mary's question with, “Well, of course it's possible that I'm still in the air. Isn't that obvious?” After all, it should be obvious to both Mary and Sam that Sam lacks the sort of justification for the claim that he is no longer in the air that would be metaphysically incompatible with his being in the air; indeed, it is doubtful that anyone could ever have such justification. (Huemer, 2007, 127-128)

Thus, an orthodox philosopher endorsing MP – and epistemic modals infallibilism – can avoid the sceptical threat by rejecting the claim that the epistemic possibilities are the possibilities logically or metaphysically compatible with the subject's evidence.

Let us now turn to the consideration of premise 2 of EPISTEMIC MODALS INFALLIBILISM ARGUMENT. According to this premise, if it is not epistemically possible for *S* that not-*p*, then *S* satisfies the epistemic norm of action with respect to *p*. Is this correct?

First, that seems to be an obvious claim. Suppose that there is no epistemic chance for *S* that not-*p*. How could it be epistemically inappropriate for *S* to act on *p*? To explain why it is *epistemically* inappropriate for *S* to act on *p*, we need to appeal to an *epistemic* possibility of error. For example, to explain why high stakes Keith should not act on the proposition that the bank is open next Saturday, we need to appeal to the possibility that the bank has changed its hours. If that is *epistemically* impossible for Keith – if Keith is in a good enough epistemic position to eliminate this possibility – then we cannot explain why it is epistemically inappropriate for him to act on this proposition by appealing to this possibility. But if it is epistemically impossible for Keith that the bank is not open, there is no other epistemic possibility of error, and acting on this proposition cannot be epistemically inappropriate.

Second, we should note that there is a strong connection between talk about epistemic possibilities and action. As Huemer notes:

[a] desideratum for an account of epistemic possibility is that such an account should illuminate the uses of epistemic possibility talk – it should be intelligible in terms of the account why, for example, it is appropriate for one who has recently lost his wallet to be concerned about the epistemically possible locations of the wallet, as opposed, say, to its logically or nomologically possible locations. The account should likewise put us in a position to understand why the judgement “The wallet might be in the car” results in a trip out to search the car, while the judgement “The wallet is definitely not at the movie theater” forestalls a similar trip to the theater. In short, we should like an account of epistemic possibility that (insofar as this is possible) (...) explains the connection between rational action and judgements of epistemic modality. (Huemer, 2007, 121)

It is natural to explain the connection between judgements of epistemic modality and judgements of epistemically rational action by saying that it is epistemically appropriate for S to act on p if (and only if) it is epistemically impossible for S that not-p.

Third, we can argue in favour of this second premise by noting that it can explain why we never balance “p” against “It might be that not p” in our reasonings, or why if “p” is a reason one has (if we satisfy the epistemic norm of action with respect to p) then “It is possible that p” is not a reason one can appropriately use. If not-p is epistemically possible for S, then S can truly say “It is possible (/ It might be) that not-p”, and if not-p is impossible for S, S can truly say “It is impossible that not-p”. Now, assume that we are committed to the principle that it is epistemically appropriate for S to act on p if and only if it is epistemically impossible for S that not-p. It follows that if “p” is a reason S has, “It is possible that p for S” is obviously false. Thus, we never balance “p” against “It might be that not-p” in our practical reasonings because this would obviously be incoherent. If it is true that it might be that not-p, then – according to this principle – it is true that we do not satisfy the epistemic norm of action with respect to p. Hence, we should not use p as a premise. And if it is appropriate to use p as a premise, then – if we assume that we are committed to this principle – we think that it is (obviously) false that it might be that not-p. Hence, we think we should not use “It might be that not-p” as a premise. In brief, on the assumption that we have an implicit commitment to this principle, to balance “p” against “It might be that not-p” is to represent ourselves as satisfying the epistemic norm of action with respect to p and, at the same time, as not satisfying this norm with respect to p.

Therefore, if we grant that the notion of epistemic possibility should be understood in terms of the possibilities that are consistent with what the subject knows, or at least, if we grant that if S know that p then p is impossible for S, we should conclude that knowledge is sufficient for

epistemically rational action.¹¹

In summary, we might want to motivate the sufficiency of knowledge for epistemically rational action on the basis of epistemic modals infallibilism. To do this, we have first to embrace a non-sceptical view of epistemic modals infallibilism. Second, we must grant that epistemically impossible cases of error are rationally excluded from practical deliberation from an epistemic point of view. Given these two claims, if S knows that p, then it is epistemically rational for S to act on p. Knowledge is sufficient for rational action.

10.2.4 Against epistemic modals infallibilism: Rysiew and Dougherty's proposals

So far, I have considered a possible argument in favour of the sufficiency of knowledge for epistemically rational action from a non-sceptical version of epistemic modals infallibilism and from the claim that epistemically impossible cases of error can be rationally excluded from practical deliberation (from an epistemic point of view). I shall now argue that this argument is not compelling and that the data supporting MP can be explained pragmatically. I shall first consider Dougherty and Rysiew's pragmatic explanations of the data, but I shall show that their proposals are not satisfactory. In the next subsection, I shall offer a new pragmatic explanation based on the certainty norm for assertion.

Before coming to that, however, it is useful to consider some reasons to think that Moore's principle, as stated by DeRose, is false. First, as Moore himself formulates it, this principle concerns knowledge with (absolute) certainty:

[O]ne, at least, of the ways in which we use expressions of the form "It is possible that p" is such that the statement in question cannot be true if the person who makes it knows *for certain* that p is false. (Moore, 1959, my emphasize):

Although Moore is unsympathetic to the claim that knowledge does not entail absolute certainty, he acknowledges that this is an open possibility. Thus, he writes:

I have just said that if a person can ever say with truth, with regard to any particular proposition p, "I know that p is true", it follows that he can also truly say "It is absolutely certain that p is true". But an objector might perhaps say: "I admit that if a person

¹¹DeRose (1991) does not argue for the claim that epistemic possibilities should be understood in terms of possibilities consistent with what the subject knows. DeRose merely argues for MP. According to him, even if S does not know that p, that does not mean that p is epistemically possible for her. Perhaps an other member of the relevant community knows that p, or there is an available way of coming to know that p (where what counts as the relevant community or the available way of knowing that p depends on the context). This can be sufficient to make S's claim that it is possible that p false. Huemer (2007) endorses the claim that if S knows that p, then not-p is not epistemically possible for S, but he rejects the claim that if S does not know that p, then p is epistemically possible for S. Perhaps S does not know that p just because she does not believe that p. In such a case, p may still be epistemically impossible for S.

could ever truly say ‘I know with absolute certainty that p is true’ then it would follow that he could also say ‘It is absolutely certain that p is true’. But what you said was not ‘know with absolute certainty’, but ‘know’; and surely there must be some difference between knowing and knowing with absolute certainty, since, if there were not, we should never be tempted to use the latter expression. I doubt, therefore, whether a mere ‘I know that p’ does entail ‘It is absolutely certain that p’.”

Moore doubts that to explain the difference between these two expressions – ‘know’ and ‘know with absolute certainty’ – we should grant that knowledge does not entail certainty. However, he offers no reason to doubt that. His reply is just that he has been and shall be concerned “only with uses of “know”... such that “I know that p” does entail “I know with absolute certainty that p””.

Now, remember that I have argued in chapter 6 (section 2.3) that knowledge does not entail certainty. If we grant that a distinction between knowledge and (absolute) epistemic certainty has to be drawn, there is no reason to think that epistemic possibilities are possibilities consistent with what the subject merely knows. Rather, it is more natural to think of epistemic necessity in terms of epistemic certainty and epistemic possibilities in terms of epistemic uncertainty.

A second reason to think that Moore’s principle is false is that there is a tension between this principle (or non-sceptical epistemic modals infallibilism), the non-sceptical orthodox view of knowledge, and the view according to which the epistemic norm of action should be understood in terms of epistemic chances of error (see defence of premise 2 of EPISTEMIC MODALS INFALLIBILISM ARGUMENT above). Indeed, as the bank cases show, some cases of error are, given the non-sceptical orthodox framework, known by the subject not to obtain (e.g., that the bank has changed its hours). Given epistemic modals infallibilism, they thereby are epistemically impossible for her. Still, these cases of error are significant when it comes to acting rationally. Therefore, a non-sceptical orthodox philosopher endorsing non-sceptical epistemic modals infallibilism should reject the claim that whether or not you satisfy the epistemic norm of action with respect to p depends on whether there is an epistemic chance that not-p for you.

However, this is not a comfortable move. As argued above, if not-p is epistemically impossible for S, how could the possibility that not-p make S’s action on p *epistemically* irrational? How could we explain the connection between judgements of epistemic modality and rational action? And, finally, how could we explain why we never balance “p” against “It might be that not p” in our reasonings? We must grant the natural idea that if not-p is epistemically impossible for S, then it is epistemically rational for S to act on p. If this is correct, a non-sceptical orthodox philosopher should reject MP and abandon epistemic modals infallibilism about knowledge.

The fact that knowledge does not entail certainty and the fact that there is a strong connection between judgements of epistemic modality and rational actions are *prima facie* reasons to think that MP is false. However, remember that there are also considerations supporting MP. If MP is false, we must explain why these considerations fail to be decisive. The relevant considerations are

the following:

First, the following sentences – often called (first-person) concessive knowledge attributions (CKAs) – seem inconsistent:

- (1) I know that p, but it may be that not-p.
- (2) I know that p, but there is a chance that not-p.

If MP is true, we can explain why uttering these CKAs seem inconsistent. If MP is true, indeed, they are inconsistent.

Second, an inference from an utterance by S, “It is possible that not-p”, or “It may be that not-p”, to the proposition that S (believes that) she does not know that p, does not seem crazy. If MP is true, this inference is not crazy because if what S says is true, then it is a good deductive inference.

Third, if we endorse MP and the knowledge norm for assertion, we can explain the oddity of the following sentences:

- (3) p but it is possible that not-p.
- (4) S knows that p but it is possible that not-p.

Indeed, given the knowledge norm of assertion, by asserting that p you represent yourself as knowing that p, and given MP, the assertion that it is possible that not-p semantically implicates that you do not know that p.¹² Further, given the knowledge norm of assertion and the factivity of knowledge, by asserting that S knows that p, you represent yourself as knowing that p, and given MP, the assertion that it is possible that not-p semantically implicates that you do not know that p. We can thus explain the oddity of third-person CKAs.¹³

If we want to reject MP, we need to provide an alternative account of these data. The main alternative approaches which have been proposed so far have been developed by Rysiew and Dougherty. According to them, we should explain these data pragmatically on the basis of Grice’s cooperation principle (CP), and without appealing to MP. Three explanations are suggested in different places.

First, Rysiew (2001, 493) writes that CKAs

are *bound* to sound odd (...) “I know that p” standardly functions to convey the speaker’s confidence as to p and (thus) to counter a doubt as to whether p. However, the concession clause, “but, maybe, however, etc...,” *raises the then-mentioned possibility to salience*. So while “I know” communicates that the speaker is confident that p and thus that he (believes he) can rule out the salient not-p possibilities, the concession clause communicates that there’s a salient not-p possibility which the speaker *can’t* rule out.

¹²See DeRose (1991, 600)

¹³See Stanley (2005a)

Why is the speaker's confidence as to p pragmatically conveyed by a first-person 'knowledge' ascription? According to Rysiew, an assertion that p already communicates that the speaker takes herself to have a justified belief in p , and hence, an assertion "I know that p " must communicate something "over and above what's communicated by " p "; otherwise, saying "I know that p would mean violating the maxim of manner (be perspicuous, relevant, brief, orderly, etc.) and thus CP ." (Rysiew, 2001, 492). Rysiew concludes that what is communicated by such a 'knowledge' ascription must be that the ascriber takes herself to be in a position to rule out the salient doubts. Yet, in asserting that it is possible that not- p , the ascriber grants that she can't rule out the salient not- p possibility. Hence the oddity.

However, there are several problems with this approach. Firstly, as already noted in chapter 4, Rysiew does not explain why, given Grice's cooperative principle, we should expect that a first-person 'knowledge' ascription typically communicates a sufficient confidence rather than an insufficient confidence. After all, given Grice's maxim of quantity, you should assert the stronger, and given that the subject does not assert that she is certain that p , why is it that she is not communicating that she *merely* knows that p , and hence, that she is not certain that p and is not confident enough? Of course, there is no denying that the subject does convey by such an assertion that she is *sufficiently* confident that p . The point is merely that Rysiew's appeal to Grice's cooperative principles do not provide a good explanation for that.

Secondly, it is unclear that this approach can explain the paradoxical aspect of CKAs. If what is typically pragmatically imparted by a first-person 'knowledge' ascription that p , in contrast to a mere assertion that p , is that the ascriber is confident enough that p and can rule out the salient not- p possibility, typically, this implicature should be cancelled without inconsistency by the assertion that it is possible that not- p . Thus, we should not expect CKAs to sound inconsistent.

Thirdly, this approach does not explain why it does not seem crazy to infer from S 's assertion that it is possible that not- p that S does not (believe that she) know(s) that p . Indeed, it would be crazy, in this framework, to infer from S 's claim that she cannot rule out a salient possibility of error, that S doesn't (believe that she) know(s) that p . In this framework, not all salient possibilities are relevant to knowledge. Yet, the pragmatic meaning of an assertion by S that it is possible that not- p is supposed to be merely that S is not in a position to eliminate this salient possibility.

Finally, utterances " p but it is possible that not- p " and third-person CKAs (" S knows that p but it is possible that not- p ") are as odd as first-person CKAs ("I know that p but it is possible that not- p "). But, even granting that knowledge is the norm of assertion, given that this approach rejects MP, it is unclear how it can explain this. Suppose, as this proposal assumes, that an assertion "I know that p " conveys a specific confidence that a mere assertion " p " does not convey. And suppose that this is what explains why "I know that p but it is possible that not- p " is odd. Then, of course, this proposal cannot be used to explain why asserting " p but it is possible that not- p " or " S knows that p but it is possible that not- p " are also odd. On this view, in contrast to an assertion "I know

that *p*”, an assertion “*p*” or “*S* knows that *p*” is not supposed to convey a specific confidence as to whether *p*.

A second approach is proposed by Rysiew (2001, 493) and Dougherty and Rysiew (2009, 126). It relies on the rule of quantity. The idea is that a subject uttering “It is possible that not-*p*” pragmatically implicates that she does not know that *p*, given that we must assert the stronger. In effect, in the present framework, knowledge is compatible with some uneliminated possibilities of error, but it is more informative to say that these possibilities of error are unlikely, if this is the case. Yet, in the first half of a CKA, the subject asserts that she knows that *p*. Hence the infelicity.

This second approach is better in that, given the rule of quantity, it is plausible that what is typically communicated by an assertion that not-*p* is possible is that the ascriber does not know that *p*, and this can explain why it is not crazy to infer that *S* does not (believe that she) know(s) that *p* from her utterance “not-*p* is possible”.

However, as the previous attempt, this proposal cannot explain the oddity of CKAs. As Dodd writes, even if an utterance of ‘*p*’ pragmatically implicates that not-*q*, this

doesn’t mean that an utterance of ‘*p* and *q*’ is infelicitous. Thus showing that ‘*p*’ implicates not-*q* shouldn’t be taken to provide an explanation of the infelicity of ‘*p* and *q*’. So I think [Dougherty and Rysiew’s] proposal of how to explain pragmatically the infelicity of CKAs fails.

Compare with “John and Mary married and had children, but they had their children first”. This assertion does not sound incoherent, even if the first half of the assertion pragmatically conveys that they first married.

Further, this proposal does not explain why an utterance “*p* but it is possible that not-*p*” and third-person CKAs sound incoherent. Assume that by asserting *p*, the subject represents herself as knowing that *p*. Then, by asserting that *p* is possible, the subject should not pragmatically convey that she does not know that *p*, since this implicature is cancelled by the first half of the assertion. On the contrary, if MP is false, this sentence should look like a banal truth.

Compare with assertions “I believe that *p*”. Due to the rule of quantity, this assertion conveys that the assertor does not know that *p*. Now assume that knowledge is the norm of assertion. By asserting that *p*, the assertor conveys that she knows that *p*. Yet, an assertion “*p* but I believe that *p*”, although unusual, does not sound inconsistent.

Likewise, by uttering “*S* know that *p*, but it is possible that not-*p*”, the subject represents herself as knowing that *p* in the first half of the utterance, and hence, if MP is false, it is unclear how the second half of the utterance could convey that she does not know that *p*.

A third proposal is developed by Dougherty and Rysiew (2009), though. It is based on the rule of relevance and the claim that it is obvious that, with respect to (almost) every *p*, not-*p* is possible for *S*. Given this assumption and the rule of relevance, we should expect that an assertion

“It is possible that p” conveys that the possibility that not-p is particularly significant – otherwise, Dougherty and Rysiew note, why mention it? Thus, by such an assertion, the ascriber represents herself as not confident enough that p, given that she represents herself as taking the possibility that not-p seriously. On the other hand, according to Rysiew and Dougherty, the first half of a CKA is assertible only if there is no significant doubt as to p. Hence the infelicity.

In Dougherty and Rysiew’s words:

either the doubt or reservation which “it’s possible that not-p” is naturally understood as indicating is significant, or it is not. If it is, there’s a norm to hedge the assertion which comprises the first half of CKAs. This may be a generic consequence of the CP, or a consequence of the Maxim of Quality. If, however, the doubt is not significant, then the Maxim of Relation recommends that one not mention it. Either way, the explanation of the oddity of CKAs is pragmatic plainly. (Dougherty and Rysiew, 2009, 127-128)

There are several problems with this third proposal, however. First, it is still unclear how this can explain why CKAs and assertions “p but it is possible that not-p” sound inconsistent. Firstly, asserting obvious truths is not automatically infelicitous, and on this view, CKAs could merely express obvious truths.¹⁴

Secondly, let us grant that mentioning that there is a possibility that not-p conveys that this possibility is significant and should be taken seriously. Still, either the uneliminated epistemic possibility that not-p is *epistemically* significant - where this means that this possibility may well be incompatible with knowledge that p - or it is significant in a different sense, say, *practically* significant – e.g., significant to decide what to do.

We can agree that *if* the utterance “not-p is possible” imparts that the possibility that not-p is epistemically significant, then there is an incoherence between uttering the semantic content of the first half of CKAs (namely, that the assertor knows that p) and the (alleged) pragmatic meaning of uttering the second half of CKAs (namely, that the assertor does not know that p).

Likewise, if we grant that knowledge is the norm of assertion, we may agree that *if* an utterance “not-p is possible” imparts that the possibility that not-p is epistemically significant, then there is an incoherence between asserting the pragmatic meaning of the first half of the assertion “p but it is possible that not-p” (namely, that the assertor knows that p) and the (alleged) pragmatic meaning of the second half of this assertion (namely, that the assertor does not know that p).

Yet, in the present framework, the significance of epistemic possibilities of error is either epistemic – when these epistemic possibilities are relevant to knowledge – or practical – when they are relevant to rational action. But if so, there is no reason to think that the significance of the not-p possibility, which is allegedly imparted by an utterance “It is possible that not-p”, should be interpreted by the audience epistemically, rather than pragmatically. And if this significance is

¹⁴See Hawthorne (2004, 25).

interpreted pragmatically, CKAs and utterances “p but it is possible that not-p” should not sound odd.

In other words, a second problem for this approach is that it needs to assume that the interpretation we naturally give to the supposed significance of the not-p possibility, when it is asserted “It is possible that not-p”, is the epistemic one. However, there is nothing in this account which explains why this should be the case. On the contrary, it is asserted in the first half of first-person CKAs, or, assuming that knowledge is the norm of action, it is conveyed in the first half of the assertion “p but it is possible that not-p” and third-person CKAs, that the assertor knows that p. This should lead the audience toward a non-epistemic interpretation of the significance of the not-p possibility imparted by the utterance of the second half of these assertions. An assertion that I know that p, or that p, should cancel the possible implication that I (believes that I) do not know that p – allegedly generated by an assertion that it is possible that not-p.

Dougherty and Rysiew (2011) see the problem, and they write:

But of course the hearer who wishes to understand the mention-worthiness of “possibly not-p” might very naturally rely on the fact that, in a CKA, the relevant kind of significance is most likely epistemic significance. (After all, the mentioned possibility immediately follows upon a claim that p is known.)

Quite the contrary, in fact. The fact that the subject just asserted that she knows that p should make it clear that she takes the significance to be of a non-epistemic sort. Further, we cannot explain in this way the oddity of “p but it is possible that not-p” because in these assertions, there is no claim that p is known. Hence, the mentioned possibility does not follow such a claim.

This second problem raises a third problem regarding the explanation why it is not crazy to infer from an assertion “It is possible that not-p” that the assertor does not (take herself to) know that p. In asserting “It is possible that p”, it is said, the assertor indicates that there is a significant doubt that p, and hence, a doubt that may well prevent knowledge. Therefore, we can safely infer from such an assertion that the subject does not (take herself to) know that p. But, again, if the significance of the possibility of error can in principle be merely practical, it would be crazy to infer from such an assertion that the assertor does not (take herself to) know that p.

A fourth problem of this proposal is that it relies on the assumption that it is obvious that, with respect to (almost) every p, not-p is possible for S. This assumption is based on the idea that epistemic possibilities for S are cases of error logically or metaphysically compatible with S’s evidence. But recall that this is a disputable claim, as AIRPORT and LOST WALLET above show. It seems that there are some metaphysically possible cases in which the subject has evidence E for p and p is false, although such cases should not count as epistemically possible for S.

Of course, Dougherty and Rysiew will presumably explain AIRPORT and LOST WALLET by invoking the pragmatic appropriateness of saying, in such cases, “It is not possible that I am still

in the air” or “It is not possible that my wallet is at the movie theatre”. According to their line of thought, these claims are, strictly speaking, false. Still, they may say that these claims convey that the possibility in question is not *significant*. Here, the significance has to be practical. But then, this raises a new problem. How are we to explain that “It is possible that not-p” conveys that not-p is epistemically significant, whereas “It is impossible that not-p” conveys that not-p is practically insignificant? Further, in such cases, “I might still be in the air / my wallet might be on the moon, but that’s too unlikely” just does not seem true.

10.2.5 Against epistemic modals infallibilism: the certainty norm of assertion

I have argued that Rysiew and Dougherty’s proposals do not succeed in explaining all the data in support of MP, and, in particular, the apparent inconsistency of CKAs. Their proposals are based on a Gricean approach. Still, a different pragmatic explanation of the data can be given. In chapter 6, I have defended a certainty norm for assertion. Endorsing this norm allows us to explain all the data offered in support of MP pragmatically.

For some data, the explanation is straightforward. Consider “p, but it is possible that not-p”. In the first half of the assertion, the subject asserts that p. Given the certainty norm of assertion, she represents herself as certain that p. Yet, in the second half of the assertion, she asserts that it is possible that not-p, or that there is a chance that not-p. On any account of epistemic possibilities, if p is certain for S, then not-p is impossible for S. Hence the infelicity.

Consider CKAs. In the first half of first-person CKAs, the subject asserts that she knows that p. Given the certainty norm of assertion, she represents herself as certain that she knows that p, and by factivity of knowledge, as certain that p. Yet, in the second half of these CKAs, she asserts that it is possible that not-p, or that there is a chance that not-p. On any account of epistemic possibilities, if p is certain for S, then not-p is impossible for S. Hence the infelicity of first-person CKAs.

This approach derives the infelicity of first-person CKAs from the infelicity of asserting “p, but it is possible that not-p”. Hence, it can also explain, given the factivity of knowledge, the oddity of third-person CKAs – why “S knows that p but it is possible that not-p” is infelicitous – without appealing to MP and the knowledge norm for assertion.

This approach does not provide a straightforward explanation why it does not seem crazy to infer from an utterance “It is possible that not-p” that the assertor does not (take herself to) know that p, though. For, after all, it is granted that the fact that p is not certain for S does not entail that p is not known by S.

In reply to this, we should note first that this approach can explain why it is not crazy to infer from an utterance “It is possible that not-p” that the assertor does not know *for certain* (or takes

herself not to know for certain) that p . Now, if this is *that* inference that we need to explain, the present approach does well, because the fact that S does not know for certain that p deductively follows from the fact that not- p is possible for S .¹⁵

A second point to note is that if an inference is not crazy but defeasible, then the fact that it is not crazy does not suggest that it relies on a semantic fact. Thus, in order to show that MP is true, we need to show that the inference from *S has asserted 'It is possible that not-p'* to *S does not (believe that she) know(s) that p* cannot be defeated by the consideration that S actually *does* (believe that she) know(s) that p , although she is uncertain of that because (she believes that) there is a chance for her that not- p .

In effect, if MP is true, S cannot know that p if there is a chance that not- p for her. If this is correct, if S believes that there is a chance that not- p for her, she should believe that she does not know that p . It would be irrational for her to believe that she knows that p , or to suspend judgement as to whether she knows that p , when she believes that there is a chance for her that not- p .

However it seems that there are cases in which S can rationally believe that she knows that p (or suspend judgement as to whether she knows that p) while believing that there is an epistemic chance for her that not- p . This belief (or the fact) that there is an epistemic chance for her that not- p can even be used to explain why S merely believes *unconfidently* that she knows, or why she suspends her judgement on this issue.

To illustrate compare the two following dialogues:

DIALOGUE 1

John - There is a small chance that the bank has changed its hours.

Mary - You mean you do not know whether the bank is open?

John - No, I think I know. But it is not certain. I prefer to check.

DIALOGUE 2

John - There is a small chance that the bank has changed its hours.

Mary - You mean you do not know for certain whether the bank is open?

John - No, I think I know for certain. But it is not certain. I prefer to check

John's answer to Mary in DIALOGUE 1 seems perfectly understandable. There seems to be cases

¹⁵It pays to note that when Moore writes on this issue, he very often uses 'know for certain', rather than the mere 'know', as in the following passage:

[If] I were to say "It is possible that Hitler is dead at this moment" this would naturally be understood to mean exactly the same as if I said "Hitler *may* be dead at this moment". And is it not quite plain that if I did assert that Hitler *may* be dead at this moment part at least of what I was asserting would be that I personally did not know for certain that he was not dead? Consequently if I were to assert now "It is possible that I am not standing up" I should naturally be understood to be asserting that I do not know for certain that I am. And hence, if I do not know for certain that I am, my assertion that it is possible that I'm not would be false. (Moore, 1959)

in which a subject can know that p although it is uncertain for her that she knows that p . But this suggests that it is uncertain for her that p , for if p were certain for her, arguably, there would be no room for doubting that p , and hence, for doubting that she knows that p . In other words, it can make sense to doubt whether one knows that p . If one doubts whether one knows that p , arguably, that is because p is not epistemically certain for one. But this means that one assumes that one can know that p even if there is a chance for one that not p . Otherwise, why doubting that one knows that p ? It should be clear that one does not know that p since there is a chance for one that not- p .

In contrast, it is difficult to make sense of John's answer to Mary in DIALOGUE 2. It seems to me that we can explain this by the claim that if it is uncertain for S that she knows (for certain) that p , then S does not know for certain that p .

In any event, DIALOGUE 1 shows that even if it is not crazy to infer from S 's assertion "It is possible that not- p " that S does not (believe that she) know(s) that p , this cannot be explained by MP. This inference is defeated in cases in which the assertor is uncertain that p but has an uncertain belief that she knows that p , or suspends judgement as to whether she knows that p .

Still, if we grant that MP is false, how are we to explain that this inference is not crazy? Although the certainty norm of assertion offers no straightforward explanation here, an appeal to the Gricean rule of quantity can help. When the assertor says "It is possible that not- p ", she does not assert what would be, if true, more informative, namely "I know that p ". Given the rule of quantity and the certainty norm of assertion, this means that the assertor does not believe with certainty that she knows that p (otherwise she would have said it). Now, it is very likely that if the subject does not believe with certainty that she knows that p , that is because she actually does not know that p . Hence the quick, but defeasible, inference from an assertion "It is possible that not- p " to "The assertor (believes that she) does not know that p ".

In sum, the claim that certainty is the norm of assertion can be used to provide a pragmatic explanation of the data advanced in favour of MP. It is worth emphasizing that, in contrast to Rysiew and Dougherty's proposals criticised above, the present approach shows why CKAs and assertions " p but it is possible that not- p " are bound to sound odd. Asserting the second half of the assertion is not liable to cancel the pragmatic implication of asserting the first half of the assertion, for the pragmatic implication is not a Gricean implicature, but an implication related to the norm of assertion. Such an implication is not cancellable.

Finally, let me stress that, contrary to Dougherty and Rysiew's third proposal, the present approach does not rely on the claim that, with respect to (almost) every p , not- p is possible for S . Again, this claim is based on the idea that epistemic possibilities for S are cases of error logically or metaphysically compatible with S 's evidence, which is a disputable claim (see PHONE AIRPORT and LOST WALLET).

To recap: We can undermine the main arguments in favour of MP pragmatically by appealing

to the certainty norm of assertion. If so, we can reject EPISTEMIC MODALS INFALLIBILISM ARGUMENT. And if knowledge is not always sufficient for action, we cannot reject the non-sceptical orthodox view of knowledge on the basis of the idea that for (almost) any proposition p , there is a practical situation PS in which a particularly strong epistemic position with respect to p is required to be in a good enough epistemic position – epistemically speaking – to act on p .

Now, we must make sense of the connection between judgements of epistemic modality and epistemically rational action, while preserving the interpretation of the epistemic norm of action in terms of epistemic chances

In the remaining of this chapter, I shall defend a certainty norm of action, similar to the certainty norm of assertion defended in chapter 6. According to this norm, it is appropriate for S to act for the reason that p if and only if p is certain for S. As I take epistemic certainty to be context-sensitive, I shall grant that epistemic (im)possibilities are context-sensitive. In this way, I will make sense of the connection between judgements of epistemic modality and action while avoiding (1) the claim that any case of error metaphysically or logically compatible with the subject's evidence is an epistemic possibility for her (see PHONE AIRPORT and WALLET above) as well as (2) the claim that any case of error metaphysically or logically incompatible with what the subject knows is epistemically impossible for her.

The natural interpretation of the epistemic norm of action in terms of epistemic chances will then be preserved. Given that we can reject MP and the sufficiency of knowledge for action, this approach does not impugn the non-sceptical orthodox view of knowledge.

10.3 The certainty account

10.3.1 Eight desiderata for an account of the epistemic norm of action

So far, I have argued against the sufficiency of knowledge for rational action. I have presented cases in which it is plausible that a subject knows that p but should not act on p because there is an epistemic possibility that not- p and it is epistemically appropriate for her to premise that not- p is possible. I have also argued that, given the certainty norm of assertion, the main arguments supporting Moore's principle – the principle according to which if you know that p , then not- p is not (epistemically) possible for you – are not compelling. If Moore's principle is false, there is no reason to think that knowledge is always sufficient for rational action. Epistemic possibilities of error which do not destroy knowledge may play a role when it comes to act rationally on a certain proposition. If this is right, the argument against the non-sceptical orthodox account of knowledge based on the sufficiency of knowledge for rational action fails.

I shall now argue in favour of a certainty norm for rational action. To do so, I shall place what I take to be reasonable constraints on the account of the epistemic norm of action we should propose,

and I shall show that the certainty account of action satisfies all these constraints.

Some of the constraints I accept rely on claims that everybody accept (e.g., that our epistemic evaluations of action shift with the practical context), others on claims that almost all philosophers find plausible and which can be accepted by default (e.g., that scepticism is false and that the orthodox account of knowledge is right), yet others on claims based on specific arguments previously proposed in various chapters (e.g., that knowledge is not always sufficient for action). If I am right, then, any account which satisfies all these constraints is better than any other which does not.

Of course, philosophers disagreeing with some of these underlying claims will reject the associated constraints. However, these claims are independently reasonable. Therefore, it would be inadequate to reject the associated constraint just because one endorses some other epistemic norm.

What are these constraints? First of all, it is accepted by everybody that orthodoxy about knowledge is the default view. Hence the first desideratum to satisfy is what I shall call the ‘orthodox constraint’:

Orthodox constraint: Knowledge and the meaning of ‘know’ do not shift with the practical context.

A second constraint that a plausible account must satisfy is to avoid leading to some sort of scepticism. That is, such an account must be compatible with the idea that we often know what we ordinarily take ourselves to know, and that there are many things with respect to which we satisfy the epistemic norm of action. Call this desideratum the ‘non-sceptical constraint’:

Non-sceptical constraint: Knowledge and epistemically rational actions are pervasive.

In chapter 7, I have considered and distinguished various notions of epistemic evaluations of action and various features of these evaluations. These distinctions have been used in various arguments in favour or against one or another account. On the basis of these distinctions, we can also lay out further constraints.

The account must be an account of the right thing, i.e. treating p as a reason (premise that p , act for the reason that p , act on p) in contrast to acting as if p . In other words, this account must say what epistemic condition a subject must satisfy with respect to a certain consideration for this consideration to be a good motivating reason to act. For example, it may be epistemically rational for me to act as if it is raining, if the proposition that it is raining is slightly probable, even though it would not be rational in such a case to defend my action by the consideration that it is raining – rather than with the consideration that there is a possibility that it is raining (see chapter 7 section 3.3). Call this desideratum the ‘right thing constraint’:

Right thing constraint: The epistemic norm of action determines the good motivating reasons.

This account must also be an account of when it is *epistemically* appropriate to act on a proposition, in contrast to when it is *practically* appropriate to act on a proposition given such or such epistemic position. Suppose an evil demon kills you if you act on a proposition if and only if this proposition is absolutely certain. It is not practically appropriate for you to act on this proposition if this proposition is absolutely certain, although that is epistemically permissible (see chapter 7 sections 3.3 and 4.1). In effect, the reason why you should not treat this consideration as a reason has nothing to do with the consequences of a possible error. Call this the ‘right sort of account constraint’:

Right sort of account constraint: The epistemic norm of action determines the good motivating reasons with respect to the possibility of error.

Further, this account must be compatible with, or even provide an explanation of, the prominence of ‘know’ in epistemic evaluations of actions, the fact that we sometimes appropriately use weaker and stronger terms, and the fact that ‘knowledge’ ascriptions are always relevant and decisive when it comes to act (see in particular chapter 8). Call this the ‘conversational constraint’:

Conversational constraint: Illuminate the way in which we use epistemic words in our epistemic assessments of action.

In addition, this account must make sense of the fact that epistemic evaluations of action vary with practical factors. This variation is clear if we consider, for example, DeRose’s bank cases (see chapter 9). Call this the ‘variation constraint’:

Variation constraint: Illuminate why our epistemic assessments of action vary with practical factors.

This account must also make sense of, or at least be compatible with a plausible explanation of, the connection between judgements of epistemic modality and rational action. Suppose John is looking for his keys. A (true) judgement “They cannot possibly be in my jacket” is associated with the rationality of John not checking in his jacket, while a judgement “There might be in my car” is associated with the rationality of checking in the car (see section 1.3. above). That is, these judgements are associated with treating “The keys are not in my jacket” as a reason to look elsewhere, and treating “They might be in my car” as a reason to check in the car. Call this the ‘epistemic chance constraint’:

Epistemic chance constraint: Illuminate the connection between judgements of epistemic modality and judgements of rational action.

Besides, I have argued that knowledge is necessary but not (always) sufficient for rational action (see chapters 8, 9 and the present chapter above). If this is right, we must provide an account of the epistemic norm of action which is compatible with this claim. Call this the ‘K_{necc} and K_{not-suff} constraint’:

Knec and Knot-suff constraint: Knowledge is necessary but not (always) sufficient for rational action.

If we can propose an account which satisfies these eight constraints, and on the assumption that there is no further reasonable constraint satisfied by some other account which is not satisfied by the proposed account, then, I suggest, the proposed account will be better than the others.

10.3.2 The eight desiderata and the main accounts: summary

I have argued against the main accounts proposed in the literature by showing that they cannot satisfy all the constraints above. It is useful to summarize why I think that, in light of these constraints, we should reject these accounts.

I have argued against the knowledge account by arguing that it must reject the orthodox constraint, or else, it cannot satisfy the variation constraint. Indeed the two ways in which the advocates of the knowledge account propose to satisfy the variation constraint – by appealing to a distinction between primary and derivative norms or to that between insufficient and sufficient reason – are unsatisfactory (see the present chapter section 1.1).

Further I have argued that if one endorses an insensitivist view of the epistemic norm of action, then one cannot satisfy the non-sceptical constraint. In effect, if one endorses an insensitivist view of the epistemic norm of action, one must grant that the epistemic norm requires to satisfy maximum epistemic standards. The argument for this is the following. If the epistemic norm is satisfied, on the assumption that it is insensitive, raising the stakes makes no difference. But for any non maximum-standards norm, we can raise the stakes so that it becomes irrational to act on the target proposition. Hence, the norm must be a maximum-standards norm. To put it in a different way: If the insensitive norm is also a non maximum-standards norm, we can exploit the fact that it is not a maximum-standards norm to show that it is insufficient in some cases, in particular, when the stakes are raised. Therefore, the insensitive norm cannot be a non maximum-standards norm (see chapter 9 section 9.4). Hence, if the advocates of the knowledge account accept the orthodox constraint, given that, on their view, knowledge is the epistemic norm of action, they should grant that knowledge and epistemically rational action require to satisfy (insensitive) maximum epistemic standards. It follows that this approach violates the non-sceptical constraint.

I did not consider non-factive accounts in which knowledge-level justification only is necessary and sufficient. However, it is very plausible that accounts of this sort will face the same criticism as the knowledge account.

I have argued against the warrant account by arguing that it is at odd with conversational facts regarding the use of ‘know’, and hence, that it does not satisfy the conversational constraint (see chapter 8 section 3). Also, it seems to me that some of the arguments used in favour of this account, for example, cases like KICKOFF (see chapter 8 section 4.2) do not satisfy the right thing

constraint, for they rely on a confusion between decisive possibilities and good motivating reasons. Further, of course, this account rejects the necessity of knowledge for rational action, and hence, the Kneec and Knot-suff constraint (see chapter 8 section 2.3 for a defence of the necessity of knowledge). Finally, we may add that it is unclear how this account can satisfy the epistemic chance constraint.

I did not consider pluralist accounts in which there are several conditioned norms (rather than a single norm with a context-sensitive content), some weaker than knowledge and some stronger than knowledge. However, it is very plausible that accounts of this sort will face the same kind of criticism as the warrant account.

10.3.3 The eight desiderata and the certainty norm

Consider the idea that certainty is the norm of action. In chapter 6, I have developed and defended a certainty account of assertion. Assume that a similar account can be proposed with respect to action. According to such an account of action,

CNAc-R. It is epistemically appropriate for S to act on p in C if and only if p is epistemically certain for S in C.

To illustrate, consider the following example. You are offered a bet on whether Dublin is the capital of Ireland. If you take the bet and Dublin is the capital of Ireland, you win one dollar. If you take the bet and Dublin is not the capital of Ireland, you are tortured to death. Presumably, you will decline taking such a bet. It seems that it is rational for you to avoid the bet because the proposition that Dublin is the capital of Ireland is not epistemically certain for you. (If asked why you decline the bet, you will presumably say something like “After all, I might be wrong”.) Even if you know and outright believe this proposition (see ch. 5 section 4.3 problem 4 and chapter 6 section 2.3), this proposition is not justified enough to meet the standards of certainty in these circumstances. In contrast, in low stakes contexts, you are willing to treat the proposition that Dublin is the capital of Ireland as a reason to act. Indeed, you know that this proposition true, and in low stakes contexts, knowledge may often be sufficient for epistemic certainty.

How does this proposal fare with respect to the eight desiderata that I have laid out above?

This proposal satisfies the orthodox constraint. Indeed, it is assumed that knowledge does not entail epistemic certainty. Thus, even if certainty shifts with the practical context, or even if ‘certain’ is context-sensitive, that does not imply that knowledge depends of the practical context, or that ‘know’ is context-sensitive.

This proposal also meets the non-sceptical constraint. First, it is not committed to maximal infallibilism about knowledge. So it does not lead to the idea that we know hardly anything. Second, as the notion of certainty is taken to be context-sensitive, in the same way as the certainty account of assertion is compatible with the claim that many assertions are epistemically appropriate,

the certainty account of action is compatible with the idea that many actions are performed for epistemically appropriate reasons.

Further, this account satisfies the right thing and the right sort of account constraints. According to the right thing constraint, the account must tell us when a consideration can be a good motivating reason. According to the right sort of account constraint, the account must tell us when a consideration is a good motivating reason *with respect to the possibility of error*. If p is certain for S , then there is not epistemic possibility for S that not- p . Hence, we can explain why it is appropriate for S to ignore the possibility that not- p in her practical decision. In contrast, if p is not epistemically certain for S , then there is a possibility of error for S . But if so, intuitively, S should not act on p , but on the mere chance that p .

This approach also satisfies the conversational constraint. As explained in chapter 6, given the certainty norm of assertion, the certainty norm of action and the factivity of knowledge, we can show why ‘knowledge’ ascriptions are always relevant and decisive for action. By asserting ‘I know that p ’, the assertor represents the proposition that she knows that p as epistemically certain for her (certainty norm of assertion), and hence, she represents the proposition that p as epistemically certain (factivity of knowledge). If p is epistemically certain for S , then it is rational for S to act on p (certainty norm of action). Given the certainty norm of action, we can also explain why terms stronger than ‘know’ are sometime used. The use of weaker terms can be explained in terms of excuses (see chapter 8 section 4.3).

This account satisfies the variation constraint, since the epistemic standards of epistemic certainty are taken to shift with the subject’s practical context. On this view, in the low stakes case, Keith knows with certainty that the bank is open. Yet, in the high stakes case, the epistemic standards for certainty are higher, and if Keith maintains his outright belief that the bank is open, he still knows, but without certainty, that the bank is open. Therefore, it is epistemically rational for Keith to act on this proposition in the low stakes case but not in the high stakes case.

This proposal can also make sense of the connection between judgements of epistemic modality and action, and hence, it satisfies the epistemic chance constraint. When the question is whether to act on p , it can be relevant to say “It is possible that not- p ” because, according to this account, if this is true, we should not act on p . And it can be relevant to say “It must be the case that p ” or “It is impossible that not- p ”, because if it is epistemically impossible that not- p , it is appropriate to act on p . Also, note that this proposal explains why, if p is a reason one has, p cannot be defeated by the mere possibility that not- p , and why it makes no sense to simultaneously treat p as a reason for acting and to raise doubts as to whether p .¹⁶ When one treats p as a reason to act, one treats p as epistemically certain. According to this view, it is appropriate to treat p as epistemically certain if and only if p is epistemically certain.

Finally, this account satisfies the Kneec and Knot-suff constraint. If p is epistemically certain for

¹⁶See Fantl and McGrath (2009, 169).

S, then S is in a position to know that p. If so, if S is rational, when she considers whether p, she believes that p and thereby knows that p. This implies that if S appropriately treat p as a reason and is motivated by p, and hence, consider whether p, S knows that p. Still, as knowledge does not entail certainty, S may know that p even if it is not appropriate for S to act on p.

In sum, this certainty account of action, in which we adopt a context-sensitive view of the epistemic standards required by certainty, is compatible with the eight constraints that I have placed on what should be a satisfying account of the epistemic norm of action. The main rival accounts proposed so far do not satisfy some of these desiderata, and there are no clear further desiderata satisfied by these other accounts which are not satisfied by this certainty account. Therefore, I invite the conclusion that we should adopt the certainty account of action.

If the certainty account of action is the correct account, then there is no good argument from the normative connection between knowledge and action against the orthodox view of knowledge. Remember that this argument starts from the claim that there are cases in which the cost of error is so high that a very strong epistemic position with respect to the relevant propositions is required before treating them as our reasons for action. Next, the argument assumes that knowledge is always sufficient for action. It is then argued that either knowledge requires a very strong epistemic position, which leads to sceptical consequences, or that the epistemic standards of knowledge shift with the practical context. As scepticism is implausible, the argument continues, we should endorse the second option: when the cost of error is raised and it becomes irrational to treat the relevant propositions as reasons, we lose knowledge. By embracing a context-sensitive certainty norm of action, we reject the assumption that knowledge is always sufficient for action. Therefore, we escape this line of argument.

10.4 Conclusion

In this chapter, I have argued against the sufficiency of knowledge for rational action by presenting cases in which a subject knows that p but cannot appropriately act on p. Further, I have undermined a possible argument in favour of the sufficiency of knowledge based on epistemic modals infallibilism. I have used the certainty account of assertion (defended in chapter 6) to explain pragmatically the data put forth in favour of the claim that if S knows that p, then not-p is impossible for S. Finally, I have defended a certainty account of action similar to the certainty account of assertion by showing that it satisfies eight reasonable desiderata. If we adopt this account, we can avoid the argument against the orthodox view of knowledge from the supposed normative connection between knowledge and action. I conclude that the non-sceptical orthodox view of knowledge is safe.

Bibliography

- Alvarez, M. (2010). *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford University Press.
- Anscombe, G. E. M. (2000). *Intention*. Harvard University Press, 2 edition.
- Arpaly, N. (2003). *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford University Press.
- Arpaly, N. and Schroeder, M. (2012). Deliberation and acting for reasons. *Philosophical*, 121(2):209–239.
- Audi, R. (2001). *The Architecture of Reason: the Structure and Substance of Rationality*. Oxford University Press.
- Austin, J. L. (1959). Other minds. In Flew, A. G., editor, *Logic and Language, Second Series*. Oxford: Blackwell.
- Ayer, A. J. (1973). Wittgenstein on certainty: Ayer. *Royal Institute of Philosophy Supplement*, 7:226–245.
- Bach, K. (2005). The emperor’s new ‘knows’. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, pages 51–89. Oxford University Press.
- Bach, K. (2008). Applying pragmatics to epistemology. *Philosophical Issues*, 18(1):68–88.
- Bach, K. (2010). Knowledge in and out of context. In Campbell, J. K., Rourke, M. O., and Silverstein, H. S., editors, *Knowledge and Skepticism*, pages 105–36. MIT Press.
- Baron, M. (2005). Justifications and excuses. *Ohio State Journal of Criminal Law*, 2(2):387–406.
- Benton, M. A. (2011). Two more for the knowledge account of assertion. *Analysis*, 71(4):684–687.
- Benton, M. A. (2014). Knowledge norms. In *Internet Encyclopedia of Philosophy*.
- Blaauw, M. (2012). Reinforcing the knowledge account of assertion. *Analysis*, 72(1):105–108.

- Botterell, A. (2009). A primer on the distinction between justifications and excuse. *Philosophy Compass*, 4(1):172–196.
- Broome, J. (1999). Normative requirements. *Ratio*, 12(4):398–419.
- Broome, J. (2013). *Rationality through Reasoning*. Wiley-Blackwell.
- Brown, J. (2006). Contextualism and warranted assertability manoeuvres. *Philosophical Studies*, 130:407–35.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reason. *Noûs*, pages 167–189.
- Brown, J. (2011). Fallibilism and the knowledge norm for assertion and practical reasoning. In Brown, J. and Cappelen, H., editors, *Assertion: New Philosophical Essays*. Oxford University Press.
- Brown, J. (2013). Impurism, practical reasoning, and the threshold problem. *Noûs*, 47(1):179–192.
- Burnett, H. (2014). A delineation solution to the puzzles of absolute adjectives. *Linguistics and Philosophy*, 37(1):1–39.
- Cappelen, H. and Lepore, E. (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Wiley-Blackwell.
- Carruthers, P. (2014). The fragmentation of reasoning. In Quintanilla, P., Mantilla, C., and Cépeda, P., editors, *Cognición social y lenguaje. La intersubjetividad en la evolución de la especie y en el desarrollo del niño*. Pontificia Universidad Católica del Perú.
- Chisholm, R. M. (1988). The indispensability of internal justification. *Synthese*, 74(3):285–96.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2:91–123.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. In Tomberlin, J., editor, *Philosophical Perspectives*, pages 57–89. Blackwell.
- Cohen, S. (2010). Luminosity, reliability, and the sorites. *Philosophy and Phenomenological Research*, 81(3):718–730.
- Conee, E. and Feldman, R. (2001). Internalism defended. *American Philosophical Quarterly*, 38(1):1–18.
- Conee, E. and Feldman, R. (2004). *Evidentialism*. Oxford University Press.

- Craig, E. (2010). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.
- Crisp, T. M. (2005). Hawthorne on knowledge and practical reasoning. *Analysis*, 65(286):138–140.
- Cruse, D. (1980). Antonyms and gradable complementaries. In Kastovsky, D., editor, *Perspectiven der Lexikalischen Semantik: Beitrage zum Wuppertaler Semantikkolloquium vom 2-3, Dec. 1977*, pages 14–25. Bouvier, Bonn.
- Dancy, J. (2000). *Practical Reality*. Oxford University Press.
- Dancy, J. (2004). *Ethics Without Principles*. Oxford University Press.
- Dancy, J. (2009). Reasons and rationality. In Robertson, S., editor, *Spheres of Reason*, pages 93–112. Oxford University Press.
- Davidson, D. (1963). Actions, reasons, and causes. *Journal of Philosophy*, 60(23):685–700.
- Davidson, D. (1980). *Essays on Actions and Events*. Oxford University Press.
- DeRose, K. (1991). Epistemic possibilities. *Philosophical Review*, 100(4):581–605.
- DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104(1):1–52.
- DeRose, K. (2009). *The Case for Contextualism*. Oxford University Press.
- Dodd, D. (2010). Confusion about concessive knowledge attributions. *Synthese*, 172(3):381–396.
- Dougherty, T. and Rysiew, P. (2009). Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research*, 78(1):123–132.
- Dougherty, T. and Rysiew, P. (2011). Clarity about concessive knowledge attributions: Reply to dodd. *Synthese*, 181(3):395–403.
- Douven, I. (2006). Assertion, knowledge, and rational credibility. *Philosophical Review*, 115(4):449–485.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67 (24):1007–1023.
- Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 49 (1):1–22.
- Engel, P. (2004). Truth and the aim of belief. In Gillies, D., editor, *Laws and models in science*, pages 77–97. London: King's College Publications.
- Engel, P. (2009). Pragmatic encroachment and epistemic value. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford University Press.

- Engel, P. (2013). In defence of normativism about the aim of belief. In Chan, T., editor, *The aim of belief*. Oxford University Press.
- Fantl, J. (2003). Modest infinitism. *Canadian Journal of Philosophy*, 33(4):537–562.
- Fantl, J. and McGrath, M. (2002). Evidence, pragmatic and justification. *The Philosophical Review*, 111:67–94.
- Fantl, J. and McGrath, M. (2007). On pragmatic encroachment on epistemology. *Philosophy and Phenomenological Research*, 75:558–89.
- Fantl, J. and McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.
- Fantl, J. and McGrath, M. (2011). Pragmatic encroachment. In Bernecker, S. and Pritchard, D., editors, *The Routledge Companion to Epistemology*. Routledge.
- Fantl, J. and McGrath, M. (2014). Do practical matters affect whether you know? In Steup, M. and Turri, J., editors, *Contemporary Debates in Epistemology*, pages 88–95. Blackwell.
- Fassio, D. (2015). Belief, aim of. In *Internet Encyclopedia of Philosophy*.
- Fassio, D. (forthcoming). Is there an epistemic norm of practical reasoning? *Philosophical Studies*, pages 1–30.
- Feldman, F. (1983). Obligations – absolute, conditioned and conditional. *Philosophia*, 12(3-4):257–272.
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60(3):667–695.
- Finlay, S. and Shroeder, M. (2012). Reasons for action: Internal vs. external. In Zalta, E.-N., editor, *Stanford Encyclopedia of Philosophy*.
- Firth, R. (1967). The anatomy of certainty. *Philosophical Review*, 76(1):3–27.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Harvard University Press.
- Foley, R. (1993). *Working Without a Net: A Study of Egocentric Epistemology*. Oxford University Press.
- Foley, R. (2002). Conceptual diversity in epistemology. In Moser, P. K., editor, *The Oxford Handbook of Epistemology*, pages 177–203. Oxford University Press.
- Foley, R. (2005). Justified belief as responsible belief. In Sosa, E. and Steup, M., editors, *Contemporary Debates in Epistemology*, pages 313–26. Blackwell.

- Frankish, K. (2004). *Mind and Supermind*. Cambridge University Press.
- Ganson, D. (2008). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139(3):441–458.
- García-Carpintero, M. (2004). Assertion and the semantics of force-markers. In Bianchi, C., editor, *The Semantics/Pragmatics Distinction*, pages 133–166. Csi.
- Gardner, J. (2007). *Offences and Defences: Selected Essays in the Philosophy of Criminal Law*. Oxford University Press.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178(3):529–547.
- Gerken, M. (2013). *Epistemic Reasoning and the Mental*. Palgrave Macmillan (Innovations in Philosophy).
- Gerken, M. (2015). The roles of knowledge ascriptions in epistemic assessment. *European Journal of Philosophy*, 23(1):141–161.
- Gerken, M. (forthcoming). *On Folk Epistemology*. Oxford University Press.
- Gibbons, J. (2013). *The Norm of Belief*. Oxford University Press.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73 (20):771–791.
- Goldman, A. I. (1999). Internalism exposed. *Journal of Philosophy*, 96(6):271–293.
- Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.
- Harman, G. (1986). *Change in View*. MIT Press.
- Harman, G. (2002). 'the internal critique. In Gabbay, D. M., editor, *Handbook of the Logic of Argument and Inference: The Turn Towards the Practical*, pages 171–186. Elsevier.
- Hawthorne, J. (2004). *Knowledge and Lottery*. Oxford University Press.
- Hawthorne, J. and Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105 (10):571–590.
- Hazlett, A. (2009). Knowledge and conversation. *Philosophy and Phenomenological Research*, 78 (3):592–620.
- Henderson, D. (2009). Motivated contextualism. *Philosophical Studies*, 142(1):119–131.
- Henderson, D. (2011). Gate-keeping contextualism. *Episteme*, 8(1):83–98.

- Huemer, M. (2007). Epistemic possibility. *Synthese*, 156(1):119–142.
- Ichikawa, J. J. (2012). Knowledge norms and acting well. *Thought: A Journal of Philosophy*, 1(1):49–55.
- James, W. (1897). The will to believe. In *The Will to Believe and Other Essays in Popular Philosophy*, pages 1–15. New York: Longmans, Green, and Co.
- Jeffrey, R. C. (1970). Dracula meets wolfman: Acceptance vs. partial belief. In Swain, M., editor, *Induction, Acceptance, and Rational Belief*, pages 157–185. D. Reidel.
- Kaplan, D. (1989). Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In *Themes from Kaplan*, pages 481–563. Oxford University Press.
- Kennedy, C. and McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81(2):345–381.
- Klein, P. D. (1981). *Certainty, a Refutation of Scepticism*. University of Minnesota Press.
- Klein, P. D. (2000). Contextualism and the real nature of academic skepticism. *Philosophical Issues*, 10(1):108–116.
- Kolodny, N. (2005). Why be rational? *Mind*, 114(455):509–563.
- Kolodny, N. and MacFarlane, J. (2010). Ifs and oughts. *Journal of Philosophy*, 107(3):115–143.
- Korsgaard, C. (1986). Skepticism about practical reason. *Journal of Philosophy*, (83):5–25.
- Korsgaard, C. (2008). *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford University Press.
- Korsgaard, C. M. (1996). *The Sources of Normativity*. Cambridge University Press.
- Kruglanski, A. W. and Webster, D. M. (1996). Motivated closing of the mind: "seizing" and "freezing.". *Psychological Review*, 103(2):263–283.
- Kvanvig, J. (2009). Assertion, knowledge, and lotteries. In Pritchard, D. and Greenough, P., editors, *Williamson on Knowledge*, pages 140–160. Oxford: Oxford University Press.
- Kyburg, H. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press.
- Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese*, 158(3):345–361.
- Lawlor, K. (2013). *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford University Press.

- Lehrer, K. (1990). *Theory of Knowledge*. Routledge.
- Lenman, J. (2009). Reasons for action: Justification vs. explanation. In Zalta, E.-N., editor, *Stanford Encyclopedia of Philosophy*.
- Levin, J. (2008). Assertion, practical reason, and pragmatic theories of knowledge. *Philosophy and Phenomenological Research*, 76(2):359–384.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8:339–59.
- Lewis, D. (1996). Elusive knowledge. *Australian Journal of Philosophy*, 74:549–67.
- Littlejohn, C. (2009). Must we act only on what we know?
- Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, C. (2014). The unity of reason. In Littlejohn, C. and Turri, J., editors, *Epistemic Norms: New Essays on Action, Belief and Assertion*. Oxford University Press.
- Locke, D. (2015). Practical certainty. *Philosophy and Phenomenological Research*, 90(1):72–95.
- Locke, J. and Nidditch, P. H. (1979). *The Clarendon Edition of the Works of John Locke: An Essay Concerning Human Understanding*. Clarendon Press.
- Lord, E. (2010). Having reasons and the factoring account. *Philosophical Studies*, 149(3):283–296.
- Marusic, B. (2011). The ethics of belief. *Philosophy Compass*, 6(1):33–43.
- McNally, L. (2011). The relative role of property type and scale structure in explaining the behavior of gradable adjectives. In Nouwen, R., van Rooij, R., and Sauerland, U., editors, *Vagueness in Communication*. Springer.
- Miller, R. W. (1978). Absolute certainty. *Mind*, 87(345):46–65.
- Moore, G. E. (1942). A reply to my critics. In Schlipp, P., editor, *The philosophy of G. E. Moore*, pages 543–667. Evanston: Tudor.
- Moore, G. E. (1959). Certainty. In *Philosophical Papers*. Allen & Unwin/ Unwin Hyman.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86 (2):279–294.
- Nagel, J. (2010). Epistemic anxiety and adaptive invariantism. *Philosophical Perspectives*, 24(1):407–435.

- Nagel, J. (2012). The attitude of knowledge. *Philosophy and Phenomenological Research*, 84(3):678–685.
- Neta, R. (2009). Treating something as a reason for action. *Noûs*, 43 (4):684–699.
- Owens, D. (2000). *Reason Without Freedom: The Problem of Epistemic Normativity*. Routledge.
- Pagin, P. (2015). Problems with norms of assertion. *Philosophy and Phenomenological Research*, 92(1).
- Parfit, D. (2001). Rationality and reasons. In Egonsson, D., Jesefsson, J., Petersson, B., and Rønnow-Rasmussen, T., editors, *Exploring Practical Philosophy*. Burlington, VT: Ashgate.
- Parfit, D. (2011). *On What Matters*. Oxford University Press.
- Pascal, B. (1670). *Pensées*. London: Dent, 1910.
- Perry, J. (1986). Thought without representation. *Proceedings of the Aristotelian Society Supplementary*, 60:262–283.
- Pidgen, C. (2013). The is-ought gap. In LaFollette, H., editor, *The International Encyclopedia of Ethics*. Wiley-Blackwell.
- Pollock, J. L. and Cruz, J. (1999). *Contemporary Theories of Knowledge*. Rowman & Littlefield.
- Pritchard, D. (2008). Certainty and scepticism. *Philosophical Issues*, 18(1):58–67.
- Radford, C. (1966). Knowledge—by examples. *Analysis*, 27(1):1–11.
- Raz, J. (1999). *Practical Reason and Norms*. Oxford University Press.
- Reed, B. (2010). A defense of stable invariantism. *Noûs*, 44 (2):224–244.
- Reed, B. (2014). Practical matters do not affect whether you know. In Steup, M. and Turri, J., editors, *Contemporary Debates in Epistemology*, pages 95–106. Blackwell.
- Ross, J. and Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*, 88(2):259–288.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35 (4):477–514.
- Rysiew, P. (2005). Contesting contextualism. *Grazer Philosophische Studien*, 69 (1):51–70.
- Rysiew, P. (2007). Speaking of knowing. *Noûs*, 41(4):627–662.
- Scanlon, T. . (2014). *Being Realistic About Reasons*. Oxford University Press.

- Scanlon, T. M. (1998). *What We Owe to Each Other*. Harvard University Press.
- Schaffer, J. (2006). The irrelevance of the subject: Against subject-sensitive invariantism. *Philosophical Studies*, 127 (1):87–107.
- Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, 96:317–333.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press.
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139(1):57–71.
- Setiya, K. (2007). *Reasons Without Rationalism*. Princeton University Press.
- Shah, N. (2003). How truth governs belief. *Philosophical Review*, 112(4):447–482.
- Shah, N. and Velleman, J. D. (2005). Doxastic deliberation. *Philosophical Review*, 114(4):497–534.
- Skorupski, J. (2010). *The Domain of Reasons*. Oxford University Press.
- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford University Press.
- Smithies, D. (2012). The normative role of knowledge. *Noûs*, 46(2):265–288.
- Sosa, E. (2000). Skepticism and contextualism. *Philosophical Issues*, 10(1):1–18.
- Srinivasan, A. (2015). Are we luminous? *Philosophy and Phenomenological Research*, 90(2):294–319.
- Sripada, C. S. and Stanley, J. (2012). Empirical tests of interest-relative invariantism. *Episteme*, 9(1):3–26.
- Stanley, J. (2005a). Fallibilism and concessive knowledge attributions. *Analysis*, pages 126–131.
- Stanley, J. (2005b). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18(1):35–57.
- Sturgeon, S. (2008). Reason and the grain of belief. *Nous*, 42(1):139–165.
- Sutton, J. (2007). *Without Justification*. MIT Press.
- Sylvan, K. (2016). Epistemic reasons i: Normativity. *Philosophy Compass*, 11(7):364–376.
- Tappolet, C. (2014). The normativity of evaluative concepts. In *Mind, Values and Metaphysics. Philosophical Essays in Honour of Kevin Mulligan*, volume 2. Reboul, A.

- Thomson, J. J. (2008). *Normativity*. Open Court.
- Turri, J. (2010). Prompting challenges. *Analysis*, 70(3):456–462.
- Unger, P. K. (1975). *Ignorance: A Case for Scepticism*. Oxford University Press.
- Velleman, D. (2000). *The Possibility of Practical Reason*. Oxford University Press.
- von Wright, G.-H. (1963). *Norm and Action: A Logical Enquiry*. Routledge and Kegan Paul.
- Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 10 (1):417–443.
- Wedgwood, P. (2002). Internalism explained. *Philosophy and Phenomenological Research*, 65(2):349–369.
- Wedgwood, R. (2012). Outright belief. *Dialectica*, 66(3):309–329.
- Weiner, M. (2005). Must we know what we say? *Philosophical Review*, 114(2):227–251.
- Williams, B. (1973). Deciding to believe. In *Problems of the Self*, pages 136–151. Cambridge University Press.
- Williams, B. (1980). *Moral Luck*, chapter Internal and External Reasons. Cambridge University Press, 1981 edition.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55 (219):213–235.
- Williamson, T. (2009). Reply to mark kaplan. In Pritchard, D. and Greenough, P., editors, *Williamson on Knowledge*, pages 333–340. Oxford: Oxford University Press.
- Williamson, T. (forthcoming). Justifications, excuses, and sceptical scenarios. In Dorsch, F. and Dutant, J., editors, *The new evil demon problem*. Oxford University Press.