



Belief, Knowledge and Practical Matters

高洁 著
Written by Jie Gao



ZHEJIANG UNIVERSITY PRESS
浙江大学出版社

This work is supported by the National
Social Science Fund of China
(No. 21CZX016).

Belief, Knowledge and Practical Matters

高 洁 著

Written by Jie Gao



ZHEJIANG UNIVERSITY PRESS

浙江大学出版社

• 杭州 •

图书在版编目 (CIP) 数据

信念、知识和实践问题 = Belief, Knowledge and Practical Matters : 英文 / 高洁著. -- 杭州 : 浙江大学出版社, 2024.2
ISBN 978-7-308-24679-8

I. ①信… II. ①高… III. ①知识论—英文 IV.
①G302

中国国家版本馆CIP数据核字(2024)第016720号

信念、知识和实践问题

Belief, Knowledge and Practical Matters

高 洁 著

Written by Jie Gao

责任编辑 周烨楠

责任校对 李瑞雪

责任印制 范洪法

封面设计 项梦怡

出版发行 浙江大学出版社

(杭州市天目山路 148 号 邮政编码 310007)

(网址: <http://www.zjupress.com>)

排 版 杭州朝曦图文设计有限公司

印 刷 广东虎彩云印刷有限公司绍兴分公司

开 本 880mm × 1230mm 1/32

印 张 8

字 数 389 千

版 印 次 2024 年 2 月第 1 版 2024 年 2 月第 1 次印刷

书 号 ISBN 978-7-308-24679-8

定 价 88.00 元

版权所有 翻印必究 印装差错 负责调换

浙江大学出版社市场运营中心联系方式: 0571-88925591; <http://zjdxcs.tmall.com>

Preface

For much of the twentieth century, most epistemologists held views according to which the epistemic realm is independent of the practical realm, and epistemic concepts are independent from practical ones. This ‘purist’ orthodoxy has been challenged since the beginning of the twenty-first century. According to pragmatic encroachment, whether an epistemic attitude towards p has some positive epistemic status (e.g., whether a belief is epistemically rational or justified, or it amounts to knowledge) partially depends on practical factors such as the costs of being wrong or the practical goals of the subject. Depending on such factors, a belief may count as justified or as knowledge in some circumstances but not in others. Among the many varieties of pragmatic encroachment, encroachment on knowledge is one of the most important and controversial.

This book takes purism about knowledge as the default position and defends it from the challenges of pragmatic encroachment. The book is divided into two parts, a negative and a positive one. The negative part critically examines existing purist strategies in response to pragmatic encroachment. The positive part provides a new theory of how practical factors can systematically influence our confidence and explores some implications of such influence. In particular, it provides a new purist explanation of the data commonly used to motivate pragmatic encroachment on knowledge. Moreover, it develops a new variety of pragmatic encroachment, not on knowledge but on credence. Thus, my aim in this book is not to provide a full-fledged defence of all kinds of purism. Rather, I argue for purism about knowledge but pragmatic encroachment about credence.

This book is a revised version of my PhD dissertation “Belief, Knowledge and Action” which was passed with no correction in August 2016 at the University of Edinburgh. In this revision, I have made some substantive changes. Chapter 1 includes a new section (section 1.4) on the problematic consequences of pragmatic encroachment on knowledge.

Most parts of chapter 2 have been revised. They draw heavily on an article I published in *Episteme*, “Against the iterated knowledge account of high-stakes cases” (2019). Chapter 5 in the original dissertation has now been split into two chapters (5 and 6). In particular, the current chapter 5 draws substantially on my 2019 article in *Philosophical Studies*, “Credal pragmatism”. In addition, throughout the book I have introduced some stylistic changes and added discussions of new relevant literature.

I would like to thank a number of people who have provided useful feedback on the original dissertation and on earlier drafts. Mikkel Gerken, my main doctoral supervisor in Edinburgh, provided detailed feedback on numerous drafts of all chapters in the dissertation. Duncan Pritchard and Allan Hazlett, my secondary supervisors, provided helpful feedback on different draft chapters of the dissertation. I also benefited from conversations with Jennifer Nagel and Nick Treanor, the two referees of my PhD viva. Other people who have commented on earlier drafts include Aidan McGlynn, Jacques Vollet and Roger Clarke. Special thanks go to Davide Fassio, who provided detailed comments on the draft of the whole book.

Finally, thanks to the editors of *Episteme*, *Synthese* (chapter 3 draws substantially on my 2017 article “Rational action without knowledge (and vice versa)”), and *Philosophical Studies* for permission to reuse in this book some material from articles mentioned above. I would also like to thank Yenan Zhou at Zhejiang University Press for her efficiency, good advice, and support.

This work is supported by the National Social Science Fund of China (No. 21CZX 016)

Contents

Introduction 1

PART I Pars Destruens

1. Knowledge and Practical Matters 9

- 1.1 Practical factor effects on knowledge ascriptions 10
- 1.2 The knowledge norm of practical reasoning 16
- 1.3 Pragmatic encroachment on knowledge 21
- 1.4 Problems for pragmatic encroachment 28
- 1.5 Moderate invariantist accounts of the practical factor effects on knowledge ascriptions 31
- 1.6 Criticisms to the knowledge norm of practical reasoning 38
- 1.7 Concluding remarks 42

2. High Stakes and Iterated-knowledge 44

- 2.1 The iterated knowledge account of high-stakes cases 46
- 2.2 Counterintuitive appropriateness 53
- 2.3 Higher-order reasoning/knowledge and practical rationality 58
- 2.4 Non-luminosity, higher-order belief and blameworthiness 67
- 2.5 Concluding remarks 71

2 Belief, Knowledge and Practical Matters

3. Rational Action Without Knowledge (and Vice Versa) 73

- 3.1 The epistemic norms of practical reasoning 73
- 3.2 Counterexample one: rational action based on acceptance 78
- 3.3 Counterexample two: rational action performed by the
sceptic 88
- 3.4 The excuse manoeuvre 92
- 3.5 Concluding remarks 93

4. Doxastic Accounts 94

- 4.1 Weatherson's doxastic pragmatism 95
- 4.2 Ganson's doxastic pragmatism 99
- 4.3 Objections to Weatherson and Ganson's accounts 102
- 4.4 Bach's doxastic pragmatism 107
- 4.5 Nagel's psychological account 109
- 4.6 Objections to Bach and Nagel's accounts 115
- 4.7 Concluding remarks 118

PART II Pars Construens

5. Credal Pragmatism 121

- 5.1 Two types of practical factors 123
- 5.2 Threshold pragmatism and credal pragmatism 125
- 5.3 Credal pragmatism vs. threshold pragmatism 131
- 5.4 Objections and replies 139
- 5.5 Three dimensions of epistemic rationality 144
- 5.6 Concluding remarks 155

6. Credal Pragmatism and the Distinction Between Dispositional Belief and Occurrent Belief 157

6.1 Dispositional belief and occurrent belief 157

6.2 Dispositional belief, occurrent belief and practical factor effects
164

6.3 Diagnoses of high-stakes cases 173

6.4 Concluding remarks 180

7. Fallibilism and the Knowledge Norm of Practical Reasoning 181

7.1 Fallibilism and the knowledge norm of practical reasoning 182

7.2 Threshold makers and communicative heuristic accounts 195

7.3 The infallibilist intuition and the knowledge norm of practical
reasoning 199

7.4 Concluding remarks 215

Concluding Remarks 217

Bibliography 219

Index 243

Introduction

According to a traditional view in epistemology, *epistemological purism*, whether a true belief qualifies as knowledge exclusively depends on truth-relevant factors, such as the quantity and quality of evidence, the reliability of belief-forming methods, the counterfactual relations to the truth of the believed proposition, and so on. In this view, if my true belief is that p qualifies as knowledge while yours is not, this must be because of some difference in some of truth-relevant factors. Epistemological purism has been by far the dominant view in the history of epistemology. However, in recent years this consensus has been broken. In the last decade, debates over the relation between on the one hand knowledge and other epistemic conditions such as justification and on the other hand practical matters have become one of the central focuses in contemporary epistemology. By ‘practical matters’ I mean practical factors, such as stakes, and practical rationality.¹

This ‘practical turn’ in epistemology develops around two controversial claims. One is that knowledge depends on practical factors as well as truth-relevant factors (Fantl & McGrath, 2002, 2007, 2009a, 2009b, 2009c, 2012; Hawthorne, 2004; Ross & Schroeder, 2014; Schroeder, 2012; Stanley, 2005; Weatherson, 2011, 2012). The other is that practical reasoning is constrained by a knowledge norm which says that knowledge that p is a necessary and/or sufficient condition for appropriately treating p as a premise in practical reasoning (Fantl & McGrath, 2002, 2007, 2009a,

¹ Following the common use in the literature, I use ‘practical factors’ in a restricted descriptive sense (as opposed to a normative sense). Much of the literature focuses on the costs of being wrong, sometimes referred to as ‘the stakes’. However, Anderson (2015) and Gerken (2011, 2017) argue that this myopic focus on stakes is a mistake. According to them, the relevant practical factors should also include, amongst other things, the costs of double-checking, the urgency of acting, the availability of alternative courses of action, the availability of further evidence, social roles and conventions associated with the action. I agree with their diagnosis.

2 Belief, Knowledge and Practical Matters

2009b, 2012; Hawthorne, 2004; Hawthorne and Stanley, 2008; Mueller, 2021; Ross & Schroeder, 2014; Stanley, 2005; Williamson, 2005).²

Against this background, the aim of this book is to explore two clusters of issues concerning the relations between knowledge, belief, practical factors and practical rationality. One cluster of issues concerns the effects of practical factors on epistemic states. In particular, do practical factors have any non-trivial effects on doxastic attitudes and knowledge?³ If yes, how? What are the implications for our understanding of epistemic rationality? The other cluster concerns the debate on the epistemic norm of practical rationality: is practical rationality governed by any epistemic norm, i.e., does appropriately treating p as a premise in practical reasoning require a satisfaction of some epistemic condition with respect to p ? If yes, what is the epistemic condition at issue?

More precisely, in this book I defend a view called *credal pragmatism*, according to which credence is sensitive to practical factors, occurrent belief depends on the actual degree of credence, and dispositional belief depends on the degree of credence one would have in normal circumstances. Credal pragmatism explains two aspects of the intuitive relation between knowledge and practical matters. The first aspect is constituted by the so-called *practical factor effects on knowledge ascriptions*, according to which the practical factors associated with a proposition, p , seem to influence our ascriptions of knowledge that p or our assessments of these ascriptions. The second aspect concerns some central motivations for the idea that knowledge is the epistemic norm governing practical

² Some philosophers take the second claim to explain the first one (e.g., Fantl & McGrath, 2002, 2009a). Others take the two claims to be independent. For example, Stanley (2005) presents the two claims as separate and independent. Hawthorne and Stanley (2008, p. 576, p. 588) argue, against Fantl and McGrath, for the separateness of the two claims. Weatherson (2012) argues for the former claim, and holds that it doesn't need any support from the latter. Other philosophers endorse the latter claim, but not the former (see e.g., Williamson (2005)).

³ There are obvious 'trivial' influences of practical factor effects on doxastic attitudes. They concern, for example, cases in which the attitudes are about practical matters, and thus a change in practical matters affects one's doxastic attitudes. I am not concerned with this type of effects here. My main concern is with effects of practical factors on the attitudes without affecting directly their contents.

reasoning. One main motivation is the prominence of the use of ‘knows’ and its cognates in folk epistemic assessments of practical reasoning—and in particular the fact that knowledge is often taken to be a necessary and sufficient epistemic condition for relying on a proposition in practical reasoning, at least in most circumstances. Credal pragmatism explains these two aspects in terms of specific relations between doxastic attitudes and practical matters. More precisely, according to this view, both aspects can be explained by appeal to specific dispositional properties constitutive of occurrent belief.

Credal pragmatism is compatible with *moderate purist invariantism*, which is the received view in contemporary epistemology. As said above, *purism* (i.e., epistemological purism) is the view that only truth-relevant factors matter for knowledge. Here I take purism to be a metaphysical claim concerning the nature of knowledge.⁴ *Invariantism* holds that the semantic content (truth-conditions and truth-values) of knowledge ascriptions does not vary with changes in the context of ascription or assessment.⁵ ‘*Moderate*’ denotes a non-sceptical position according to which a suitably wide range of ordinary knowledge-claims is true. Moderate purist invariantism (henceforth, moderate invariantism) is a conjunction of purism, invariantism and non-sceptical moderatism. In sum, this view takes knowledge to depend exclusively on truth-relevant factors and to require a contextually invariant epistemic standard that we can meet quite easily and very often do, and takes knowledge ascriptions to univocally refer to knowledge so conceived.

The book is constituted by two parts, a negative (*Pars destruens*) and a positive one (*Pars construens*). The negative part (chapters 1–4) clears the path for credal pragmatism by challenging the knowledge norm of practical reasoning and by advancing a range of problems affecting

⁴ Gerken (2017) distinguishes *semantic purism*, according to which ‘knows’ is not semantically sensitive to practical factors and *metaphysical purism*, according to which knowledge itself is not sensitive to practical factors. Here, following Fantl and McGrath (2009a), I use ‘purism’ only to refer to the latter kind of purism identified by Gerken. Note that within the taxonomy of Gerken, the conjunction of semantic purism and metaphysical purism equals to purist invariantism.

⁵ See MacFarlane (2014) for discussions of these technical terms.

4 Belief, Knowledge and Practical Matters

some prominent moderate invariantist approaches alternative to credal pragmatism.

The main goal of chapter 1 is to provide a general overview of the issues that constitute the background for the views and arguments defended in this book. In particular, I provide a thorough discussion of two aspects of the relation between knowledge and practical matters: one is constituted by the practical factor effects on knowledge ascriptions; the other is the intuitive normative role of knowledge in the regulation and assessments of action and practical reasoning. Given the focus of this book—a defence of a specific moderate invariantist view—I will tackle these issues from a moderate invariantist perspective.

Chapter 2 critically examines Williamson (2005)’s account of the practical factor effects on knowledge ascriptions. Timothy Williamson endorses both the knowledge norm of practical reasoning and moderate invariantism. He must explain away the intuition according to which it seems appropriate to rely on a proposition in practical reasoning when stakes are low, but not when stakes are high. His account features an error theory about the intuitive judgments about high-stakes cases. According to Williamson’s account, an alleged failure to acknowledge the distinction between knowing and knowing that one knows can explain the intuitive judgments in question. In this chapter, I provide three objections to Williamson’s account. This result undermines the prospects of defending moderate invariantism while maintaining a knowledge norm of practical reasoning. A better strategy is to refute the knowledge norm of practical reasoning.

Chapter 3 criticises the idea that practical reasoning is governed by an epistemic norm. I provide original counterexamples to epistemic norms in general. These counterexamples are based on cases in which it is intuitively appropriate for the subject to rely on propositions that the subject doesn’t believe.

Chapter 4 criticises some so-called doxastic accounts. These accounts explain the practical factor effects on knowledge ascriptions in terms of the influence of practical matters on belief. In particular, I consider the accounts of Weatherston (2005), Ganson (2008), Bach (2005, 2008, 2010) and Nagel (2008, 2010a). Though I criticise these accounts, the positive view that I defend in this book can be classified as an alternative type of

doxastic account.

The second, positive part of the book (chapters 5–7) defends credal pragmatism. Moreover, on the basis of this view, it proposes novel accounts of the intuitive relations between knowledge and practical matters. Chapter 5 compares credal pragmatism with threshold pragmatism and argues that the former is more plausible than the latter. This chapter also considers the issue of whether our doxastic attitudes' sensitivity to practical factors can be considered rational, and if yes, in what sense.

Chapter 6 explores some implications of the sensitivity of credence to practical factors for the distinction between two varieties of full belief: dispositional belief (which is the type of belief relevant for knowledge) and occurrent belief. On the basis of this distinction, I provide an account of the practical factor effects on knowledge ascription.

In Chapter 7, based on credal pragmatism, I develop a novel account of the data used to motivate the knowledge norm of practical reasoning. These data concern several features about knowledge ascriptions, including i) why in folk epistemological practices knowledge is often taken to be a necessary and sufficient epistemic condition for relying on a proposition in practical reasoning; ii) concessive knowledge attributions; and iii) the infallibilist intuition that knowledge excludes error possibilities. I show that the account of the above phenomena is coherent with the fallibilist picture of knowledge.

PART I

Pars Destruens

Knowledge and Practical Matters

This chapter provides a general overview of a set of issues concerning the intuitive relation between knowledge and practical matters, which constitute the background for the views and arguments defended in this book. In particular, I provide a thorough discussion of two aspects of that relation, mentioned in the Introduction: one is constituted by practical factor effects on knowledge ascriptions; the other is the intuitive normative role of knowledge in the regulation and assessments of action and practical reasoning. Given the focus of this book—the defence of a specific moderate invariantist view, I will tackle these issues from a moderate invariantist perspective.

This is the plan of the chapter. Sections 1.1 and 1.2 present two aspects of the relation between knowledge and practical matters. Section 1.1 introduces the intuitive data suggesting the existence of practical factor effects on knowledge ascriptions. Section 1.2 discusses the role of knowledge ascriptions in ordinary epistemic assessments of practical reasoning and the idea that knowledge is the norm of practical reasoning. Section 1.3 presents and critically discusses two main arguments against purist moderate invariantism based on these aspects of the relation between knowledge and practical matters. Section 1.4 presents problematic consequences of pragmatic encroachment. Section 1.5 introduces the two main moderate invariantist accounts of the practical factor effects on knowledge ascriptions: doxastic accounts and pragmatic accounts. This section will focus in particular on the latter type of accounts the former will be carefully examined in chapter 4. Pragmatic accounts will receive relatively little attention in the rest of the book, but they are sufficiently important and discussed in the literature to deserve acknowledgment and critical discussion here. Section 1.6 considers some popular objections provided by moderate invariantists to the knowledge norm of practical

reasoning. Section 1.7 sums up the main upshots of my discussion in this chapter.

1.1 Practical factor effects on knowledge ascriptions

Our intuitive judgments about certain pairs of cases seem to suggest practical factor effects on knowledge ascriptions. In particular, such cases highlight shifting patterns of knowledge ascriptions due to varying practical factors concerning the subject's and the speaker's practical interests.¹ Consider the following pair of cases:

BANK

Low Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says, "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning."

High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, "I guess you're right.

¹ Gerken (2017) defines a broader notion of practical factor effects on knowledge ascription in which practical factors concern the subject's, the speaker's, the hearer's, or the evaluator's practical interests.

I don't know that the bank will be open tomorrow." (Stanley, 2005, pp. 3–4, adapted from DeRose, 1992, p. 913)

This pair of cases is designed in such a way that Hannah in *Low Stakes* (henceforth the LS-subject) and Hannah in *High Stakes* (henceforth the HS-subject) share the same strength of epistemic position with regard to the proposition that the bank will be open tomorrow (henceforth q) and the subjects in both cases believe that q .² The two cases vary along two factors. One factor is the stakes associated with q for the subject and the hearer. The other factor is the conversational salience of alternatives associated with q . It is only in *High Stakes* that the stakes are high and an alternative, that the bank might change hours, becomes conversationally salient. Intuitively, the self-knowledge ascription made by the LS-subject is true, while the self-knowledge denial made by the HS-subject is true as well.

In the two bank cases, the knowledge ascriber is identical to the putative-knower; therefore, the stakes for the ascriber and the putative-knower are the same person. Let's consider two cases in which the knowledge ascriber is different from the putative-knower:

Low Attributor–High Subject Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their

² The term 'strength of epistemic position' originates from Reed (2010). It can be understood as a placeholder to be filled by one's preferred account of justification, warrant, etc. It is worth noting that some pragmatic encroachers such as Stanley would disagree with the claim that subjects in low- and high-stakes cases share the same strength of epistemic position. This is because these philosophers take one's strength of epistemic position as dependent on one's evidential support, and evidence as equal to knowledge. Since these philosophers take knowledge to be dependent on practical factors, also evidence, epistemic support and strength of epistemic position are so dependent on these factors. Here I will simply assume a sense of strength of epistemic position which is independent of practical factors. In the bank case, this strength is constituted by the support to the target belief provided by Hannah's memory, which is exactly the same in the two circumstances.

account, it is very important that they deposit their paychecks by Saturday. Two weeks earlier, on a Saturday, Hannah went to the bank, where Jill saw her. Sarah points out to Hannah that banks do change their hours. Hannah utters, “That’s a good point. I guess I don’t really know that the bank will be open on Saturday.” Coincidentally, Jill is thinking of going to the bank on Saturday, just for fun, to see if she meets Hannah there. Nothing is at stake for Jill, and she knows nothing of Hannah’s situation. Wondering whether Hannah will be there, Jill utters to a friend, “Well, Hannah was at the bank two weeks ago on a Saturday. So she knows the bank will be open on Saturday.”

High Attributor–Low Subject Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah calls up Bill on her cell phone, and asks Bill whether the bank will be open on Saturday. Bill replies by telling Hannah, “Well, I was there two weeks ago on a Saturday, and it was open.” After reporting the discussion to Sarah, Hannah concludes that, since banks do occasionally change their hours, “Bill doesn’t really know that the bank will be open on Saturday.” (Stanley, 2005, pp. 4–5)

In *Low Attributor–High Subject Stakes*, the intuition recognised by most philosophers is that Jill’s third-person positive knowledge ascription is false because the subject Jill attributes knowledge to is in a high-stakes situation. In *High Attributor–Low Subject Stakes*, the suggested intuitive reaction is that Hannah’s negative third-person knowledge ascription is true, even though the subject to whom Hannah attributes is in a low-stakes situation.³ This seems to suggest that both the attributor and the subject’s high stakes can influence our intuitive judgment about third-

³ As I will discuss below, empirical studies provide inconclusive results on whether folks hold the claimed intuitions with respect to the cases at issue.

person knowledge ascription.

Other similar examples involving third-person knowledge ascriptions (*Low Attributor-Low Subject* and *High Attributor-Low Subject*) discussed in the literature include Stewart Cohen's *airport* cases (Cohen, 1999, p. 58), Jeremy Fantl and Matthew McGrath's *train* cases (Fantl & McGrath, 2002, p. 67, 2009a, p. 32, fn.2) and Gerken's *boat* cases (Gerken, 2017, chapter 2).⁴ Given our reactions to the above cases in which the attributor is distinct from the putative-knower, it seems that the stakes for the attributor and the subject both matter to our judgments about knowledge ascriptions.

Intuitive judgments about the pairs of cases introduced above are considered non-controversial by the wide majority of philosophers in the debate (e.g., Bach, 2005, 2008, 2010; Brown, 2013; S. Cohen, 1999; DeRose, 2009; Fantl & McGrath, 2002, 2009a; Gerken, 2017; Nagel, 2008, 2010a; Stanley, 2005).⁵ Other cases are more controversial yielding clashes of intuitions among philosophers. One type of controversial case, so-called *Ignorant High Stakes*, separates the *de facto* stakes from the reasonably presupposed stakes:

Ignorant High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. But neither Hannah nor Sarah is aware of the impending bill, nor of the paucity of available funds. Looking at the lines, Hannah says to Sarah, "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning."

The peculiar feature of *Ignorant High Stakes* is that the subject is not aware of the *de facto* high stakes. Stanley takes it to be intuitive that in such cases

⁴ In Gerken's boat cases, there is no variation in the conversational salience of alternatives. Only a variation in the stakes of the ascriber/speaker is left.

⁵ However, as I will discuss below, empirical studies provide inconclusive results on whether folks hold the claimed intuitions with respect to the cases at issue.

the subject's self-ascription of knowledge is false—the subject doesn't know (Stanley, 2005, p. 5). However, other philosophers don't have the same intuitive reaction (see e.g., Gerken, 2017, p. 36).

Another type of stakes cases that generated similar interpretational controversies is constituted by *non-linguistic* cases, in which the subject does not make any linguistic knowledge ascription. Some non-linguistic cases involve an inner mental knowledge ascription instead of a conversational knowledge ascription. For example, Fantl and McGrath consider a non-conversational version of *High Stakes* where the subject entertains the following internal monologue:

Keith will in fact be willing to say this to himself all alone:
 “Gosh, if we wait till Saturday and the bank is closed, we will
 be in deep trouble. Do I know the bank will be open tomorrow?
 No, I guess I don't. I should go check inside.” (Fantl & McGrath,
 2009a, p. 49)

Some other non-linguistic pair of cases even don't include a mental knowledge ascription. They typically involve a low- and a high-stakes cases in which neither linguistic knowledge ascription nor non-linguistic mental knowledge ascription is presented. Examples include Brian Weatherson's *genie* case (see Weatherson, 2012, pp. 82–83), Jacob Ross and Mark Schroeder's *sandwich* case (see Ross & Schroeder, 2014, p. 261) and Chandra Sripada and Jason Stanley's *pine nuts* case (see Sripada & Stanley, 2012, pp. 11–12).

Until now I have been talking about intuitive judgments of philosophers. A number of empirical studies in experimental philosophy have been conducted concerning intuitive reactions of folks to stakes pair cases. Evidence both supporting and undermining the existence of practical factor effects on knowledge ascriptions has been reported. These experiments have also provoked debates over methodological issues. Studies including Buckwalter (2010), May et al. (2010), Feltz and Zarpentine (2010), Knobe and Schaffer (2012), Rose et al. (2019), Francis et al. (2019) report null results of the stakes effects on knowledge ascription. Some of these studies are criticised by DeRose (2011), Sripada and Stanley (2012) and Pinillos and Simpson (2014) with regard to methodological

issues and statistical power of null results (see Buckwalter 2014 for a response to critics). Other studies, including Pinillos (2012), Sripada and Stanley (2012), Pinillos and Simpson (2014), Dinges and Zakkou (2020) provide positive evidence for the stakes effects on knowledge ascriptions. In particular, those studies provide positive evidence for alleged intuitions about non-linguistic cases and ignorant cases. In addition, there has been a debate between Buckwalter and Schaffer (2015) and Pinillos and Simpson (2014) over methodological issues about the experiments in Pinillos (2012) and Sripada and Stanley (2012). Some other studies in experimental philosophy, including Buckwalter and Turri (2017), Pinillos (2012), Pinillos and Simpson (2014) and Turri, Buckwalter and Rose (2016), indicate an important correlation between knowledge and actionability (i.e., how a person should act), supporting the idea that practical factor effects in stakes pair cases strongly depend on the practical factors concerning action. Besides practical factor effects on knowledge ascriptions generated by stakes, Shin (2014) has found that similar effects can be driven by urgency. According to the results of Shin's experiments, participants are more willing to ascribe knowledge to a subject in an urgent condition than a counterpart subject under a non-urgent condition.⁶

Although neither in 'armchair' philosophy nor in experimental philosophy a consensus has been reached over whether there is any practical effect on knowledge ascriptions, we should not neglect alleged intuitive asymmetries and positive evidence of the effects. Hence, in the rest of our discussion, I will assume that the data considered by philosophers constitute a *prima facie* evidence for the existence of practical factor effects on knowledge ascriptions.

By acknowledging the practical factor effects on knowledge ascriptions, a challenge emerges for moderate invariantists. Given that in the relevant stakes-involving cases it is stipulated that the subject's belief condition and epistemic factors relevant to that belief remain constant across low- and high-stakes cases, moderate invariantism delivers the counterintuitive verdict that the subjects in both low- and high- stakes cases know that *p*. The challenge for moderate invariantists, then, is to explain why it seems

⁶ See Gao (2015) for further discussion of recent experimental philosophy studies in epistemology, with a focus on studies concerning pragmatic encroachment.

infelicitous to ascribe knowledge that q to the putative-knowers in cases such as *Highs Stakes*, *Ignorant High Stakes*, *Low Attributor-High Subject Stakes* and *High Attributor-Low Subject Stakes*, and why it seems felicitous to deny knowledge that q to the relevant putative-knowers even though they know that q .

1.2 The knowledge norm of practical reasoning

Our ordinary epistemic assessments reveal intuitive connections between knowledge and practical reasoning. On that basis, some philosophers have argued that knowledge plays an important normative role in practical reasoning (Fantl & McGrath, 2002, 2009a; Hawthorne, 2004, pp. 21–32, chapter 4; Hawthorne and Stanley, 2008; Stanley, 2005).⁷ First, it has been highlighted that we often invoke knowledge in criticising others for acting on inadequate grounds. In one example due to Hawthorne and Stanley (2008), if one doesn't buy health insurance on the grounds that he is healthy enough, his loved ones can criticise him since he doesn't know that he will not get sick. Likewise, if on the way to a restaurant, I just go down a street on a hunch and then discover that the direction is wrong, it would be very natural for my partner to complain, "You shouldn't have gone down this street, since you did not know that the restaurant was there." To take another example, before the result of a lottery is announced, it seems inappropriate for me to sell my lottery ticket for a penny on the basis that I will lose, since I do not know that the ticket is a loser. These kinds of considerations support the following norm given by Hawthorne and Stanley (2008):

Action-Knowledge Principle (AKP)

Treat the proposition that p as a reason for acting only if you know that p .

⁷ Besides practical reasoning, it has also been argued that knowledge sets the normative standard for appropriate or proper belief and assertion. The endorsement of knowledge norms is often taken to be one of the core commitments of *knowledge-first epistemology*. See Williamson (2000, 2011, 2013) for a defence of knowledge-first epistemology and Gerken (2018) and McGlynn (2013, 2014) for criticisms.

This norm articulates a necessary condition for appropriately treating a proposition as a reason for acting. The norm neatly explains our negative epistemic assessments of actions identified above. In addition, it fits well with the suggestion that, in a decision-theoretic framework, only one's beliefs which amount to knowledge should shape one's decision tables (cf. Dutant, forthcoming; Schulz, 2017; Weatherson, 2012).

Second, we often appeal to knowledge to justify our decisions to do, or refrain from doing, certain things. For one thing, we often defend our past actions by saying “but I knew that p ” or our ongoing actions by saying “but I know that p ”. For example, when I am asked why I went down that direction rather than the other, I reply that I knew that it was the shortest direction to the restaurant (Stanley, 2005; Hawthorne & Stanley, 2008). For another thing, we can rationalise our decisions or recommend certain actions by ascribing the relevant knowledge to the subject. As Lackey (2010, p. 363) notes, once one comes to know that one's lottery ticket is a loser—one has learnt the result through radio announcement—then one can go ahead and tear up one's ticket or throw it away. These kinds of considerations suggest that knowledge provides sufficient epistemic grounds for rational action.

Fantl and McGrath have defended a variety of sufficiency conditions tying knowledge to action (labelled ‘Action’ in 2009a, p. 49, ‘KJ’ in 2009a, p. 66 and ‘Actionability’ in 2012, p. 65):

Action

If you know that p , you are proper to act on p when the question of whether p is relevant to the question of what to do.

KJ

If you know that p , then p is warranted enough to justify you in ϕ -ing, for any ϕ .

Actionability

You can know that p only if p is actionable for you.⁸

⁸ Fantl and McGrath (2012, pp. 65–66) interpret ‘ p to be actionable for you’ as the idea that epistemic shortcomings in your relationship to p do not stand in the way of your

Ross and Schroeder (2014) formulate a sufficiency condition in the following way:

Knowledge-Action Principle (KAP)

For any agent S and proposition p , if S is in a choice situation in which S could not rationally act as if p , then S does not know that p (where to act as if p is to act in the manner that would be rationally optimal on the supposition that p is true).

Some philosophers defend biconditional principles involving both a necessity and a sufficient condition. The best-known version has been proposed by Hawthorne and Stanley (2008, p. 578):

Reason-Knowledge Principle (RKP)

Where one's choice is p -dependent, it is appropriate to treat the proposition that p as a reason for acting *iff* you know that p .

Three remarks about RKP are in order here. First, according to Hawthorne and Stanley, a choice between options x_1, \dots, x_n is p -dependent *iff* the most preferable of x_1, \dots, x_n conditional on the proposition that p is not the same as the most preferable of x_1, \dots, x_n conditional on the proposition that not- p (ibid.). Such a condition is needed for there are many cases where p is completely irrelevant to the issue at hand, so it seems odd to say that it is appropriate to treat the proposition that p as a reason for acting, even if one knows that p .

Second, the notion of appropriateness in RKP is supposed to be understood as rational permissibility, rather than in terms of obligation. As Hawthorne and Stanley (2008) point out, "it would be overly

reliance on p as a basis for action. They also suggest a more precise characterization of actionability in terms of epistemic certainty: p is actionable for you *iff* either p is epistemically certain for you or your lacking epistemic certainty for p does not stand in the way of p 's being among your justifying practical reasons. Fantl and McGrath (2009) provide a slightly different explanation of actionability in terms of justifying practical reasons: A justifying practical reason is a practical reason that doesn't merely support doing a given action; it supports it strongly enough so that the action is justified for you.

demanding to require someone to treat all of their relevant knowledge as reasons for each action undertaken.”

Third, there are two readings of ‘treating p as a reason for action’ but only one of them is compatible with the intuition that RKP is supposed to capture. One is to take ‘treating p as a reason for action’ as ‘using p as a premise in practical reasoning’. This interpretation is in accordance with the formulation endorsed by Hawthorne (2004) and is the most intuitive understanding. According to the other stronger reading, nothing less than knowledge can be considered as a reason. This view implies that if p is S ’s reason for acting, S knows that p . However, if treating p as a reason for acting already entails that p is known, S ’s reasoning cannot be assessed normatively according to whether S knows that p . Therefore, this notion cannot be the one used in RKP. Thus, those who may find the notion of ‘treating p as a reason for acting’ vague can safely substitute the notion with ‘using p as a premise in practical reasoning’.

Assuming that RKP expresses a conceptually or metaphysically necessary truth, if knowledge and practical reasoning are related as RKP suggests, then the contrast between intuitive judgments about the low- and high-stakes cases discussed in section 1 is exactly what we should expect (Stanley, 2009, p. 5). If knowing that p is necessary for appropriately treating p as a premise in practical reasoning, then the lack of knowledge of high-stakes subjects provides the best explanation for why it seems irrational for those subjects to rely on p as a premise in practical reasoning. If one’s knowledge that p grants that one is in a good enough epistemic position to act upon p , then from the fact that it is irrational for high-stakes subjects to rely on p as a premise in practical reasoning, it follows that these subjects don’t know that p .

Timothy Williamson also endorses the existence of a tight normative connection between knowledge and practical reasoning. Williamson (2005, p. 227) argues that knowledge entitles one to rely on a proposition in practical reasoning. In particular, in cases in which one has good evidence that p is true but p is, in fact, false, one would not be entitled to use p as a premise in practical reasoning, but only excused to do so. Williamson also advances the following constraint on knowledge ascriptions:

KPR

A first-person present-tense ascription of ‘know’ with respect to a proposition is true in a context *iff* that proposition is an appropriate premise for practical reasoning in that context. (ibid.)

However, differently from Hawthorne and Stanley, from KPR and the normative connections between knowledge and practical reasoning, Williamson doesn’t conclude that knowledge is sensitive to practical factors. His view will be discussed in greater detail in chapter 2.

A further bi-conditional knowledge norm of practical reasoning has recently been defended by Mueller (2021). Andy Mueller suggests that practical reasoning is an intellectual performance aimed at determining which course of action to take. Mueller considers Sosa (2011)’s evaluative framework according to which belief formation is a sort of intellectual performance. As other performances, this can be assessed along three normative dimensions: accuracy, adroitness, and aptness. Mueller argues that practical reasoning can be analyzed in a similar way, along the same three dimensions. In this framework, practical reasoning is accurate if it leads to an intention to act in a way that can realize the reasoner’s end given the reasoner’s actual circumstances; it is adroit only if the exercise of one’s intellectual abilities determines an intention to act; and it is apt only if practical reasoning is accurate because it was adroit. He argues that, if an epistemic norm of practical reasoning is supposed to provide an epistemic condition constraining apt practical reasoning, the following knowledge norm seems to be a good candidate:

Knowledge norm for apt practical reasoning

Practical reasoning in which p is treated as a reason can be apt only if one knows that p . (Mueller, 2021, p. 5403)

According to Mueller, since aptness is a different notion from rational permissibility, entitlement or epistemic success, the above norm is a new variant of the knowledge norm. In addition, Mueller insists that this knowledge norm can remain neutral on whether in the bank cases the HS-subject doesn’t know q , assuming the LS-subject knows it. In his account,

it's open to hold that high-stakes scenarios call for special abilities instead of the ones sufficient for apt practical reasoning in the low-stakes scenario. But it is also equally plausible to maintain that HS-subject's reliance on q in practical reasoning could be indeed apt, and yet it might still seem defective due to a failure to assess risk properly.

1.3 Pragmatic encroachment on knowledge

Pragmatic encroachment on knowledge (also called *interest relative invariantism*, *subject sensitive invariantism* or *impurism*; henceforth, pragmatic encroachment) implies the denial of purism.⁹ Pragmatic encroachers about knowledge hold that practical factors can have a direct impact on knowledge (Fantl & McGrath, 2002, 2007, 2009a, 2009b, 2009c, 2012; Hawthorne, 2004; Ross & Schroeder, 2014; Schroeder, 2012; Stanley 2005; Weatherson, 2011, 2012).¹⁰ That is, even if two subjects both believe that p and the truth-relevant factors for p are held fixed for these two subjects, one can know that p whereas the other does not know that p due to a mere difference in practical circumstances.¹¹

Two main kinds of arguments for pragmatic encroachment can be devised from the two aspects of the connection between knowledge and practical matters that have been discussed so far (see also Brown (2013)

⁹ A terminological remark. In this book, I shall use 'pragmatic encroachment' to denote practical encroachment on knowledge. Practical encroachment broadly construed can also concern other epistemic properties, such as epistemic justification, epistemic rationality, doxastic attitudes, uses of 'knows', etc. For an introduction on various types of pragmatic encroachment, see Gao (forthcoming).

¹⁰ Fantl and McGrath (2002, 2009a) also hold pragmatic encroachment on epistemic justification. Weatherson (2012) disagrees.

¹¹ In the literature, pragmatic encroachment and interest relative invariantism are also often referred to as subject sensitive invariantism. But strictly speaking, subject sensitive invariantism is different from interest relative invariantism: The former doctrine holds that other factors pertaining to the subject in addition to practical factors—such as the salience of alternatives—are also relevant in determining whether the subject knows or not. For a version of such subject sensitive invariantism, see Hawthorne (2004, chapter 4).

and Fantl & McGrath (2012)).¹² According to the first kind of argument, pragmatic encroachment provides the best explanation of practical factor effects on knowledge ascriptions suggested by our intuitive judgments about a wide range of cases. These include most of the cases discussed in section 1.1, including those where the variations in the salience of alternatives are eliminated and non-linguistic pair cases. If we take *BANK* as an example, the argument can be formulated as follows:

- 1) It is intuitive that, in *Low Stakes*, Hannah knows that *q*.
- 2) It is intuitive that, in *High Stakes*, Hannah does not know that *q*.
- 3) The best explanation of 1 and 2 is pragmatic encroachment: whether one knows partly depends on practical factors.
- 4) If pragmatic encroachment is the best explanation of 1 and 2, we should adopt pragmatic encroachment.¹³

According to pragmatic encroachment, the stakes at issue have to be those pertaining to the subject rather than those of the attributor or the hearer.¹⁴ Hence, these versions of pragmatic encroachment can readily explain intuitive judgments about *Ignorant High Stakes* and *Low Attributor-High Subject Stakes*, but not *High Attributor-Low Subject Stakes*. Although the main versions of pragmatic encroachment can explain

¹² Here I will not consider another type of argument for pragmatic encroachment that appeals to knowledge-based decision theory advanced by Weatherston (2012).

¹³ This argument for pragmatic encroachment can be found in Stanley (2005).

¹⁴ This is the case according to the most popular versions of pragmatic encroachment, including Stanley's interest-relative invariantism, Hawthorne's subject sensitive invariantism and Fantl and McGrath's impurism. In principle, pragmatic encroachment remains open to the question of whose practical interests matter. Given the definition of pragmatic encroachment, knowledge may depend on the practical interests relative to persons different from the relevant subject. In this regard, Grimm (2015) defends a form of pragmatic encroachment according to which practical interests, either for the subject or for the evaluator, i.e., the ascriber of knowledge, or for certain third parties who might rely on the belief in question, can raise the epistemic standard relevant to knowledge. Hannon (2015) suggests that epistemic standards relevant to knowledge are fixed by the social context.

most of the practical factor effects in pair cases, they cannot deal with all of them. Some supplementary theory, such as an error theory about some of our intuitive judgments is required in order to fully address the relevant intuitive data.¹⁵ However, recall that moderate invariantism, without any additional error theory about our intuitive judgments (where ‘error theory’ here is conceived broadly, as a theory that postulates false intuitions, as opposed to postulating specific performance error), can only accommodate the intuitive judgments about *Low Stakes*. In comparison, pragmatic encroachment (or at least the main versions of it) has an explanatory advantage over moderate invariantism in terms of the number of cases directly accommodated.¹⁶ But it is not clear that pragmatic encroachment also has any advantage for what concerns the simplicity of the theory, for both views require some additional account to deal with cases that cannot be directly explained.

The second kind of argument for pragmatic encroachment appeals to the knowledge norm of practical reasoning. Assume that the variation in the salience of alternatives is eliminated from *Low Stakes* and *High Stakes*, leaving stakes as the only variant between the two scenarios. One argument uses the bi-conditional knowledge norm and judgments about the propriety of Hannah’s practical reasoning in the variants of *Low Stakes* and *High Stakes*. It can be formulated as follows:

- 1) In *Low Stakes*, but not *High Stakes*, it is appropriate for Hannah to treat *q* as a premise in her practical reasoning. (Judgments about the propriety of Hannah’s practical reasoning)
- 2) It is appropriate to treat *q* as a premise in one’s practical reasoning *iff* one knows that *q*. (The knowledge norm of practical reasoning)

¹⁵ It is important to note that the supplementary theory does not need to be an error-theory. There may be ways to try to extend pragmatic encroachment in a way that preserves the residual intuitive judgments.

¹⁶ Such reliance on intuitive judgments about knowledge is very similar to the ‘methodology of the straightforward’ that has been used to motivate epistemic contextualism. See Gao, Gerken and Ryan (2017) for a critical reflection on this methodology.

- 3) Hannah knows that q in *Low Stakes* but not *High Stakes*. (From 1 and 2)
- 4) The two scenarios differ only in the stakes. (Assumption)
- 5) Knowledge is sensitive to the stakes (and hence to practical factors). (From 3 and 4)

Another argument that uses the knowledge norm of practical reasoning doesn't appeal to the intuitiveness of relevant pairs of cases. Rather, it appeals to some general principles about knowledge and practical reasoning. Consider the following case. A subject in an ordinary situation may know that p even though her epistemic position could be strengthened, say, by gathering further evidence for p . If this were not true, then knowledge would require one having always a maximal epistemic position with respect to the target proposition, which is far too demanding, at least if we want to maintain a non-sceptical perspective about knowledge. Furthermore, as the stakes rise, a subject needs a stronger epistemic position in order to rely on the proposition that p in her practical reasoning. The higher the stakes, the stronger must be the epistemic position. We can arbitrarily rise stakes to a level at which it will not be rational for S to rely on p . Think for example of a proposition you take to know (e.g., that the Champions League final will be tomorrow); then consider whether you would be disposed to bet 10\$, then 100\$... continue like this until you reach a point at which for some proposition you take to know, you would not be disposed to bet that sum (see Weatherson (2012) for similar examples). Let us call the level at which it is not rational for S to rely on p , HS. The argument runs as follows:

- 1) Epistemic position E is sufficient for S to know p in an ordinary situation, but insufficient for S to rely on p in HS. (Assumption)
- 2) If E is insufficient for S to rely on p in her practical reasoning, then S doesn't know that p . (The sufficiency condition of the knowledge norm of practical reasoning)
- 3) E is sufficient for S to know that p in an ordinary situation but not in HS. (From 1 and 2)
- 4) The ordinary situation and HS differ only in the stakes.

(Assumption)

5) Knowledge is sensitive to the stakes (and hence to practical factors). (From 3 and 4)

As Brown (2013, p. 246) observes, there can be another argument for pragmatic encroachment if we substitute the sufficiency direction of the knowledge norm of practical reasoning with the necessity direction:

1) Epistemic position *E* is sufficient for *S* to know *p* in an ordinary situation, but insufficient for *S* to rely on *p* in HS. (Assumption)

2*) If *E* is sufficient for *S* to rely on *p* in her practical reasoning, then *S* knows that *p*. (The necessity condition of the knowledge norm of practical reasoning)

3*) The best explanation for why *S* in HS is not in a good enough epistemic position to rely on *p* in her practical reasoning is that *S* doesn't know that *p*. (From 2*, IBE)

4*) *E* is sufficient for *S* to know that *p* in an ordinary situation but not in HS. (From 1 and 3*)

5*) The ordinary situation and HS differ only in the stakes. (Assumption)

6*) Knowledge is sensitive to the stakes (and hence to practical factors). (From 4* and 5*)¹⁷

A further argument for pragmatic encroachment has been suggested by Fantl and McGrath (2009a, 2009c). Their argument relies on *fallibilism* about knowledge, which roughly is the idea that one may know that *p* even though her epistemic position is not maximal, and thus could be further strengthened. More precisely, fallibilism refers to a cluster of views that assert the compatibility of knowledge with some sort of 'epistemic lack'. Different understandings of this lack have been suggested. Some examples include: epistemic chance (given one's evidence) being less than maximal (probability 1), having a degree of justification that falls short of epistemic

¹⁷ Versions of this argument can be found in Hawthorne (2004) and Stanley (2005).

certainty, or short of entailing the truth of the known proposition.¹⁸

Fantl and McGrath argue that purism, fallibilism and the sufficiency condition of the knowledge norm of practical reasoning constitute a trilemma, and that in order to solve this trilemma we should drop purism. Their argument is in the same spirit of the above argument exploiting the sufficiency condition of the knowledge norm of practical reasoning. More specifically, it is based on a notion of fallibilism defined in terms of epistemic chance. According to this view, a subject may know that p even though there is a small, non-zero epistemic chance for her that not p . It is a consequence of this kind of fallibilism that there will be cases like *Low Stakes* in which the subject knows that p although there is a small epistemic chance for her that not- p . As long as the stakes are low enough, fallible knowledge places the subject in a good enough epistemic position to act on p . However, Fantl and McGrath argue, practical features of the situation could change in such a way to cause a raise in stakes. The stakes could raise high enough so that the chance that not- p , though small, will undermine the rationality of acting as if p were true. If we combine this result with the sufficiency direction of the knowledge norm for practical reasoning, it will follow that in *High Stakes* situations in which the chance of p is not enough to make it rational to act as if p , the subject does not know that p . But since low-stakes and high-stakes scenarios differ only in the stakes, it follows that whether one knows depends on the stakes, and hence on practical factors. According to Fantl and McGrath, the weakest of the three initial ideas is purism. Thus, abandoning purism is the best solution to the trilemma.¹⁹

¹⁸ See chapter 7 for a detailed overview of different kinds of fallibilism.

¹⁹ Fantl and McGrath's original argument is as follows (Fantl & McGrath 2009a):

- 1) If you know p , p is warranted enough to be a reason you have to ϕ . (the knowledge-reasons link, KR)
- 2) If p is warranted enough to be a reason you have to ϕ , p is warranted enough to justify you in ϕ -ing.
- 3) So, if you know p , p is warranted enough to justify you in ϕ -ing. (the knowledge-justification link, KJ)
- 4) Whether p is warranted enough to justify you in ϕ -ing can vary between a low-stakes case in which you know that p and an appropriately chosen

Another argument for pragmatic encroachment based on decision theory is due to Weatherson (2012). Weatherson's argument is based on a knowledge-based decision theory. According to this theory, the structure and information in a decision table must match one's knowledge. More specifically, a state should appear on the table if and only if one doesn't know that it doesn't obtain. If one knows that a certain relevant state in the decision table will not happen, one can ignore the consequence of possible actions under that state.²⁰ Given this theory, Weatherson invites us to imagine a case in which a subject intuitively seems to know that p . For example, you hear your partner's voice in the nearby room, and thereby come to know that your partner is at home. Then suppose that the subject of this case is offered a bet with very high stakes (e.g., 1000 years of torture if the subject bets for the truth of p and loses, but gets a candy if she wins). Intuitively, it is not rationally permissible for her to take the bet. However, if she knows that p , then the state of not- p can be eliminated from the decision table. If the not- p state is not on the table, then the dominating option (which delivers the better outcome under all the relevant states in the table) is to take the bet. If it is rationally permissible to take dominating options (as standard decision theory prescribes), then it should be rationally permissible for the subject to take the bet. But intuitively, it is not.

There are several ways to solve this problem. One straightforward way is to deny that the subject in this case ever knew that p . But this would lead to a radical form of scepticism, since for nearly all propositions we think to know it would seem irrational to bet with such incredibly unfavourable odds. The solution favoured by Weatherson is to accept that knowledge

high-stakes case, holding fixed your warrant for p across the cases.

- 5) So, whether you know p can vary with the stakes, holding fixed your warrant for p . (pragmatic encroachment)

Fantl and McGrath (2012) advance a similar argument based on fallibilism about certainty according to which knowledge that p does not require absolute epistemic certainty that p . For a criticism to this argument, see Ye (2022).

²⁰ See Fassio and Gao (2021) for a more detailed discussion of various kinds of knowledge-based decision theory.

that p can be sustained in normal circumstances, but it is defeated when the subject is offered the bet. More precisely, according to Weatherston when the subject is offered the bet, stakes are so high that the possibility that not- p should be included in the table as a possible state with non-zero probability. Given the assumption that a state should appear on a decision table if and only if one doesn't know that it doesn't obtain, in such high-stakes circumstances the subject doesn't know that p . In conclusion, a mere difference in stakes can deprive a subject of her knowledge.

In this section, we considered several types of arguments for pragmatic encroachment on knowledge. Upholders of moderate invariantism have provided several responses to these arguments. Concerning the first type of argument, moderate invariantists have provided different kinds of explanations of the intuitive asymmetry between low- and high-stakes cases. These explanations mainly fall into two categories: doxastic and pragmatic. For what concerns the second type of argument, a prominent line of response has been to reject the knowledge norm of practical reasoning (see e.g., Brown, 2008a, 2008b; Gerken, 2011, 2017). Concerning the latter type of argument based on knowledge-based decision theory, Brown (2012a) has provided a convincing objection to it. In the rest of this chapter (sections 1.5–1.6), I will examine more carefully the main moderate invariantist responses that have been pursued in the literature. But before that, I would like to discuss a range of problems for pragmatic encroachment.

1.4 Problems for pragmatic encroachment

There are mainly two types of problematic consequences of pragmatic encroachment. The first set of problems concerns counterintuitive verdicts about how knowledge could be produced. The second set of problems concerns infelicitous assertions that would be deemed to be true if knowledge were sensitive to stakes.²¹

²¹ Ichikawa et al. (2012) argues that pragmatic encroachment is in tension with core tenets of belief-desire psychology. This may constitute a further objection to pragmatic encroachment.

Problem 1: Counterintuitive verdicts about knowledge

According to Gillian Russell and John Doris (2009), a consequence of pragmatic encroachment is that factors like indifference and wealth could produce knowledge. They imagine a variant of the bank case in which, as in the original *High Stakes* case, if the subject fails to pay the deposit this will lead to a significant financial loss for her. However, in this new case the subject doesn't care about losing money. This indifference could be due to various reasons: maybe the subject is so wealthy that the material loss is trifling for him, or maybe she is so depressed that becomes totally indifferent to any material possession. Since in this case whether the bank will open on Saturday is not a serious practical question for the subject, according to pragmatic encroachment a memory that the bank was opening two weeks ago is sufficient for her to meet the standard for knowledge. An awkward consequence of this case is that being indifferent to the stakes can lead one to have more knowledge.

Problem 2: Troubles from ordinary language

Blome-Tillmann (2009a) argues that pragmatic encroachment leads to intuitively problematic consequences when we consider modal and temporal embeddings. If knowledge were sensitive to stakes, then the following sentences involving *modal embeddings* would be true:

“I know that the bank is open Saturday, but had more been at stake, I wouldn't have known.”

“I don't know that the bank is open Saturday, but had less been at stake, I would have known.”

Yet these sentences sound very counterintuitive.

Concerning sentences involving *temporal embedding*, imagine that in a variant of *High Stakes* the stakes were high on Thursday but they became low on Friday. Then, if pragmatic encroachment were right, the subject could say:

“I didn't know *q* on Thursday, but on Friday I did.”

Again, this claim sounds very odd. It seems surprising that, even though the subject didn't gain any evidence relevant to whether q is true on Friday, she could gain knowledge just because it became less important for her whether q . In general, lowering stakes doesn't seem like a good way of gaining knowledge.

The stakes-sensitivity of knowledge also predicts that the following *problematic inferences* could constitute sound reasoning:

“Less is at stake than I thought. Good, then I know that p .”

“More is at stake than I thought. That means I don't know that p .”

Moreover, if pragmatic encroachment is right, the following *first-third-person knowledge ascriptions* should be classified as reasonable:

“I know that p , but he has more at stake. So he doesn't know that p .”

“I don't know that p , but he has less at stake, and he knows that p .”

Again, these assertions sound deeply infelicitous, and the corresponding thoughts would sound clearly unreasonable.²²

Blome-Tillmann also considers cases in which two agents have exactly the same evidence, but for one (call him John) stakes on p are low while for the other (call him Paul) they are high. In such cases pragmatic encroachment entails that one can truly assert:

“John and Paul have exactly the same evidence for p , but only John has enough evidence to know p , Paul doesn't.”

Again, such an assertion sounds absurd.²³

²² Problematic inferences and first-third-person knowledge ascriptions are discussed by McGrath (2017).

²³ For responses to Problem 1 and Problem 2, see Weatherson (2011).

1.5 Moderate invariantist accounts of the practical factor effects on knowledge ascriptions

As we have seen in section 1.3, one main type of argument for pragmatic encroachment is based on intuitive judgments about practical factor effects on knowledge ascriptions. In addressing this type of argument, most moderate invariantists deny the intuition that the HS-subject doesn't know p and focus on providing accounts that explain away this intuition. In this regard, moderate invariantists have advanced two types of accounts of the practical factor effects: *doxastic accounts* and *pragmatic accounts*.²⁴ Under the label of doxastic accounts, I include a number of specific accounts according to which intuitive data about knowledge ascriptions can be explained in terms of the influence of practical matters (which may be rational practical dispositions, as in Bach (2005, 2008, 2010), Ganson (2008), Weatherson (2005), or simply practical factors, as in Nagel (2008, 2010a)) on belief.

Amongst doxastic accounts there are so-called *doxastic pragmatist* accounts, which hold that perceived high stakes raise the threshold on credence necessary for forming a normal or rational outright belief. Since the credence of the subject remains fixed across low- and high-stakes cases, but the threshold for outright belief (the type of belief required for knowledge) goes up in *High Stakes*, the HS-subject doesn't believe that q . This explains the intuitive judgment that the HS-subject doesn't know that q (Bach, 2005, 2008, 2010; Ganson, 2008; Weatherson, 2005). Another type of doxastic account, due to Nagel (2008, 2010a), holds that practical factors psychologically affect beliefs, sometimes leading to belief suspension or revision. The positive view I will defend in chapters 5–7 is a doxastic account alternative to doxastic pragmatism. I postpone further discussions of doxastic accounts to chapter 4, which is fully dedicated to the examination of these views.

²⁴ A less prominent type of account is constituted by psychological bias accounts, including epistemic focal bias account (Gerken, 2013, 2017), egocentric bias account (Nagel, 2010b) and source bias account (Turri, 2015; Turri & Friedman 2014). See Gerken (2017, chapter 10) for discussions of these alternative accounts.

Here I would like to briefly discuss the other prominent kind of accounts in more detail. According to pragmatic accounts, the divergence between intuitive judgments in low- and high-stakes cases is due to the pragmatics of knowledge ascription, i.e., the variability of the conversational propriety of ascribing knowledge, as opposed to the semantics of knowledge ascription (Bach, 2001; Black, 2005; Brown, 2006; Gerken, 2017; Hazlett, 2009; Lutz, 2014; Pritchard, 2010; Rysiew, 2001, 2007). Depending on the type of pragmatic mechanism invoked, pragmatic accounts can be distinguished into four approaches. These approaches appeal to, respectively: i) conversational implicature, ii) conventional implicature, iii) conversational impliciture, and iv) directive force. In the rest of this section, I will only discuss the most prominent pragmatic approach that has been widely absorbed and criticised in the literature, i.e. the conversational implicature approach.²⁵

According to the conversational implicature approach, which is also referred to as ‘warranted assertability manoeuvre’, the infelicity of a knowledge ascription is closely related to an implicature in the knowledge ascription. In *High Stakes*, Hannah conversationally implicates something false with a positive knowledge ascription (Black, 2005, p. 334; Brown, 2005, pp. 280–281, 2006, pp. 425–426; Lutz, 2014, p. 1729; Rysiew, 2001, pp. 486–487). By falsely denying knowledge, she implicates something true (Brown, 2006, p. 426; Lutz, 2014, p. 1737; Rysiew, 2001, pp. 486–487)²⁶.

For example, according to Rysiew (2001, 2007), one of the proponents of this view, the word ‘know(s)’ has a stable semantic meaning. In a relevant alternatives semantics, first developed by Dretske (1970) and adopted by Rysiew, *S* knows that *p* iff *S* can rule out all the relevant not-*p* possibilities—where what is ‘relevant’ is invariant across contexts. However, Rysiew also identifies another kind of ‘salient’ not-*p* possibilities that are occasion-sensitive and vary with conversational settings—where ‘salient’ refers to those counter-possibilities (possibilities that not-*p*) which

²⁵ For problems with approaches i) and iii), see Blome-Tillmann (2013). The directive force account proposed by Gerken is still under development.

²⁶ Note that some advocates of the conversational implicature approach, most notably Black (2005, p. 336), Hazlett (2009, pp. 616–619) and Pritchard (2010, pp. 89–90), argue that we should reject the intuition that denials of knowledge are true in *High Stakes*.

the parties in a given situation ‘have in mind’ (Rysiew, 2001, p. 488).²⁷ In Rysiew’s view, the salient not-*p* possibilities are semantically irrelevant but conversationally relevant to the use of ‘know(s)’. In particular, an utterance of “S knows that *p*” pragmatically imparts that *S*’s epistemic position with respect to *p* is good enough to eliminate the salient not-*p* possibilities. Given that the set of salient not-*p* possibilities can encompass a wider range of error possibilities than the set of relevant not-*p* possibilities, we end up with situations where a literally true knowledge ascription seems to be false because the sentence implicated by that ascription is false.

According to Rysiew’s explanation of the intuitive judgments of the bank cases, in *Low Stakes* the relevant possibilities match with the salient ones. Hannah’s utterance “I know that the bank will be open tomorrow” both semantically expresses and conversationally implicates a truth. In *High Stakes*, however, due to the high stakes at issue, the unconsidered and epistemically irrelevant possibilities that the bank might have changed the hours raised by his wife becomes part of the salient possibilities. Hence, by uttering a truth, “I know that the bank will be open tomorrow”, Hannah would falsely implicate that her evidence is strong enough to rule out all those possibilities. By contrast, by uttering the falsehood, “I don’t know that the bank will be open tomorrow”, Hannah can impart that his epistemic position is not so strong to rule out the possibility that the bank has recently changed the hours and hence will be closed on Saturday.

As Rysiew, Brown (2006, pp. 424–425) also employs the possible world semantics to formulate the requirement for knowledge: there is a *context-invariant* range of possible worlds across which the subject’s belief must match the facts in order to constitute knowledge. According to this view, in the *High Stakes* case, had Hannah positively ascribed knowledge to herself, she would have implicated that her belief that the bank will be open matches the facts across a *wider* range of possible world in which the bank has recently changed its Saturday hours. Other accounts of conversational implicatures in knowledge ascriptions are also available. For example, Lutz (2014, p. 1728) holds that in a context in which a bit of practical reasoning (or subsequent action on that reasoning) is salient, a

²⁷ Rysiew in (2001, p. 488) uses the terms ‘relevant’ and ‘salient’, while in his (2007, p. 637) he discusses possibilities that are ‘considered’ and ‘worth taking seriously’.

knowledge ascription implicates that one is rational to take p as true in that practical reasoning.

The conversational implicature approach has been the target of several objections. The main problems with this approach are that it is of very limited scope and ill-motivated. As many have pointed out, we not only have the intuition that it is felicitous for Hannah to utter “I don’t know that q ”, but also have the intuition that it is natural for Hannah to believe that she doesn’t know that q (see e.g. Baumann (2011, pp. 160–161), Blome-Tillmann (2013, p. 4312), Fantl and McGrath (2009a, p. 42), Roeber (2014, p. 256)). The approach at issue only covers the former linguistic part of the data, concerning conversational felicity of knowledge ascriptions, but not the latter part of the data regarding what a subject actually believes. It follows that the explanation provided by this approach cannot apply to non-linguistic cases.

In response, proponents of the conversational implicature approach might appeal to a further error theory. They may say that we are prone to mistake what is conveyed pragmatically for what is expressed semantically, and this is what we do in both linguistic and non-linguistic contexts (see e.g., Rysiew (2001, pp. 502–503, 2007, p. 648)). However, this response is unsatisfactory unless it also explains why we are prone to make such mistakes in some specific cases involving knowledge ascriptions but not in others. Baumann (2011) criticises the possibility of a ‘warranted believability manoeuvre’, accounting for the alleged confusion between truth (related to the semantic content) and what is warranted to believe (related to the pragmatic implicature), on the basis of important asymmetries between thought and language. If Baumann is right, then some alternative explanation of the non-linguistic cases distinct from the conversational implicature approach is called for. However, as Blome-Tillmann (2013, p. 4307) argues, once such an explanation is in place, it would also provide an explanation of the linguistic cases, hence rendering the conversational implicature approach redundant.

As for the objection that the conversational implicature approach is ill-motivated, it has been argued that the suggested pragmatic implicatures cannot be motivated by independent general conversational principles (Blome-Tillmann, 2013; DeRose, 2002, 2009; Dimmock & Huvenes, 2014; Petersen, 2014). This constitutes a serious problem for the approach at

issue, because any pragmatic explanation of our apparently semantic intuitions must be given in terms of general conversational principles. In response, Brown (2006) and Rysiew (2007) have illustrated that the relevant implicatures of knowledge ascriptions can be explained in terms of Gricean maxims. For example, as argued by Brown, moderate invariantists can explain how the relevant pragmatic implicatures are generated by appealing to Grice's rule of relevance according to which utterances should be relevant to the conversation. According to Brown, in *High Stakes*, given a mention of the practical importance of the issue and of the error possibilities, a very strong epistemic position with respect to *q* is made conversationally relevant. As a result, a positive knowledge ascription that *q* pragmatically conveys that one is in a very strong epistemic position with respect to *q* and a negative knowledge ascription implicates that one is not in a very strong epistemic position with respect to *q* (Brown, 2006, p. 426).

Another widely discussed objection concerns the cancellability of implicatures. According to this objection, cancellability is the best test for implicature; nonetheless, the putative implications postulated by proponents of the pragmatic approach are not cancellable (Cohen, 1999; DeRose, 2009; Dimmock & Huvenes, 2014; Roeber, 2013). First, there is an ongoing debate on whether the implicature of a positive knowledge ascription that *q* in *High Stakes* is cancellable. Utterances such as "I know that the bank will be open, but I cannot rule out that it has changed its hours", or "I know that the bank will be open, but we need to investigate further" sound infelicitous (Cohen, 1999, p. 60; Dimmock & Huvenes, 2014, p. 3249).²⁸

One line of response is to admit the uncancellability of the pragmatic implication, but deny that it constitutes a real problem for the conversational implicature approach (Brown, 2006, p. 428; Lutz, 2014, section 4.1; Rysiew, 2001, p. 496, 2007, p. 646). For example, Rysiew

²⁸ However, these claims about the intuitive felicity of such assertions could be contentious. For example, some moderate invariantists might not find the concessive knowledge ascriptions in question problematic. Similar concessive knowledge ascriptions in third-person or third-person past tense might sound less infelicitous or even felicitous to those philosophers.

argues that what is pragmatically conveyed by a knowledge ascription is universally held, which renders uncomfortable cancellations expectable. However, it has been objected that the majority of implicatures (including universal implicatures) are cancellable. Hence, it would be *ad hoc* to claim that knowledge ascriptions fall into a special category of non-cancellable implicature without further explanations (Blome-Tillmann, 2013; DeRose, 2009; Dimmock & Huvenes, 2014). Another line of response is to argue that some particular formats of implications are in fact cancellable (Rysiew, 2001, p. 495; Brown 2006, p. 428; Lutz, 2014, p. 1735). But as it has been objected by Dimmock and Huvenes (2014, section 5.2), even if implicatures of a specific type are cancellable, there are reasons to doubt that the alleged implicatures of knowledge ascriptions belong to that type.

Moreover, it has been argued that the alleged implicature of the negative knowledge ascription uttered by Hannah in *High Stakes* is clearly not cancellable (Roerber, 2013, p. 24; see also Blome-Tillmann (2013, fn.32) and Peterson (2014)). Although there might still be some space for proponents of the conversational implicature approach to argue for the cancellability of the implicature of positive knowledge ascriptions, it is hard to see how cancellations of implicatures can make sense with negative knowledge ascriptions at all. Recall that, according to the conversational implicature approach, the negative knowledge ascription in *High Stakes* implicates, for example, that the subject is not in a good enough epistemic position to act on the key proposition. However, the utterance “I don’t know that the bank will be open, but it won’t harm if we go home now and pass the bank tomorrow” in *High Stakes* sounds utterly nonsense no matter how to put it. For it is hard to see what reason Hannah could have for thinking both that she does not know that *q* and that she can reasonably act as if *q* is true, i.e. go straight home. Thus, the cancellation at issue here does not work.²⁹

Concerning the explanation of the felicity of negative knowledge ascriptions in *High Stakes*, two further problems are worth mentioning.

²⁹ According to a weaker understanding of the cancellability test argued by Blome-Tillmann (2008), it is sufficient to have the implicature comfortably cancelled in some context. With this understanding at hand, it can be shown that there are some contexts in which the alleged implicature of the knowledge denial is cancellable.

First, DeRose (1999, p. 200, 2002, p. 192, 2009, p. 114) argues that the explanation of an appearance of truth in the case of a negative knowledge ascription in *High Stakes* is different in kind from the explanation of an appearance of falsehood in the case of a positive knowledge ascription. According to Keith DeRose, the former kind of explanation seems much more problematic than the latter kind, for a false assertion will remain unwarranted despite whatever true implicatures it might generate. In reply, Brown (2006, section 3) comes up with influential examples in which an utterance may seem true since it pragmatically conveys a truth, although it is literally false, such as in cases of implicature.³⁰ For example, “I have not eaten” asserted in circumstances in which the speaker has not eaten recently seems true although it is literally false given that the speaker has eaten at some time or other in the past. The reason is that the utterance pragmatically conveys the truth that the speaker has not eaten recently.

Second, Iacono (2008) points out that the conversational implicature account of negative knowledge ascriptions is incompatible with epistemic norms of assertion. Take the knowledge norm of assertion for example. According to this norm, a proper assertion that *p* requires that one knows that *p* and hence the truth of *p*. Thus, there cannot be false but conversationally proper assertion. This is inconsistent with what is prescribed by the conversational implicature approach, according to which the negative knowledge ascription is false, though it sounds felicitous. Other epistemic norms of assertion considered by Iacono include the truth norm, the reasonable to believe norm and the belief norm. For example, against the compatibility of the pragmatic approach with a reasonable-to-believe norm of assertion (RTB), Iacono argues as follows: to account for low-stakes knowledge ascriptions, the moderate invariantist has to assume that RTB is met, i.e., that it is reasonable to believe that the low-stakes subject knows. But then that has to be assumed for high-stakes knowledge ascriptions as well, or at least the moderate invariantist is so committed.

³⁰ Implicature is conceptually independent of what’s said. However, in implicature what is meant is built up from the explicit content of the utterance by conceptual strengthening, which yields what would have been made explicit if the appropriate lexical material had been included in the utterance (Bach, 2001).

So, the invariantist will have to say that we violate RTB when we assert that the high-stakes subject does not know that *p*.

There are some further influential objections to the conversational implicature approach. For example, Dimmock and Huvenes (2014, pp. 3244–3247) argue that the approach has difficulties in explaining certain retraction judgments. More specifically, it seems natural and appropriate for the LS-subject who now finds herself in a high-stakes situation to retract her positive knowledge ascription previously made in the low-stakes situation (“I said I knew, but I was wrong”). However, the knowledge ascription made in the low-stakes situation is supposed to be true and convey nothing false.

To sum up, in this section I have examined the most prominent pragmatic account of the practical factor effects on knowledge ascriptions, the conversational implicature approach. The present discussion is not supposed to be exhaustive. The aim here was merely to provide a general survey of this type of account rather than a thorough examination. However, I hope I have conveyed that we have good reasons to doubt that pragmatic accounts can successfully explain the relevant data and to look for some alternative account of them.

1.6 Criticisms to the knowledge norm of practical reasoning

Now let’s consider the other type of arguments for pragmatic encroachment, based on the knowledge norm of practical reasoning. In addressing this type of argument, most moderate invariantists deny the knowledge norm. And since the knowledge norm of practical reasoning is primarily motivated by folk epistemic assessments of rational action, critics of this norm have tried to undermine this motivation. It has been argued that data about our ordinary use of ‘know’ only provide a very fragile basis for concluding that practical reasoning is governed by such a norm, since sometimes we use ‘know’ in a loose sense, meaning ‘truly believe’ (Littlejohn, 2009, pp. 470–471; see also Hawthorne (2000, p. 202)). In addition, it is natural to switch from the use of ‘know’ to a wide range of other epistemic and doxastic vocabularies, such as ‘certainty’, ‘having

good reason to think', 'lacking enough evidence', etc. (Gerken, 2011, 2015, 2017).

Of course, in addition to objections that try to undermine the motivations for the knowledge norm of practical reasoning, there are also direct objections to the norm. Against the necessity direction of the knowledge norm, it has been argued that in some cases the knowledge norm seems to fail to deliver the right verdict. For example, it seems that the knowledge norm cannot accommodate intuitions in Gettier-style cases in which the subject is reasonable in acting on a justified true belief, even in the absence of knowledge (Anderson, 2015, p. 345; Baumann, 2012, pp. 10–11; Brown, 2008a, section 5, 2008b, pp. 171–172; Gerken, 2011, pp. 535–536, 2017, section 6.3.b; Hill & Schechter, 2007, p. 115; Littlejohn, 2009, p. 469; Locke, 2015, p. 82; Neta, 2009, pp. 687–688).

In response, proponents of the knowledge norm of practical reasoning argue that in these kinds of cases, the agents are only *excused* for treating the propositions at issue as reasons (Hawthorne & Stanley, 2008, p. 586; Williamson, 2005, p. 227). However, the appeal to excuses has been criticised (see Brown (2008b, p. 173), Gerken (2011, pp. 539–540), Locke (2015, p. 83, fn.23) and Neta (2009, p. 688); in response, see Boulton (2017), Kelp & Simion (2017), Littlejohn (forthcoming) and Williamson (forthcoming)). There is an ongoing debate on whether this so-called excuse-maneuver could avoid the above objections. However, we can at least take the objection as a *prima facie* reason to question the necessity claim.

The sufficiency direction of the knowledge also has been criticised. Brown (2008b), Lackey (2010), Reed (2010) and Roebler (2018) came up with a variety of counterexamples in which a subject knows something but it would be inappropriate for her to act on that knowledge. Consider one of the most discussed cases, Brown's *surgeon* case:

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 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, 2008a, p. 176)

Intuitively, the claim of the nurse is felicitous. This puts pressure on SUFF. Although the relevant evaluation explicitly concerns action, it seems that it reflects a judgment about the underlying reasoning. That is, the surgeon should not treat the proposition that the diseased kidney is the left one as a granted premise in her practical reasoning and thereby remove the left kidney straightaway before double-check, even though intuitively the surgeon knows that it is the left kidney that should be removed. Similarly, Lackey argues that sometimes one cannot treat a piece of isolated, second-hand knowledge as a premise in practical reasoning. For example, it seems that an oncologist should not report a diagnosis of pancreatic cancer to her patient if all her evidence for that diagnosis only consists of isolated, second-hand knowledge and hence she should not treat that knowledge as a premise in her practical reasoning (Lackey, 2010, pp. 364–366).

However, it is controversial whether those are good counterexamples against the sufficient direction of the norm. First, these cases are open to interpretations according to which should one act on the knowledge in question, one would violate some norm associated with one's social role rather than a norm governing practical reasoning and action themselves (McGlynn, 2014, p. 136; Neta, 2009, p. 698; Weatherson, 2012). In response, Gerken (2012, 2017, chapter 6) argues that the subjects' rational beliefs about social roles are best seen as contributing to an increased warrant-demand on action, and hence the social role is one of the relevant determiners of the strength of epistemic position required. In addition, there are other counterexamples that do not involve social roles and

conventions (see e.g., Gerken (2017, chapter 6) and Lackey (2010, p. 370)).

Ichikawa (2012) argues that these cases only show that knowledge does not provide a sufficient reason for action, not that knowledge does not constitute a reason tout court to operate. In response, Gerken (2017, chapter 6) argues that the point of those cases is that knowledge does not constitute a partial and pro tanto reason which in conjunction with other known propositions suffices for operating. My own position here is more sympathetic to the critics of the knowledge norm of practical reasoning. However, here I want to maintain a position similar to the one I took with respect to the objection to the necessity claim of the knowledge norm. Thus, I only assume that the counterexamples to the sufficiency claim provide some *prima facie* reasons to doubt that claim.

Given these difficulties, some philosophers have opted for other epistemic norms that not only are compatible with the original data motivating the knowledge norm, but also provide good explanations for some of the cases in which the knowledge norm delivers the wrong verdict. Littlejohn (2009, 2012) argues that the norm of practical reasoning is justified belief.³¹ Neta (2009) argues that it is justified belief that one knows that *p*. Gerken (2011, 2015, 2017) suggests a *warrant account* according to which it is belief that *p* warranted to a degree that is adequate relative to the deliberative context. Here the deliberative context concerns circumstances that the subject rationally believes or presupposes herself to be in, and depends on a variety of practical factors, such as alternative courses of action, availability of further evidence, considerations of urgency and stakes, social roles and conventions associated with the action. In spite of the divergence among these proposals, all of them hold that the norm of practical reasoning is belief plus some other property. As in the case of the knowledge norm, these norms can come in *necessity* and *sufficiency versions* depending on whether the relevant doxastic property is necessary or sufficient for appropriateness.

Among the advocates of the alternative epistemic norms of practical reasoning, Gerken is the only one who has explicitly used his epistemic norm of practical reasoning, i.e., the warrant account, to defend moderate

³¹ Littlejohn has abandoned this view and now defends a knowledge norm. See for example Littlejohn (2013).

invariantism (Gerken, 2011, 2015, 2017, 2018). By contrast, accounts of Littlejohn and Neta can be used by pragmatic encroachers for developing variants of their views for justification or justification for knowledge respectively.³²

In chapter 3, I will raise objections to epistemic norms in general. By appealing to a series of counterexamples, such as cases in which it is rational to rely on acceptance rather than belief in practical reasoning, I will argue that neither knowing that p nor believing that p are necessary or sufficient conditions for appropriately treating p as a reason for action.

1.7 Concluding remarks

This chapter introduces and critically discusses two aspects of the intuitive relation between knowledge and practical matters: one is constituted by practical factor effects on knowledge ascriptions; the other is the role of knowledge ascriptions in ordinary epistemic assessments of practical reasoning and the idea that knowledge is the norm of practical reasoning. We have seen that based on these aspects of the relation between knowledge and practical matters, there are two types of arguments for pragmatic encroachment and against moderate invariantism. After presenting problematic consequences of pragmatic encroachment, I considered some prominent moderate invariantist responses to the arguments for pragmatic encroachment. Concerning arguments based on the practical factor effects on knowledge ascriptions, moderate invariantists have to explain away the intuitive asymmetry between low- and high-stakes cases; in particular, the intuition that the high-stakes subject doesn't have the target knowledge. There are two prominent moderate invariantist accounts: doxastic and pragmatic accounts. This chapter focused on pragmatic accounts and some important objections to them. Chapter 4 will examine doxastic accounts. Concerning arguments

³² In other places, Neta criticises pragmatic encroachment (see his 2007a, 2007b, 2012). In particular, he provides an influential argument against pragmatic encroachment, the Main Street/State argument (Neta, 2007a). But this doesn't mean that his epistemic norm of practical reasoning does not have the potentiality to be used for arguing for pragmatic encroachment.

based on the knowledge norm of practical reasoning, the main strategy adopted by most moderate invariantists is to refute the knowledge norm of practical reasoning. However, there is another way for moderate invariantists to respond to arguments based on the knowledge norm. Williamson defends moderate invariantism while also maintaining a knowledge norm of practical reasoning. In the next chapter, I will examine Williamson's approach.

High Stakes and Iterated-knowledge

In the previous chapter, I introduced a general overview of a set of issues concerning the intuitive relation between knowledge and practical matters. These issues constitute the background of the debate opposing, amongst others, upholders of pragmatic encroachment and moderate invariantists. This chapter critically explores a specific moderate invariantist approach in this debate, that of Williamson.

Williamson is one of the main proponents of so-called knowledge-first epistemology. According to this approach, the notion of knowledge is not analysable in terms of further epistemic concepts. On the contrary, knowledge itself is supposed to play a foundational role in epistemological theorizing, grounding other notions such as belief, evidence and justification, and providing normative standards for assertion, action and other attitudes. Coherently with this approach, Williamson endorses the claim that knowledge is the epistemic norm of action.

As we have seen in the last chapter, a range of arguments for pragmatic encroachment relies on the acceptance of this norm. Most advocates of the knowledge norm of action are also pragmatic encroachers and vice versa. Unlike them, Williamson is a moderate invariantist. He endorses the knowledge norm but denies pragmatic encroachment.¹ As we saw in the previous chapter, a potential problem for this approach is that it seems hard to accept the knowledge norm of action while also maintaining that a high-stakes subject knows *p*. According to the knowledge norm, one may rely on what one knows in one's practical reasoning. But in high-stakes cases it doesn't seem appropriate for the subject to act on *p*, treating *p* as a premise in practical reasoning. From this it seems to follow that the

¹ To my knowledge, the only other philosopher who accepts the knowledge norm and rejects pragmatic encroachment is Turri (2010).

high-stakes subject doesn't know *p*. Williamson suggests a very clever and intriguing strategy to avoid this problem.

In this chapter, I will focus on Williamson (2005)'s arguments addressing pragmatic encroachment. Williamson proposes two separate accounts, each of which deals with a particular argument for pragmatic encroachment. He explains the mistaken judgment that the high-stakes subject does not know the relevant proposition in terms of i) a product of psychological bias and ii) a failure to acknowledge the distinction between knowing and knowing that one knows.

The psychological bias account applies to the first argument for pragmatic encroachment that appeals to the asymmetry between our intuitive judgments about certain relevant pair cases such as the bank cases. According to this account, it is natural for us to assign more weight to considerations telling against knowledge ascription when possibilities of error are made psychologically salient in the high-stakes case. Features such as the high practical costs of error for the subject or the ascriber, or the possibilities of error described in vivid and convincing detail can make the possibilities of error psychologically salient (Williamson, 2005, p. 226)². When we are struck by the potential disastrous consequences of believing falsely, we are led to focus on considerations that tell against the ascription of knowledge to the subject—more specifically on the inadequacy of the subject's epistemic position in eliminating specific salient possibilities of errors. Psychological biases can result not only in a tendency to withdraw the positive ascription of knowledge, but also can lead to an inclination to deny knowledge to the subject (*ibid.*, pp. 234–235).

² It seems that Williamson thinks that either explicitly mentioning error possibilities or raising the stakes can give rise to a psychological salience of error possibilities. But many experiments have failed to detect an effect of salience of error possibilities when one factor is controlled independently from the other (see e.g. Buckwalter (2010), Feltz & Zarpentine (2010) and May et al. (2010)). According to DeRose (2011, pp. 89–91), high stakes combining with mentioning error possibilities would give rise to a robust intuitive judgment about high-stakes cases against moderate invariantism. Merely raising the stakes without mentioning the error possibilities or vice versa tend to be ineffective in triggering the claimed intuitive judgments. See also Dinges (2016), Hawthorne (2004, p. 164) and Nagel (2010b) for analyses of the mechanisms regulating the rise of salience of error possibilities.

Nonetheless, even if the psychological bias account can deal with the first kind of argument for pragmatic encroachment, it cannot deal with the other kind of argument based on the knowledge norm of practical reasoning. This is because this account doesn't touch the problem of whether it is appropriate for the low-stakes subject or the high-stakes subject to rely on the key proposition in practical reasoning. For this reason, Williamson proposes a second account able to deal with this problem. Williamson accepts the knowledge norm of practical reasoning and provides an error theory for why it seems inappropriate for the high-stakes subject to rely on the key proposition in practical reasoning. Basically, the idea is that what explains the intuitive judgment about the rationality of action of the high-stakes subject is a lack of second-order knowledge. The two accounts together are supposed to rebut all the considered threats to moderate invariantism posed by pragmatic encroachment.

The aim of this chapter is to provide three criticisms to Williamson's account. Section 2.1 presents the argument against moderate invariantism based on the knowledge norm of practical reasoning and Williamson's response to that argument. Sections 2.2–2.4 provide three objections to Williamson's account: i) Williamson's account delivers counterintuitive verdicts about what it is appropriate for a subject to do in high stakes; ii) contrary to what Williamson claims, S^* doesn't need higher-order knowledge in order to be regarded as appropriately relying on p in practical reasoning; iii) Williamson's account doesn't provide a good explanation of why S^* would be blameworthy if she were relying on p in her practical reasoning. Section 2.5 draws some conclusions from the discussions in this chapter.

2.1 The iterated knowledge account of high-stakes cases

Following Williamson, let's focus on the following formulation of the knowledge norm of practical reasoning:

KNP

One knows that p iff p is an appropriate premise for one's

practical reasoning.

Suppose that for a low-stakes subject S it is not very important to be right about whether some proposition p is true, but for a high-stakes subject S^* it is very important to be right about whether p is true. Also suppose that S and S^* believe p on the same, good but not very robust, epistemic grounds (e.g., a memory about what happened two weeks ago).^{3,4} Intuitively, S can readily rely on p in practical reasoning, but S^* can appropriately rely on p only after taking some extra-precaution (e.g., checking the relevant information, acquiring more evidence, and so on). Thus we have:

- 1) p is an appropriate premise for S 's practical reasoning.
- 1*) p is not an appropriate premise for S^* 's practical reasoning.

By applying KNP to S and S^* respectively, we get:

- 2) S knows p iff p is an appropriate premise for S 's practical reasoning.
- 2*) S^* doesn't know p iff p is not an appropriate premise for S^* 's practical reasoning.

From 1 and 2, we have:

- 3) S knows p .

And from 1* and 2*, we get:

- 3*) S^* doesn't know p .⁵

³ Examples include DeRose's bank case and Cohen's airport case (Cohen 1999, p. 58; DeRose 1992, p. 912).

⁴ Here I assume a sense of 'epistemic grounds' which includes only truth-relevant factors.

⁵ Williamson also considers another argument for contextualism based on a meta-linguistic knowledge norm of practical reasoning relative to first-person present-tense

3 and 3^* together are incompatible with moderate invariantism. According to the latter view, if S knows p and S^* is as epistemically well-positioned with respect to p as S is (i.e., they share the same epistemic grounds), then S^* also knows p .

In order to avoid a derivation of 3^* , moderate invariantists have to reject either 1^* or 2^* . Based on intuitive judgments supporting the truth of 1 and 1^* , most moderate invariantists reject KNP, and hence 2^* (e.g. Brown, 2008; Gerken, 2011; Roeber, 2018). Although Williamson is also a moderate invariantist, he also accepts KNP. As he says, “without KNP the concept of knowledge would lose some of its significance: one reason why it matters whether you know something is that, if you do, you are entitled to use it in ways in which you would not otherwise be so entitled.” (Williamson, 2005, p. 228) Furthermore, Williamson adopts an anti-sceptical position according to which the epistemic standard of ‘know’ can be met quite easily, i.e. most of our knowledge ascriptions made in ordinary contexts are true. This commits him to accept 3 . Thus, he must reject 3^* and explain

ascriptions of ‘know’. This knowledge norm is formulated as follows:

KNP^{*}

A first-person present-tense ascription of ‘know’ with respect to a proposition is true in a context *iff* that proposition is an appropriate premise for practical reasoning in that context.

Suppose C is a context in which p makes little practical difference; C^* is a context in which p makes an enormous practical difference to the subject. The argument for contextualism can be constructed as follows:

4) “I know p ” is true in C *iff* p is an appropriate premise for practical reasoning in that context.

4^{*}) “I don’t know p ” is true in C^* *iff* p is not an appropriate premise for practical reasoning in that context.

1) p is an appropriate premise for S ’s practical reasoning.

1^{*}) p is not an appropriate premise for S^* ’s practical reasoning.

5) “I know p ” is true in C .

5^{*}) “I don’t know p ” is true in C^* .

5 and 5^{*} imply that ‘know’ is sensitive to the ascriber’s context. Hence, we have an argument for contextualism (Williamson, 2005, pp. 227–228).

away the apparent ignorance of the high-stakes subject. Since 3^* is derived from 1^* and 2^* , and Williamson rejects 3^* and maintains KNP, he must deny 1^* .

In his response to the above objection to moderate invariantism, Williamson (2005) proposes an error theory of our intuitive judgment in 1^* . According to this error theory, although p is an appropriate premise for both S and S^* , neither of them knows that she knows p , and thus neither knows that p is an appropriate premise. However, given the high stakes, S^* needs to know that p is an appropriate premise in order to be regarded as appropriately relying on p in practical reasoning. Mere knowledge of p is not sufficient (Williamson, 2005). Since Williamson's account appeals to the lack of iterated knowledge (knowledge that one knows) in explaining away the intuitive judgment about high-stakes cases, we can name his account *the iterated knowledge account of high-stakes cases*.

Williamson's account of our assessments in 1 and 1^* is based on two key steps. First, by appealing to the non-luminosity of knowledge, Williamson argues that even though S and S^* know p , neither of them knows that she knows p . Second, Williamson argues that in order to be regarded as appropriately using p as a premise in practical reasoning, S^* (but not S) must possess *second-order knowledge*. Let us consider each step in more detail.

Concerning the first step, a condition is luminous just in case whenever one is in it, one is in a position to know that one is in it. According to Williamson, only trivial conditions are luminous; for instance, those that obtain in all cases or in none. By constructing a sorites series between a case in which the condition clearly obtains and one in which it clearly fails to obtain, Williamson argues that luminosity must fail close to the boundary between cases where the condition obtains and cases where it does not, just on the obtaining side (Williamson, 2000, chapter 4). Neither knowing a proposition nor being an appropriate premise for practical reasoning are trivial conditions. Hence they are non-luminous conditions. It follows that in some cases one is not in a position to know that one knows q even if one knows q . Likewise, in some cases one is not in a position to know that the fact that q is an appropriate premise even if q is an appropriate premise. Given KNP, q is an appropriate premise *iff* one knows q . Hence, when q is an appropriate premise but one is not in a

position to know that q is an appropriate premise, one in effect knows q without being in a position to know that one knows q (Williamson, 2005, pp. 230–231).

For what concerns S and S^* , due to the setting of the cases, according to Williamson, although both of them have the relevant first-order knowledge that p , neither of them is in a position to have second-order knowledge. This is because their knowledge falls into cases close to the boundary between knowledge and ignorance, just on the knowledge side. This seems to be plausible given the specific setting of the cases. All low-high-stakes case pairs are conceived in a way that the epistemic grounds shared by the low-stakes and the high-stakes subjects are not very strong, merely sufficient to convey the intuition that the subject in the low-stakes case knows.

As for the second step, according to Williamson, the fact that a subject relies on an appropriate premise without being in a position to know that it is appropriate provides some potential reason to question or criticise the decision. How harsh we should be with the subjects in such cases depends on how much is at stake. As Williamson says, “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.” (ibid., p. 230) Thus, given the practical situation of S^* (high stakes on whether p), in order to be regarded as appropriately treating p as a premise in her reasoning, it is not sufficient for S^* to be merely in a position in which it is appropriate for her to rely on p ; rather, S^* should also know that p is an appropriate premise.

Williamson seems to hint that the second-order knowledge requirement on high-stakes subjects is related to a corresponding requirement to engage in second-order reasoning about whether to trust the first-order practical reasoning. For example, he writes (where ‘Hi’ refers to S^* and ‘Lo’ to S):

Since the stakes are higher for Hi than for Lo, the lack of second-order knowledge is more serious for Hi than for Lo. That the plane stops in Chicago is an appropriate premise for practical reasoning for both of them (given [KNP]). However,

Hi has far more reason than Lo has to check on such practical reasoning, to engage in second-order practical reasoning about whether to trust the first-order practical reasoning. Since Hi is in no position to know that the first-order premise that the plane stops in Chicago is appropriate, the second-order premise that the first-order premise is appropriate is, although true, inappropriate (given [KNP]). Thus second-order reasoning is in no position to give a clean bill of health to first-order reasoning based on the premise that the plane stops in Chicago. Although that applies to both Hi and Lo, Hi needs the bill of health more. (ibid., pp. 232–233).

What the above quoted passage seems to suggest is that, given the high stakes situation that S^* faces, S^* has far more reason than S to check whether she can rely on her first-order reasoning, i.e. whether she can trust that p is an appropriate premise for her practical reasoning and thus, given KNP, whether she knows p . Williamson suggests that this second-order reasoning requirement is prudential: it would be highly imprudent for S^* to directly rely on p without engaging in a second-order reasoning about whether it is appropriate to use p as a premise in practical reasoning, and without reaching a positive answer to that question through this second-order reasoning (ibid., p. 233).

Combining the non-luminous condition with the second-order knowledge requirement, Williamson explains the intuitive judgment in i_1^*). According to the second-order knowledge requirement, S^* needs to have second-order knowledge of p in order to be regarded as appropriately using p as a premise in her practical reasoning. But according to the non-luminosity condition, S^* is not in a position to have second-order knowledge of p . Thus we deem that all things considered it is inappropriate for S^* to use p as a premise in her practical reasoning, even though p is indeed an appropriate premise for her practical reasoning.

Williamson also adds that if the stakes are high enough, a prudent subject should engage in even higher order reasoning—third, fourth, etc.—about whether to trust the previous-order reasoning. He says:

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. (ibid., p. 233).

It follows that the subject would need this higher-order knowledge in order for her reliance on p in her practical reasoning to be deemed as appropriate. Thus Williamson's explanation can be generalised to cases in which i) both the low-stakes subject and the high-stakes subject have n iterations of knowledge of p , while neither the low-stakes subject nor the high-stakes subject is in a position to know that she has $n+1$ iterations of knowledge, for fixed n ; and ii) stakes for both the low-stakes subject and the high-stakes subject are high, but for the high-stakes subject the stakes are even higher. In such scenarios, the high-stakes subject will always need more iterations of knowledge of p than the low-stakes counterpart. The exact value of n is determined by how much is at stake.

In addition, Williamson shows how a failure to have $n+1$ iterations of knowledge that q in deliberation could end up leading to a self-denial of knowledge that q . Williamson invites us to consider a dialogue in which one interlocutor, A , first asks another, B , (who could also be herself) whether q is the case. Then, provided a positive answer, A asks whether B can provide warrant for the answer she just gave. A continues asking the same question for each positive answer. Sooner or later, B would run out of warrant. Williamson argues that when this happens, previous positive answers will in turn be destabilised in a domino effect. Similar consequences apply when one considers whether one has warrants for various levels of higher-order knowledge. When one finds out that she lacks warrant for some higher-order knowledge of q , all lower-order knowledge of q will be in jeopardy as well. This shows that a failure of some higher-order reasoning in providing justification for the lower-order reasoning would in the end hinder one from relying on the target proposition in her first-order practical reasoning (ibid., pp. 233–234).⁶ This allows moderate invariantists to deal with cases in which high-stakes

⁶ It seems that it is open to moderate invariantists to build variations in the required number of iterations of knowledge into appropriateness itself. Accordingly, KNP would be substituted with a revised version: p is an appropriate premise in practical

subjects have second-order (or even higher-order) knowledge. In such cases, the seeming lack of first-order knowledge and of warrant to rely on it can be accounted in terms of a lack of some higher-order knowledge.

2.2 Counterintuitive appropriateness

In this section I argue that Williamson's account delivers very counterintuitive verdicts about what it is appropriate for a subject to do in high stakes. According to Williamson, both S and S^* satisfy KNP's conditions for *appropriate* use of p as a premise in practical reasoning. Two clarifications are in order about the notion of 'appropriateness' used in KNP. First, there is an issue about how to understand KNP as an epistemic norm. According to one obvious and widely acknowledged understanding, KNP is an epistemic norm in the sense that it demands that a certain epistemic condition with respect to q be satisfied in order to rely on q in practical reasoning. In a narrower sense, KNP is an epistemic norm in the sense that its normative source comes from an epistemic standard as opposed to a different normative standard (prudential, moral, aesthetic, etc.).⁷ The reason for thinking that the source of normativity

reasoning *iff* one has n -iterations of knowledge of p (n is a natural number, the exact number of n is determined by how much is at stake). Hence, in some cases, p is an appropriate premise in practical reasoning *iff* one knows that p , in others *iff* one knows that one knows that p , and so on, depending on the stakes. Such a move would provide moderate invariantism a systematic response to arguments from practical differences to shifting semantic standards for epistemic terms. However, Williamson considers this idea but does not favour it for two reasons (ibid. 231–232). First, the revised version of KNP doesn't have any advantage over the original version of KNP with regard to incorporating epistemic accessibility of appropriateness, for the non-luminosity condition applies to higher-order knowledge as well. Second, the revised version of KNP mixes considerations at different levels, which creates complications. The first-order knowledge concerns truths about the external world, while higher-order knowledge is about one's own epistemic states. Since moderate invariantists already have a response to practical arguments for shifting epistemic standards under the assumption of the original version of KNP, there seems to be no particular reason to prefer a more complicated version over KNP.

⁷ For a discussion of different sources of normativity see, for example, Broome (2013,

is non-practical is suggested by the following case. Consider a case in which a demon has informed you that the next time you use the believed proposition that $2 + 2 = 4$ in practical reasoning, he will subject you to a painful death. It seems that it would then be practically irrational to deploy $2 + 2 = 4$ in your practical reasoning (Crisp, 2005). It's likely that Williamson has the narrower sense in mind.⁸ Here, for the sake of argument, I will assume this interpretation.

Second, Hawthorne and Stanley (2008) explicate the notion of appropriateness in KNP in terms of permissibility as opposed to obligation. As they point out, "it would be overly demanding to require someone to treat all of their relevant knowledge as reasons for each action undertaken". (p. 578) Thus, the claim at issue is that it is epistemically permissible for S^* to treat p as a premise in reasoning. Put in another way, KNP holds that knowledge guarantees a good enough epistemic position to treat p as a premise in whatever practical reasoning when p is practically relevant.

With these clarifications in mind, it doesn't sound quite right to say that S^* would do anything epistemically permissible if she were using p as a premise in practical reasoning. We might be able to see the point more clearly by looking at a concrete example. Since Williamson uses Cohen (1999)'s airport case as the target example in his paper, the case I suggest is a modification of that one:

Trustful Airport. Mary and John are at the Los Angeles 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 another 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." Mary and John have to deliver an organ for an urgent transplant on a patient in Chicago. They are aware of the fact that in some rare cases the itinerary could contain a misprint

pp. 26–27, chapter 7).

⁸ Compare to his discussion of the norm of assertion in Williamson (2000, chapter 11).

or the schedule could have been changed at the last minute. Still, they rely on the information they overheard from Smith without any further check about the itinerary. It turns out that the itinerary used by Smith is reliable and provides the correct information.

It seems that Mary and John should collect more evidence and should have a stronger epistemic position in order to be justified enough to rely on that information in their practical reasoning. But according to moderate invariantism, since Smith knows that the plane stops at Chicago (henceforth *r*), and Mary and John acquire a true belief of *r* based on Smith's reliable testimony, Mary and John also know *r*. Then according to Williamson's account, it is appropriate for Mary and John to use *r* in their practical reasoning (assuming that they acquire knowledge by that testimony).

If we take seriously Williamson's account, we should be able to distinguish at least two kinds of evaluations about the subjects' practical reasoning. One evaluation would be about the epistemic permissibility of relying on *r* in one's practical reasoning; the other would concern other evaluative standards relevant in judging the subject's decision-making, e.g. prudence. In terms of the habit of decision-making exhibited, it is indubitable that Mary and John are utterly imprudent in relying on *r* without searching for any further evidence for *r*.

Now, when there are multiple evaluative standards according to which an action can be assessed and they deliver opposite verdicts, normally we can easily tell those standards apart from one another and acknowledge a conflict between the respective evaluative judgments (at least from a third-person perspective fully informed about the facts). For instance, we can easily distinguish epistemic assessments from moral assessments in the following case concerning assessments relative to assertion: While I may know that the fugitive is in the basement, and so satisfy the epistemic standard required to appropriately assert that the fugitive is in the basement, this assertion would violate a moral rule if my behaviour hinted at the presence of the fugitive to the enemy soldiers (McKenna, 2015, p. 4). Another clear case in which prudential and epistemic standards intuitively diverge is Crisp's evil demon example presented above.

Consider another case sharing the same structure of evaluation as the case of Trustful Airport. Suppose that Jimmy stops at the red light of a very busy crossroad. Suppose that he is in a country in which people respect traffic regulation only 'moderately'. So he knows that without paying careful attention to vehicles passing by, he could end up having an accident. Nonetheless when the traffic light turns to green Jimmy immediately starts the engine and drives off without checking any further. Luckily, Jimmy passes the crossroad without incurring any accident. In this case, we can easily distinguish two levels of assessment, one positive and one negative. The positive assessment is that Jimmy hasn't done anything wrong and should not be subject to any sanction considered by the law. On the other hand, Jimmy's action was imprudent. He should have checked more carefully whether some vehicle was coming from the other direction—or, as Williamson may suggest, he should have engaged in second-order reasoning about whether it was appropriate for him to proceed.

However, in Trustful Airport, a similar distinction between different standards seems to be absent. Intuitively we (as third-person observers fully informed of the facts) don't hold Mary and John as epistemically appropriate but prudentially irresponsible. Rather, it seems that our assessments about their decision to rely on the information they overheard are completely negative. In this case, we can only recognise a unique negative evaluation (be it prudential or epistemic), not two evaluations driving in different directions as it would be the case if Williamson were right. Thus holding, as Williamson does, that Mary and John have met the knowledge norm of practical reasoning is at most a theoretical speculation not supported by any intuitive judgment.

One might challenge this argument by asking why normative standards should fit with our intuitive judgments. In some cases, our intuitive judgments might be unclear or even speak against the verdicts provided by the relevant normative standards. This objection might seem quite superficial. Denying the evidential value of intuitive judgments without providing any substantive reason for why they go wrong in these circumstances is dogmatic. In response, one may argue that intuitive judgments go wrong in such cases precisely because these are borderline cases in which the subject is blind about whether she knows or not, as

Williamson's account predicts. In this vein, Hawthorne and Stanley recognise something parallel to the unclear intuition about the borderline cases:

Suppose someone knew that they had turned the coffee pot off, but having left the house are a little bit anxious about whether it is off. It is far from clear that we should craft our normative theory of action to deliver the conclusion that one ought in such a circumstance to go back and check. In general, it should be noted that 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 obtains. (Hawthorne & Stanley, 2008, pp. 585–586).

This may explain why our intuitive judgments are particularly unclear in these cases. However, this possible reply misses the point. First, even if the subject in a borderline case were not in the position to recognize the double evaluation, we as third-person assessors fully informed about the facts and the epistemic position of the subject should have no problem in discerning these evaluations. As external assessors of the case, we are not in a borderline situation in which it is difficult to discriminate whether we possess the relevant information (stakes are low for us, the description of the situation is clear, and it is stipulated that the evidence in possession of the subject would be sufficient to know in a corresponding low-stakes case). So we should acknowledge the two opposite assessments predicted by the account. But we actually don't: like the subject, we take it to be plainly inappropriate for the high-stakes subject to rely on *p*, both from an epistemic and a prudential perspective.

Second, even from a first-person perspective, we should distinguish between on the one hand knowing that there is a norm and what it requires in general, and on the other hand knowing whether the conditions in the current circumstances conform with the norm's demands. For example, consider someone driving through a crossroad who knows that the law requires stopping at red lights but is not in the position to discern whether the light is green or not due to scarce visibility conditions. In this

case, one knows that there is a norm requiring from her a specific thing but cannot know whether her action fulfils that requirement. Similarly, Williamson's account predicts that the high-stakes subject does not know whether she knows q and whether it is appropriate for her to rely on q . However, it does not predict that the subject is blind about whether the epistemic and prudential standards are in force in the context. But our intuitive judgment about relevant cases is not merely that the subject is uncertain about the assessments according to these standards. Rather, in these cases, intuitively there is only one standard in force involving a negative assessment on relying on q , also from the subject's perspective. In other words, while Williamson's account predicts that the *assessments* relative to these norms (whether the subject is acting appropriately in the circumstance) may not be transparent to the subject, it does not say anything about whether the *enforcement* of the norms is transparent to her (whether the subject knows, for example, that there is a knowledge norm governing the use of propositions as premises in her reasoning). But in order to avoid my objection, one would need the latter type of blindness, not the former.

2.3 Higher-order reasoning/knowledge and practical rationality

The second objection focuses on the higher-order reasoning/knowledge requirement in Williamson's account. More precisely, I will question the following claim:

HORK

Engaging in higher-order reasoning and/or having iterations of knowledge of q is necessary and sufficient in order for a high-stakes subject to be regarded as appropriately using q as a premise in her practical reasoning.

If HORK is false, Williamson's account of our intuitive judgments about the problematic cases in terms of second-order reasoning/knowledge is undermined.

My objection to HORK relies on a preliminary clarification of the relation between higher-order reasoning/knowledge and degrees of first-order warrant. It's not clear how we should understand this relation. According to one understanding, having higher-order knowledge of q is related to the strength of epistemic position with respect to q , measurable in terms of degrees of first-order warrant.⁹ For instance, one may argue that the closer to epistemic certainty the first-order knowledge is, the safer the corresponding higher-order beliefs are, and the higher the order of knowledge the subject is in a position to have. Since under such an understanding, having further iterations of knowledge is made possible by having stronger first-order warrant, one may question the relevance of talking of higher-order knowledge. Indeed, first-order warrant may be doing all the justificatory work, and an account of the relevant cases that focuses directly on degrees of first-order warrant (e.g. Gerken 2011, 2015, 2017) would be a simpler, straightforward alternative.

In fact, I think that higher-order knowledge should be carefully distinguished from degrees of first-order warrant. First-order warrant and knowledge iterations are very different properties. This is evident if one observes that sometimes certain methods are good for the acquisition of one of these properties but not the other (e.g., higher-order reasoning can increase the number of knowledge iterations without adding first-order warrant). Higher-order knowledge is about lower-order attitudes, not directly about the strength of one's epistemic position with respect to the target proposition. Performing second-order reasoning and consequent acquisition of second-order knowledge about q cannot alone strengthen the first-order warrant with respect to q . Rather, first-order warrant seems

⁹ Since the notion of degrees of warrant is used interchangeably with one's strength of epistemic position with respect to a proposition, one has to distinguish it from the talk of epistemic probability. According to Williamson, knowing p implies that one's epistemic probability of p is 1. But, as it has been clearly argued by Brown (2010), it doesn't follow that knowing that p implies that one's strength of epistemic position with respect to p reaches the maximum. According to Williamson, knowledge can be acquired merely based on evidence in terms of factive states such as seeing. But intuitively, as Brown argues, one's epistemic position would be stronger if, in addition, one's belief were based on evidence through other means, such as confirmation from others, auditory as well as visual information or consultation with an expert.

to depend exclusively on first-order evidence supporting q .

It is worth noting that the above considerations stand even if one conceives first-order warrant in terms of reliability or safety of one's first-order belief. Williamson observes that belief's safety from error only grants that one is *in a position to* have a certain number of knowledge iterations (Williamson 2000, p.116, Appendix II). Nonetheless, for actually possessing this iterated knowledge, the subject should draw deductively the pertinent conclusions from what she knows. This implies that when the safety of first-order belief necessary for one's being in a position to have these knowledge iterations is already in place, higher-order inferential reasoning can increase the number of knowledge iterations without adding reliability to first-order belief and first-order warrant.

Moreover, to the extent that degrees of first-order warrant are much more fine-grained than the number of iterations of knowledge, it follows that the two properties have different extensions. Assume, as is plausible in many ordinary cases, that degrees of first-order warrant are related to degrees of safety of first-order belief: the stronger the warrant for believing q , the more remote the cases in which q is false from those in which q is true, and the safer the belief that q . Now, in the Williamsonian framework, different degrees of safety of first-order belief can grant the same safe margin from error sufficient for, and only for, putting one in a position to have a specific number of knowledge iterations K^n . This implies that a specific number of knowledge iterations is compatible with different degrees of first-order warrant. On a topological conception of safety like that discussed by Williamson (2000, section 5.3), we can think of margins for error as limits of regions in an n -dimensional Euclidean space. The distance between the margin for safely believing q and the margin for safely believing q can occupy several points in the region, each of which corresponds to a different degree of reliability or first-order warrant. This implies that the reliability and first-order warrant of one's belief can be strengthened without necessarily enabling the subject to be in a position to have more iterations of knowledge.¹⁰

Furthermore, there are also reasons to separate the strength of first-

¹⁰ If one finds a topological conception too abstract, one is free to think of more concrete examples in which new evidence is sufficient to strengthen one's belief

order warrant from first-order belief's reliability (and consequently from the knowledge iterations one is in a position to have). Consider a trivial condition *C*. Examples include conditions one is always in (e.g. I exist), and conditions which cannot be unsafely believed because they obtain in every possible world (e.g. necessary truths and tautologies). According to Williamson, *C* is a luminous condition, one that a subject is in a position to know that it obtains whenever it obtains. Consider a belief about condition *C*. Such a belief enjoys maximal safety. Williamson's anti-luminosity argument cannot prevent possible infinite knowledge iterations about this condition. However, it seems that no matter how good the subject's evidence in support of the obtaining of *C* already is, new evidence would further strengthen her first-order warrant for that condition. For example, no matter how well supported my maximally safe belief in the law of excluded middle is,¹¹ if I come to know that some famous mathematician discovered a new theorem which further confirms the truth of the law, my first-order warrant for my belief will be further strengthened by this new evidence.¹²

It follows from the above discussion that there are cases in which a subject is in a position to acquire infinite knowledge iterations but doesn't have a maximal first-order warrant. We can also conceive opposite cases in which a subject has very robust first-order warrant for *q* but does not

reliability but not enough to pass the threshold for making the belief reliably reliable. For example, I can discriminate a slight increase in the heat of a surface between time t_1 and t_2 . Furthermore, my perception of the heat at both times is of a degree sufficient to grant the safety of my belief that from t_1 the surface has been more than 30 degrees Celsius, but insufficient to grant a safely safe belief in the same proposition. Thus, the reliability of my belief at t_2 is higher than the reliability of my belief at t_1 , but both of them are not enough to make my belief reliably reliable. See Williamson (2000, section 5.3) for further discussions and examples.

¹¹ I am here assuming that this truth is presented to me in a simple tautological guise, i.e., I have the concepts to formulate the tautology. See Williamson (2000, pp. 107–108).

¹² The latter consideration relies on the crucial point that while safety from error is a modal property of belief, warrant is a property related to the actual support that a certain body of evidence provides to a proposition. Some Bayesian epistemologists have suggested specific methods to measure this kind of support. See Joyce (2005) for an overview.

have higher-order knowledge of q . We have already appreciated the fact that for actually possessing higher-order knowledge, the subject should deduce the pertinent conclusions from what she knows. One can have a very strong first-order warrant but fail to have higher-order knowledge simply because one doesn't engage in deductive second-order reasoning and thereby doesn't form the relevant higher-order beliefs.¹³ In addition, there are also cases in which it is at least physically impossible for the subject in possession of a strong first-order warranted belief that q to form higher-order knowledge that q . Such cases are possible with subjects who have difficulties in forming higher-order attitudes and engaging in higher-order reasoning, such as young children and subjects who lack the concept of knowledge.^{14, 15}

Once we accept that having higher-order knowledge can diverge in important ways from having strong first-order warrant, we are in a better position to assess HORK. First, engaging in higher-order reasoning and having multiple iterations of knowledge seem to be unnecessary for being regarded as appropriately using q as a premise in practical reasoning in high stakes situations in which the subject's epistemic position with respect to q is extremely strong or reaches the maximal degree. Consider an analogue of the bank case in which the evidence held by Keith that

¹³ This is easy to imagine if we consider cases in which the only way in which a subject can acquire higher-order beliefs is through higher-order reasoning, but the subject didn't yet engage in any such higher-order reasoning.

¹⁴ Empirical studies under the label "false belief test" in developmental psychology support the view that young humans and chimpanzees seem to have difficulties in employing higher-order thoughts. Studies show that young humans and chimps are not reliable in correctly predicting other subjects' actions although it is made clear that the other subjects have been misled in a way as to hold false beliefs. Many explanations of this phenomenon have been proposed. But the claim that the tested subjects do not fully possess the concept of belief still appears to be a plausible explanation. See Wellman et al. (2001) for a useful meta-analysis. For a critical discussion of these studies and their importance for the present debate, see McGlynn (2017).

¹⁵ Williamson recognizes this type of case, see Williamson (2000, p. 95, pp. 107–108, p. 115). We can also conceive cases in which it is metaphysically impossible for the subject in possession of a strong first-order-warranted belief that q to form higher-order knowledge that q . Consider the case in which q is the proposition that the subject has no second-order attitudes.

the bank will be open tomorrow (henceforth *b*) is that just a minute ago he read the bank's opening hours in front of the bank and also got a confirmation about that from the staff working at the reception desk.¹⁶ In this case, it seems that the strength of Keith's epistemic position with respect to *b* is pretty robust. Given the massive evidence Keith has for the proposition on which he acted, should we really blame him for not asking himself whether he really knows *b*, whether he knows that he knows *b*... and for not forming a second, third... *n*-order belief about the matter? There is a clear intuition here that Keith is fully blameless in this situation. It seems obvious that Keith cannot be criticised for directly relying on *b* in his practical reasoning without first engaging in higher-order reasoning and forming higher-order knowledge of *b*. His first-order warrant is more than sufficient to justify him to rely on *b*.

This point is also supported by empirical studies on strategy selection in decision-making. Psychological studies show that very often our heuristic mechanisms systematically evaluate whether our epistemic position is good enough for action. In these studies, there is a widespread consensus that the amount of cognitive effort to allocate for a given task is not typically calculated by personal-level conscious reflection on the merits of the various alternative ways of acting. As Jörg Rieskamp and Philipp Otto note, if strategy selection is always a consequence of applying a meta-strategy, one could "run into a recursive homunculi problem of deciding how to decide" (Rieskamp & Otto, 2006, p. 207). Rather, as the different rival theories on strategy selection agree, variations in cognitive effort are automatically triggered by decision environment.

According to the so-called 'adaptive toolbox' approach, we have a broad repertoire of strategies, from systematic and heuristic ones to deliberative and controlled ones. Different strategies are used for different problems, and the 'selection' of the strategy is thought to be largely driven by the environment and computed automatically (Gigerenzer & Todd, 1999). According to the alternative 'evidence accrual' approach, what changes

¹⁶ I conceive the case as one in which the source of information is more reliable than the original bank case. If one doesn't find the case helpful, one is free to change the case in accordance with what she thinks is necessary in order to select the reliable source.

in different circumstances is not the particular strategy, but the evidence threshold for the quantity of information that one option needs to reach in order to be favoured over other options (Lee & Cummins, 2004; Newell, 2005); the evidence threshold would typically be computed automatically given influences of various environmental features. Selecting a more demanding strategy from the toolbox or raising the evidence threshold to a higher level happens when one's epistemic position is not obviously strong enough for given purposes (see e.g. Epley & Gilovich (2005) and Newell & Lee (2011)).¹⁷

These psychological theories agree that in a perceived high stakes case, when people are close to being epistemically certain that p , they tend to act on p without having a second-order reflection on their epistemic standing on p . The stability and frequency of our reliance on heuristic mechanisms make it plausible that those mechanisms are evolutionarily selected and rationally acceptable. Even though in high stakes situations one may well engage in a second-order reasoning, one can also reasonably rely on a proposition in practical reasoning without deliberating on whether to trust one's first-order reasoning. These empirical data also suggest that whether a second-order reasoning is rationally or prudentially required does not depend on what is at stake, but rather on the salience of the question whether one knows that proposition.

HORK has other implausible consequences. First, according to Williamson's account, a progressive increase of stakes would give rise to an increasing demand of higher-order reasoning and knowledge iterations in order for the subject to be regarded as appropriately relying on a proposition in practical reasoning. Now, while it still seems somewhat intuitive that in some high-stakes contexts one should engage in second-order reasoning,¹⁸ it sounds rather odd to also demand the subject to engage in a third-order reasoning about whether to trust the second-

¹⁷ I am indebted to Nagel (2010a, pp. 411–412) for discussion of the above empirical studies.

¹⁸ After all, in such circumstances, it is reasonable for one to be reflectively conscious of one's epistemic position with respect to q when the practical costs for being wrong about q are severe, and a failure to engage in a second-order evaluation would be reckless and irresponsible.

order reasoning about whether to trust the first-order reasoning. After all, the decision to be made in such cases is about the first-order question whether to treat q as a premise in practical reasoning. A third-order reasoning seems to be simply off the track in adjudicating the initial question concerning what to do. In this respect, appealing to higher-order reasoning in accounting for intuitive judgments about paradigmatic high-stakes cases appears to be a misplaced strategy.

It is also excessively pretentious to ask high-stakes subjects to engage in many levels of reasoning and possess many iterations of knowledge. For one thing, when ordinary people engage in deliberation about whether to perform a certain action, they rarely or never raise questions about whether they know the premises in their deliberations, let alone questions about even higher-order knowledge (cf. Greco 2014, p. 170). Concerning the maximal iteration of knowledge one can properly attribute, there is empirical evidence that neurotypical adults can only reliably track higher-order *interpersonal* mental state attribution, as in “Mary thinks that Tom doubts that Vivian hopes to stay”, at most for five levels (Kinderman et al. 1998; Stiller & Dunbar 2007). Presumably, higher-order *intrapersonal* knowledge attributions are subject to similar or even stricter natural limits. Engaging in very high orders of reasoning about knowledge iterations is definitely impossible for normal human beings given our inability to properly track many levels of reasoning.

Having higher-order knowledge and engaging in higher-order reasoning may not be necessary for one to be regarded as appropriately using q as a premise in practical reasoning in high stakes, but would it be at least sufficient? There are reasons to think that the answer is no. One may say that higher-order knowledge of q is indirectly relevant to q . But still merely having higher-order knowledge doesn’t necessarily strengthen the warrant for q . As we saw above, there are trivial and hence luminous and maximally safe conditions for which the anti-luminosity argument cannot prevent infinite knowledge iterations. Now, suppose a subject S^* has several knowledge iterations of q about a luminous condition C . No matter how many knowledge iterations S^* has, we can conceive cases in which stakes on q are so high that S^* would not be regarded as appropriately using q as a premise in practical reasoning. Many share the intuition that it would be inappropriate and imprudent to bet on any proposition

whatsoever when stakes on q are sufficiently high, including propositions we take to be tautologies.¹⁹

There are cases in which engaging in higher-order reasoning and acquiring higher-order knowledge would even be regarded as inappropriate and detrimental for subjects in high stakes. These cases constitute counterexamples to both necessity and sufficiency directions in HORK. Consider, for example, cases in which an exclusive focus on first-order considerations is mandatory in order to complete important tasks. Take a case in which a pilot is trying to land a full-loaded aircraft on a river with broken engines. Any reflection about first-order practical reasoning (e.g., whether she really knows that pushing further a lever would have such and such consequences) could distract her from completing crucial procedures for a safe landing.²⁰ Similar considerations apply to higher-order knowledge: imagine the pilot has not yet formed a higher-order belief about what she knows concerning emergency landings and is in the position to form it only by engaging in higher-order reasoning. Since she shouldn't engage in higher-order reasoning, she may also lack the corresponding second-order knowledge.

In sum, in this subsection I have argued that engaging in higher-order reasoning and/or having higher-order knowledge seems neither necessary nor sufficient for the subject to be regarded as appropriately relying on the target proposition in practical reasoning in certain high stakes situations. This constitutes a serious problem for Williamson's account. As shown in this section, what really matters in determining whether it is appropriate to use q as a premise in one's practical reasoning in a context is the degree of first-order warrant for that proposition, and this is obviously related to the number of knowledge iterations.

¹⁹ See e.g., Hawthorne (2004, p. 29, fn. 72), Hawthorne and Stanley (2008, p. 587), Fantl and McGrath (2009, p. 189), Reed (2010, pp. 228–229).

²⁰ The case is inspired by a similar one in Markovits (2011, p. 157) which in turn draws on a real case. In the same article you can find other similar cases.

2.4 Non-luminosity, higher-order belief and blameworthiness

For the sake of argument, let us grant with Williamson that it is reasonable to regard S^* 's failure to engage in a second-order reasoning as imprudent and that having or lacking higher-order knowledge plays an important role in explaining why we would blame S^* for relying on p in her practical reasoning. In this section, I want to question that a lack of second-order knowledge *for the reasons provided by Williamson's anti-luminosity argument* can do this explanatory work. For Williamson, the lack of second-order knowledge is due to the failure of a safety requirement on belief. This cannot explain why S^* should be considered imprudent for failing to have second-order knowledge. In general, blaming someone for being imprudent requires some wrongdoing for which the subject is fully responsible, and thus which is recognisable from her own perspective. A proper explanation of this blame in the second-order knowledge failure would require that this failure be due to a lack of reasonable second-order belief—a belief that one knows q . Unfortunately, Williamson's anti-luminosity argument doesn't show failure of second-order knowledge due to failure of reasonable second-order belief.

Let me elaborate this point a little further. In the anti-luminosity argument, Williamson aims to show that any non-trivial condition is not luminous. The original argument considers cases involving indiscriminate transitions between a situation in which a condition clearly obtains (e.g., feeling cold) and one in which it clearly fails to obtain (i.e. feeling hot). In these cases there is a point at which the subject is confident that, for example, she feels cold at a time α_i , which is true but unsafe, i.e. such that at a very close time α_{i+1} the subject is still confident that she feels cold but it is not true that she feels cold. Luminosity fails precisely in these circumstances, in which one's confidence about a relevant proposition (e.g., that one feels cold), although true, is not safe from error, i.e. it is not reliably based. The argument applies to every other non-trivial condition, including knowledge itself.

This specific feature of the argument is particularly important for our discussion of Williamson's account: if S^* 's knowledge that p is non-

luminous due to the reason provided by the anti-luminosity argument (i.e., lack of safety), it is compatible with the conclusion of the anti-luminosity argument that S^* is confident enough to believe that she knows p . Indeed, Williamson's argument is compatible with cases in which the subject is in a position to reasonably believe that she knows p , but not in a position to know that she knows p , due to the fact that that belief is unsafe.

It seems uncontroversial that one can be blameless even if one's second-order belief falls short of knowledge, provided that the belief meets minimal rationality demands. An obvious example is provided by the barn *façade* case, in which a failure of knowledge is due to a modal environmental condition. Most philosophers agree that in such cases the subject's belief is justified, or at least reasonable and excusable, and therefore it is not worthy of blame. Prudential blameworthiness seems not to be related to the failure of external conditions on knowledge, such as the lack of safety or sensitivity. Rather, for being blameless to ϕ in this sense, many hold that it is sufficient to have reasonable belief that one should ϕ , though one shouldn't ϕ . A subject who acts on what she reasonably believes is fully excusable, and thus not blameable as imprudent.

If this is correct, assuming KNP, the subject cannot be considered imprudent in cases in which she reasonably believes that she knows q but doesn't know that she knows q . Since a failure of the safety condition doesn't undermine any of the conditions for having a reasonable second-order belief, the non-luminosity of knowledge in high stakes has no direct relevance to whether one is blameable as imprudent or not. So a subject in a high-stakes situation who reasonably believes that she knows q is prudentially blameless in holding the second-order belief and in using q as a premise in her practical reasoning, even though her second-order belief is unsafe (and thus the subject doesn't know that she knows q). Therefore, the reason why luminosity fails for knowledge (according to the anti-luminosity argument) doesn't explain why S^* would be blameworthy if she were relying on p in practical reasoning.²¹

In the case of S^* , our intuition that it would be blameworthy for S^* to

²¹ Similar arguments apply to other iterations of knowledge as well. By discriminating reasonable belief from safe belief, we can see that iterations of knowledge are not

treat p as a reason in practical reasoning is not due to her lack of second-order knowledge. At most it could be due to the fact that S^* cannot reasonably believe that it is appropriate to rely on p . But Williamson's anti-luminosity argument doesn't show that in cases in which luminosity fails, one is also not in the position to reasonably believe oneself to be in the relevant condition.

It might be argued that in typical high-stakes cases exemplified in the literature, should the subjects believe that their first-order beliefs constitute knowledge, those second-order beliefs would be not only unsafe, but also unreasonable, and thus blameworthy. These cases are set up in such a way that it seems clear to the subjects themselves that the evidential grounds for their first-order beliefs are not very robust, not enough to possess iterated knowledge of the relevant proposition. If these subjects were considering whether they know the relevant proposition, the reasonable attitude for them to take would be agnosticism. Thus, although we can in principle separate the belief's properties of being unreasonable and of being unsafe from error, in practice it seems very hard to distinguish them in those high-stakes cases. As a result, in all these cases the subjects would be considered blameworthy for not suspending judgment on whether they know the relevant propositions.

Nonetheless, even admitting that lack of reasonability and of safety from error cannot be easily separated in standard high-stakes cases, the point still remains that in those high-stakes cases the blameworthiness for relying on p in practical reasoning is due to the unreasonableness, not the lack of safety, of the second-order belief that p . Williamson's account would still be wrong in virtue of appealing to the wrong *explanans* (viz., anti-luminosity and lack of safety), even if it were eventually able to provide the right verdict for most high-stakes cases.

Furthermore, it is not difficult to find high-stakes cases in which it seems fully reasonable for the subject to believe that one knows the relevant proposition, even though that belief is unreliable. Williamson's account fails to provide the correct verdicts about these cases, since the lack of second-order knowledge doesn't prevent the reasonability and blamelessness of holding second-order belief and of relying on the relevant

directly related with prudential blameworthiness.

proposition in practical reasoning.

Consider a specific example similar to the case of feeling cold in the original anti-luminosity argument. Suppose that an investor allocates a certain capital in an investment with the expectation of a future financial return. The return on her investment has the property of growing slowly but stably in value when economic circumstances are favourable, but of collapsing below the initial value in unfavourable circumstances. Though circumstances are favourable now, a downturn might occur in the near future. Assume that the investor is a normal risk-averse subject: her marginal utility decreases with an increase in monetary value. The investor is reasonably disposed to keep her investment until it passes a certain threshold, after which it would be too risky not to sell (i.e., the expected disutility of the loss would be higher than the expected utility of the gain). Suppose also that throughout the process the investor regularly considers whether she knows that the investment is not risky at the moment. The case is designed in such a way that the risk of not selling the investment grows very slowly, almost imperceptibly, but steadily, as happens in Williamson's original case with the feeling of heat. As in that case, the investor's power of discriminating the value is limited: the investor cannot discriminate any significant change in risk within a few hours (though she can over several days). The subject is very confident that the investment is not risky at the beginning, and then her confidence gradually diminishes with the growth of risk. In this respect, the subject's confidence at each stage of this process should be considered as reasonable as that of the subject in Williamson's original case. An application of the anti-luminosity argument shows that there is a time t at which the subject's belief that she knows that the investment is not risky is true and reasonable but unsafe, i.e. such that it is too close to a case in which the subject loses knowledge that the investment is not risky.²² In this case, it seems that at time t the subject is not blameworthy for taking herself to

²² Note that in the close case in which the subject loses knowledge that the investment is not risky, the loss of knowledge is due to the unsafety of the belief that the investment is not risky, not the falsity of that belief. Otherwise, at time t , the subject would already lose knowledge that the investment is not risky since her belief in the relevant proposition at time t would be unsafe.

know that the investment is not risky and for not selling. The situation is analogous to the original ‘feeling cold’ case, where the subject is not unreasonable or blameworthy for believing that she feels cold when this belief is unsafe just for an *indiscriminable* margin. Nonetheless, at time *t* our investor lacks knowledge that she knows that the investment is not risky. Williamson’s account predicts the wrong result about this case, since the lack of second-order knowledge doesn’t prevent the reasonability and blamelessness of holding second-order belief that one knows and of relying on the relevant proposition in practical reasoning.²³ This type of case shows that Williamson’s account is doubly wrong: not only does it identify the *explanans* with the wrong property (iterated knowledge instead of reasonable belief), but it also fails to predict rational behaviour in a range of cases in which conditions for blameworthiness come apart from those for safety.

2.5 Concluding remarks

In conclusion, Williamson’s account is problematic, and thus insufficient to defend moderate invariantism against a prominent argument for shifting epistemic standards. Of course, Williamson might insist that, despite these problems, his approach is still better than other non-sceptical moderate invariantist approaches. For example, he might appeal to other virtues of KNP to defend this principle. He could argue that KNP is arguably simpler than other principles and it accommodates the role of knowledge ascriptions in our ordinary epistemic assessments of practical rationality in more natural terms. Problems in dealing with abnormal cases such as high-stakes cases can be balanced by these and other theoretical virtues of KNP. However, a discussion of these alleged further

²³ Let me just observe here that this example is a high-stakes case in the sense used in the relevant literature (e.g., Fantl & McGrath, 2002; Stanley, 2005; Weatherston, 2005), one in which a lot turns on whether the relevant proposition (that the investment is not risky) is true. The reader should be careful not to confuse the investment’s risk (measured by the ratio between the disutility of the loss and the utility of the gain) with the stakes on whether the investment is not risky. While the former grows with time, the latter are high at all times in the process.

virtues is beyond the scope of this chapter. My more modest aim here was to show three specific problems affecting Williamson's account.

The previous discussions suggest that maintaining the knowledge norm of practical reasoning while also accepting moderate invariantism may be problematic. A more promising way to defend moderate invariantism could then be to reject the knowledge norm of practical reasoning. In the next chapter, I will propose a series of counterexamples against epistemic norms of practical reasoning in general, and hence against the knowledge norm as well. By undermining the knowledge norm, the argument for pragmatic encroachment based on this norm would also be undermined.²⁴

24 A significant part of this chapter is adapted from Gao (2019a).

3

Rational Action Without Knowledge (and Vice Versa)

In chapter 1, I have reviewed some popular criticisms to the knowledge norm of practical reasoning. In this chapter, I aim to provide further objections to the knowledge norm. These objections are more general than those explored in chapter 1. They do not only apply to the knowledge norm, but also to other epistemic norms of practical reasoning. Section 3.1 introduces the relevant norms and provides some preliminary remarks. Sections 3.2 and 3.3 propose two types of counterexamples to the knowledge norm. Section 3.2 considers counterexamples to the claim that knowledge constitutes a necessary condition for rational action. One such example involves a typical scenario of scientific enquiry in which scientists can appropriately treat as reasons for action propositions of a theory they believe to be false but good approximations to the truth for present purposes. Cases based on a variant of Pascal's Wager and actions performed by a sceptic also illustrate the point. Section 3.3 suggests cases against the sufficiency direction of the knowledge norm. I show that in certain circumstances, it can be unreasonable for a scientist to reason from propositions of a theory she knows to be true. Section 3.4 considers and addresses a possible reply to my criticism. Section 3.5 draws conclusions from the discussions in this chapter.

3.1 The epistemic norms of practical reasoning

While my criticisms can be extended to every formulation of the knowledge norm of practical reasoning, in this chapter, I will focus on a

specific version of it, suggested by Hawthorne and Stanley (2008):¹

Reason-Knowledge Principle (RKP)

Where one's choice is *p*-dependent, it is appropriate to treat the proposition that *p* as a reason for acting *iff* you know that *p*.

RKP can be split into the two following conditionals:

NEC

Where one's choice is *p*-dependent, if it is appropriate to treat *p* as a reason for acting, then *S* knows that *p*.

SUFF

Where one's choice is *p*-dependent, if *S* knows that *p*, then it is appropriate to treat *p* as a reason for acting.

As we saw in section 1.6, the claim that knowledge is the norm of practical reasoning has been followed by a wave of criticism. It has been remarked that data about our ordinary use of 'know' only provide a very fragile basis for concluding that practical reasoning is governed by such a norm, since sometimes we use 'knowing' in a loose sense, meaning 'being certain' or 'truly believing' (Gerken, 2011, 2015; Littlejohn, 2009). It has been argued that the same alleged data used to motivate the knowledge norm can be explained assuming other epistemic norms as well (Gerken, 2011, 2015; Littlejohn, 2009; Neta, 2009). Moreover, in some cases the knowledge norm seems to fail to deliver the right verdict. For example, it seems that the knowledge norm cannot accommodate intuitions in Gettier-style cases in which the subject is blameless in treating a justified true belief that *p* as a reason for acting, even in the absence of knowledge (Brown, 2008a, 2008b; Gerken, 2011; Littlejohn, 2009).

Given these difficulties, some philosophers have opted for other weaker principles that not only are compatible with the original data motivating the knowledge norm, but also provide good explanations for the cases in

¹ This norm has been introduced in section 1.6, where I also discuss some other objections to it.

which the knowledge norm delivers the wrong verdict. Other alternative epistemic norms of practical reasoning proposed so far include the following:

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, 2015, 2017)

The Reason-Justified True Belief Principle

Where your choice is *p*-dependent, it is appropriate to treat the proposition that *p* as a reason for acting *iff* you are justified in believing *p* and *p* is true.² (Littlejohn, 2009, 2012)

JBK-Reasons Principle

Where *S*'s choice is *p*-dependent, it is rationally permissible for *S* to treat the proposition that *p* as a reason for acting if and only if *S* justifiably believes that she knows that *p*. (Neta, 2009)

In spite of the divergence among these proposals, all of them hold that the norm of practical reasoning is belief plus some other property. Let's call these epistemic norms of practical reasoning *doxastic norms*. As in the case of the knowledge norm, these norms can come in *necessity and sufficiency versions* depending on whether the relevant doxastic property is necessary or sufficient for appropriateness.

Against all these views, I doubt that there is an epistemic norm concerning the appropriateness conditions for treating a proposition as a reason for acting.³ In this chapter, I provide two counterexamples to the knowledge norm. I show cases in which it is appropriate for a subject to treat *p* as a reason for action even if the subject does not know

² Littlejohn has recently abandoned this view and now defends a knowledge norm. See Littlejohn (2013) for his recent view.

³ Brown (2008a, 2008b) questions the existence of such norms as well.

that p . These are cases in which actions are grounded in acceptance and performed by a sceptic. These cases show that knowledge is not necessary for appropriately treating a proposition as a reason for action. Furthermore, I argue that, under a certain interpretation of epistemic norms, the first case constitutes a counterexample also to SUFE, according to which knowing that p is epistemically sufficient for appropriately treating p as a reason for acting. In addition, these cases are also good counterexamples against the alternative doxastic norms mentioned above including belief as a requisite, since in both types of counterexamples the subject does not even hold a belief about the relevant proposition. My final conclusion is that, even if knowledge, as well as justified belief, warranted belief, and similar doxastic attitudes, play an important role in the rationalization of many of our actions, these attitudes are not necessary for appropriately treating a proposition as a reason for action. Moreover, according to a specific understanding of epistemic norms, they are even not sufficient. In many circumstances, different mental attitudes, such as acceptance, provide us with appropriate bases for action. Such cases show that there is no epistemic norm governing practical reasoning.⁴

Before proceeding further, three clarifications are in order. First, the upshot of my arguments is not that there is no norm at all governing practical reasoning. My arguments are consistent, for example, with the existence of other non-epistemic norms governing practical reasoning. The aim of my arguments is rather to provide counterexamples to the claim that practical reasoning is governed by an *epistemic* norm—a norm whose satisfaction condition is constituted by an epistemic notion such as (justified or warranted) belief or knowledge. This is precisely the crux of the debate on epistemic norms of practical reasoning introduced above.⁵

⁴ At least if such a norm is conceived as an exceptionality principle valid for every possible premise of a practical reasoning. This is precisely how philosophers engaged in this debate conceive such a norm. The arguments in this chapter are compatible with epistemic norms ranging on some proper subset of such premises.

⁵ For example, Gerken specifies that the type of norms relevant to the present debate only speak to the epistemic conditions under which p may serve as a premise in practical deliberation or as a reason for action. According to Gerken, authors engaged in this debate are interested in the distinctively epistemic conditions on rational use of p (Gerken, 2011, p. 531, fn. 3). For similar remarks see, for example Brown (2012a,

Second, my examples aim to show that there are cases in practical reasoning in which it is appropriate to reason from premises that are not known or believed. This is compatible with the claim that other premises used in the same reasoning are known or believed. In all my examples below, the subject is rational in using premises she doesn't know or believe in reasoning in which other premises are known. This is sufficient to show that there are no universally valid epistemic norms like RKP.⁶

The third clarification concerns the specific sense in which these norms for practical reasoning count as epistemic (see also my related discussion in section 1.2). According to one obvious understanding, these norms are epistemic in virtue of the fact that they demand that some epistemic condition with respect to p be satisfied for it to be appropriate to use p as a premise in practical reasoning. According to another understanding, such norms are epistemic because they assess whether it is *epistemically appropriate* to use p as a premise in practical reasoning—where 'epistemically' characterizes the type of appropriateness and differentiates it from other types of appropriateness: prudential, rational, moral, aesthetic, etc.⁷ While there is agreement on the fact that such norms are epistemic in the former sense, there is no consensus on whether they are also epistemic in the latter sense. Philosophers such as Fantl and McGrath (2009a) and Gerken (2011) answer affirmatively to this question, whereas others, such as Brown (2008a, 2008b), remain neutral on this issue, characterizing epistemic norms exclusively in the former

p. 125). I note also that, though the upshot of my argument is negative, one could eventually draw positive conclusions from it about which non-epistemic conditions can rationalize a practical reasoning.

⁶ It is worth mentioning here that the specific focus of this chapter is on epistemic norms of action. The chapter does not address further issues concerning norms of belief. See Benton (2014, section 3) for an overview of recent discussions on norms of belief and for relevant references. The aim of the present chapter is not to demote knowledge (or other epistemic attitudes), but rather simply to argue—against RKP and other epistemic norms of action—that believing (and thus knowing) that p is neither necessary, nor sufficient to reasonably take p as a reason for action.

⁷ In this sense, epistemic norms would depend on a genuinely epistemic normative source. For a discussion of different sources of normativity see, for example, Broome (2013, pp. 26–27, chapter 7).

sense. Still others, such as Hawthorne (2004) and Hawthorne and Stanley (2008), think that epistemic norms of practical reasoning are standards of rational appropriateness broadly conceived, not strictly epistemic. While my objections to NEC will be effective against both understandings of epistemic norms, my objection to SUFF will be specifically directed to views which do not conceive epistemic norms as standards of strictly epistemic appropriateness, such as the views of Hawthorne and Stanley.

3.2 Counterexample one: rational action based on acceptance

For a long time, philosophers of mind used to explain action within a belief-desire framework. According to this model, when we act we seek to realize our intentions and satisfy our desires in the light of what we believe. Similarly, in our practical reasoning we would reason from desires, beliefs and intentions to action. However, this philosophical orthodoxy has been called into question: some philosophers have argued that other attitudes can motivate action and figure as premises in practical reasoning. A mental attitude often discussed in the literature that plays an important role in our practical reasoning is acceptance. In what follows, I will illustrate the close tie between acceptance and practical reasoning and how it poses a serious challenge to the knowledge norm and other doxastic norms of practical reasoning.

Before discussing the relation between acceptance and practical reasoning, it is necessary to clarify the notion of acceptance relevant for the present discussion and how it differs from that of belief. According to some stipulative notions of acceptance, belief is a kind of acceptance. For example, David Velleman equates accepting that p with regarding p as true. Since believing necessarily involves regarding a proposition as true, it is a kind of acceptance. For Velleman, supposing, assuming, and propositional imagining are other kinds of acceptance (Velleman, 2000, pp. 249–250). Similarly, for Crispin Wright, there is acceptance in all cases where the agent acts in a way as if she believes that proposition. Also according to this notion, belief is a type of acceptance. Other attitudes that fall into the category of acceptance include acting on the assumption that p , taking

for granted that p and trusting that p for reasons that do not bear on considerations regarding the truth of p (Wright, 2004, pp. 177–180).

On the contrary, according to a narrower and more natural reading of acceptance that I use here, acceptance and belief are two different kinds of mental attitudes. Many have argued that acceptance in this narrower sense is an attitude widely adopted in our ordinary, religious, scientific and technological practices.⁸ Here is an example adapted from Bratman (1992, p. 5). I am in Rome on a June day and I am planning my journey to visit the city. I do not actually have a belief about whether it will rain or not, nor do I have sufficient reason to believe that it will not rain—e.g., the weather forecast for that day is not available and according to the records there have been some showers in June in past years. Nevertheless, in my present circumstances taking for granted that it will not rain simplifies my planning in a way that is useful. On the basis of that acceptance, I decided to leave the umbrella at my hotel. Below I will consider other examples of acceptance.

It has been argued that acceptance differs from belief in at least three respects.⁹ First, acceptance and belief differ from each other in terms of the requirements of rationality governing the two attitudes. Reasonable belief is peculiarly responsive to truth-conducive, epistemic factors; believing a proposition requires regarding it as true with the aim or

⁸ For instance, Alston (1996), Audi (2008), Rey (2007), Sperber (1996) and van Leeuwen (2014) all argue that religious attitudes are acceptance or acceptance-like rather than belief or belief-like. Cohen (1992), Maher (1990), Mosterín (2002), and van Fraassen (1980) have argued that it is reasonable for scientists to merely accept the content of their scientific theories but not believe them. Bratman (1992) and Cohen (1989) discuss ordinary examples of acceptance such as the one considered immediately below.

⁹ Philosophers who have defended the distinction between belief and acceptance include Alston (1996), Audi (2008), Bratman (1992), Buckareff (2004), Dub (2015), Engel (1998), Mosterín (2002), Rey (1988), Stalnaker (1984), Tuomela (2000), van Fraassen (1980), Velleman (2000). Other proposals in characterizing an acceptance-like mental state include de Sousa's 'assent' (1971), Dennett's 'opinion' (1978), Sperber's 'reflective belief' (1996, 1997) and Frankish's 'superbelief' (2004). There are important differences between these authors in the ways they draw the distinction between belief and acceptance, but the three essential differences identified below are common to most of them.

commitment of getting its truth-value right.¹⁰ On the contrary, acceptance doesn't involve commitments to the truth of the accepted proposition.¹¹ There are no rational requirements to accept a proposition only if it is true. Accepting a proposition only involves treating it *as if* it were true, regardless of whether it is true or not.¹² This doesn't mean that acceptances are not the object of rationality requirements and are not liable to criticisms. Standards for rational acceptance concern non-epistemic factors, such as instrumental, ethical and prudential considerations. Acceptances are assessed according to whether it is useful or convenient for the agent to accept *p* given her practical purposes, whether accepting *p* maximizes one's expected utility, and so on.¹³

¹⁰ For a defense of similar claims see for example, Engel (2013), Shah (2003), Steglich-Petersen (2006), Velleman (2000) and Wedgwood (2002, 2013).

¹¹ We use 'commitment' in at least two senses. On the one hand, commitment refers to an attitude of endorsement; on the other hand, it refers to a norm or a requirement. In what we may call the 'requirement' sense, a commitment is a requirement or a norm that an agent is committed to respect. Such a norm would take the following form: accept that *p* only if *p* (or only if there is evidence for *p*). This type of commitment is much discussed in the literature on the normativity of belief (cf. McHugh & Whiting (2014), Fassio (2015)). There are no requirements of this sort on acceptance in the narrow sense relevant here: acceptance doesn't involve a commitment to endorse a proposition only if its content is true, as belief does. If acceptance involves some commitment to the truth, it is in a different sense, which we may call 'endorsement': this is an endorsement or intention that the agent deliberately takes toward a proposition, making as if that proposition were true for practical purposes.

¹² For example, Vahid (2006, pp. 323–324) argues that while belief involves regarding *p* as true for its own sake (or for the sake of getting its truth value right), other attitudes involve regarding *p* as true for the sake of something else. For example, assuming involves regarding *p* as true for the sake of argument (i.e. in order to see what it entails), and imagining involves regarding *p* as true for motivational purposes. An analogous claim can be made for acceptance. In the sense used here, acceptance is regarding *p* as true for the sake of practical purposes.

¹³ For similar characterizations of acceptance, see Bratman (1992) and Cohen (1989). My examples in the text will focus on a specific practical functional role of acceptance, that of helping us to make our practical reasoning more economical and faster in some circumstances, allowing us to avoid the use of more complex believed propositions. I just note here that there can be other practical purposes making rational the use of acceptances as premises in reasoning. See the quoted references for other examples.

Second, while belief is context independent or context invariant, acceptance is context dependent. When we believe something, we regard it as true no matter what our practical situation. By contrast, what one accepts can vary from context to context depending on the particular practical demands of the situation. These demands can sometimes make it reasonable for an agent to accept a proposition in a given context, even though she would not reasonably accept the same proposition in another context. While in planning my journey for a visit to Rome I reasonably accept that it will not rain, if I were figuring out what odds I would put in a bet on the weather I would not rely on that acceptance (Bratman 1992, p. 5). Contextual dependence of acceptance explains why, while belief is subject to an ideal of agglomeration across contexts—one should be able to or aim to integrate one's various beliefs into one consistent and coherent larger view—one may accept certain things which do not cohere with her other beliefs, for mere practical reasons present in a specific context.

Third, acceptance and belief differ from each other with respect to voluntary control. Normally, believing a proposition is an involuntary mental state. It is a disposition I find myself having, a product of my automatic cognitive mechanism. Furthermore, one cannot form or revise a belief at will regardless of the evidence (or at least not in normal circumstances). By contrast, accepting a proposition is, or is generated by, a mental act involving voluntarily taking on a positive attitude toward a proposition depending on practical considerations.¹⁴

By bearing this in mind, we can find a series of cases in which it is reasonable for the subject to accept that *p* under practical pressure, and appropriate to treat *p* as a reason for action, in spite of not having good

¹⁴ Mosterín (2002, pp. 317–319) makes a similar point in terms of the ways of processing information involved in forming belief and acceptance. He observes that belief is typically tied to unconscious processing of information whereas acceptance is generated by conscious, explicit, linguistically articulated decision-driven processing of information. Dub (2015) has argued that delusions constitute pathological cases of acceptance that are formed involuntarily. If this is correct, acceptances are not always voluntarily formed and under the control of the will. Nonetheless, the possibility that acceptance can be controlled voluntarily still constitutes a genuine difference between acceptance and belief. My examples below will refer exclusively to voluntary cases of acceptance.

reasons to believe it (or even having good reasons to believe the contrary). All the situations in which it is rational to act on a proposition that is accepted but not believed constitute counterexamples to NEC and necessity versions of other doxastic norms.

A variant of Pascal's Wager provides us with a good example here. Considerations about how it might be beneficial to live as if God exists cannot ground beliefs about God's existence. Nonetheless, they are indeed good practical reasons for accepting that God exists.¹⁵ Suppose one rationally decides to wager for the existence of God purely on the basis of a calculation of expected utility. This person would thereby accept but not believe that God exists. She would then take the proposition that God exists as a premise in her practical reasoning and simply endorse the policies and assent (at least externally) to the doctrines of the Church. She doesn't need to make that calculation over and over again each time this proposition matters to her practical decisions; and she doesn't even need to treat the believed proposition "if God exists and I don't behave in accordance with God's doctrines, then I will receive severe punishment after I die" as a reason for action every time she engages in a relevant practical reasoning. Rather, in many circumstances she will simply treat the accepted proposition that God exists as a reason for her action. For example, she may be motivated by this acceptance to spend more time in the church and follow the precepts of religion.¹⁶ Moreover, though she

¹⁵ Notice here an important difference with respect to the original Pascal's Wager case. Pascal's God demands that we believe in him—mere acceptance is not sufficient. According to Pascal, the reason for going to church is to cause oneself to believe that God exists. In the present example I consider a case in which *S* doesn't take God to ask him to believe in His existence. Instead, *S*'s acceptance of God is merely motivated by how practically beneficial it would be to live as if God existed. It is also worth mentioning here that a rational agent convinced by this type of Pascalian-like reasoning should have some minimal degree of credence that God exists. This is because part of the Pascalian reasoning relies on attributing at least a small chance that God exists. This subjective chance allows accepting that God exists to maximize expected utility. This wouldn't be the case if the agent's credence that God exists were zero.

¹⁶ Maybe such actions will cause her to believe that God exists at some future time, but before that time, most of her actions will be based on acceptance, not belief.

merely accepts that God exists, it seems that it is not inappropriate or irrational for her to treat that accepted proposition as a reason in her practical reasoning.

Another common situation in which it is appropriate to act on mere acceptance comes from scientific practices. There are cases in which the scientists' actions are based on some background assumptions that they know to be false. For example, nobody in the scientific community believes in the validity and completeness of Newton's theory of motion. But because of its convenience for making calculations in certain contexts, it is warranted for scientists to use Newton's laws as premises in their reasoning, acting as if such laws were true, at least as long as the margins of error permit it. When used in practical reasoning, these laws can provide sufficiently precise predictions given specific practical purposes. This is compatible with scientists knowing that Newtonian laws are false.¹⁷

The point can be generalized to the majority of natural laws. It is widely accepted in the scientific community that no contemporary physical theory is actually true. Science is far from having reached conclusive results. However, in practice scientists accept the available natural laws, using them as premises in at least some of their reasoning in order to calculate, design experiments and so on. Accepting natural laws in order to use them in one's reasoning is very convenient in specific circumstances: it helps in achieving reasonably accurate conclusions in a simpler and faster way, even though the scientist is well aware that these laws are false.¹⁸

A specific example could be useful here. A scientist, Mary, must deliberate about which specific act of computation she should perform in order to calculate the amount of fuel needed to get to the moon and back in a lunar module. Mary needs to calculate the amount of fuel quickly. She doesn't have time to use General Relativity, which (let's say) she actually knows to be the true theory. She can calculate the amount of fuel more quickly by using Newton's laws, which Mary believes to be false but a good approximation to the truth for her present purposes. While Mary

¹⁷ Thanks to Jonas Christensen for suggesting this case to me.

¹⁸ As a matter of fact, the use of acceptance is quite widespread in many scientific practices. See also J. Cohen (1992, p. 88).

could well use as a premise in her reasoning something she knows—e.g., the complex proposition that [$F=ma$ is the Newton's law necessary for calculating the needed amount of fuel, and $F=ma$, though false, provides a good approximation given her present practical purposes]—we can well conceive circumstances in which Mary does not use this complex proposition as a premise in her reasoning, but rather reasons as follows:

- 1) I must calculate the vector sum of the force of O.
- 2) The vector sum of the force of an object is equal to the mass multiplied by its acceleration.
- 3) Therefore, I shall multiply the mass of O by the acceleration of O.

The use of 2 in Mary's reasoning instead of more complex propositions doesn't look rationally impermissible. Reasoning directly from the accepted Newton's law helps her in achieving the desired results in a faster and simpler way. It seems perfectly natural and reasonable to reason like this in similar circumstances. Indeed nobody would challenge reasoning 1–3 as inappropriate or rationally impermissible, and if Mary were asked why she drew conclusion 3, she could well cite in her defense the accepted proposition 2 instead of other known propositions.¹⁹ This example seems to be a quite realistic representation of how many scientists engaging in practical reasoning use as premises in their reasoning some accepted proposition that they believe to be false when the desired results must not

¹⁹ I am not denying here that the belief that 2 is a good approximation to a desired result plays a certain indirect role in the overall explanation, motivating and making reasonable for Mary to endorse the acceptance and use it as a premise in her reasoning. Nevertheless, in the described case Mary doesn't use the complex belief as a premise in her reasoning, but the acceptance (i.e. Newton's law). The knowledge norm is still compromised, for this norm concerns directly the attitudes that one is rational to use as premises in reasoning, not the motivations of the subject to endorse certain types of attitude and use them as premises in reasoning.

be overly precise.^{20, 21}

If one were not convinced by Mary's case, here is another more familiar example from our philosophical practice. When we calculate the subjective probability of some propositions by updating evidence using Bayesian conditionalization, we know (or at least believe) that there are more precise rules for updating evidence (e.g., Jeffrey conditionalization), and thus that the proposition expressing Bayesian conditionalization is literally false. Nevertheless, if we are in contexts in which we are not concerned with a high level of accuracy (for example, if we are trying to solve basic exercises in a Decision Theory course), we make as if the

²⁰ Two things are worth remarking here: the first is that acceptance of natural laws known to be false in one's reasoning is not something specific to scientific practice. Laws are directly used as premises in reasoning by, for example, engineers and teachers. The second remark is that in Mary's example I focused on a case in which it is urgent to make a decision. This should provide a further reason for Mary not to engage in complex reasoning and instead rely on acceptance. Other factors may influence the preference for acceptance over more complex beliefs. Another is, for example, the presence of multiple consecutive deliberations involving common premises—consider a variant of Mary's case in which she has to repeat the calculation several times; in such case it is simpler to rely on the accepted Newton's law than to repeatedly rely on complex beliefs and more elaborated inferential patterns.

²¹ An anonymous referee for *Synthese* considers the possibility of a pragmatic explanation of the case: while in conversation Mary can express her reasoning as in 1–3, that may just be shorthand for a more complex reasoning understood in the conversational context, involving only beliefs as premises. Notice however that as I described the case, there is no conversational context in which Mary talks about her inference. Rather, the case involves a genuine inferential transition from premises 1 and 2 to conclusion 3. This excludes the possibility of explaining the case in terms of conversational implicatures. Another reason to think that a pragmatic account of Mary's reasoning is implausible is the following. To someone who criticises Mary for relying on a false premise it seems to make perfect sense to answer by saying that, of course, Newton's law is false, but making as if it were true and relying on it in her reasoning makes things simpler and reaches a conclusion whose accuracy is sufficient for present practical purposes. This possible answer seems perfectly fine in this context, but it is incompatible with a pragmatic account, according to which the only correct answer to the challenge should be the resolution of a conversational implicature (for example: "Of course, I was not speaking literally. What I really meant was..."). The fact that the former answer seems perfectly appropriate shows that a pragmatic account about this and similar cases is inadequate.

proposition expressing Bayesian conditionalization was true. When we reason from this proposition in such contexts, we deliberately overlook the fact that it is inaccurate and we move automatically to a conclusion, as we would do in reasoning from a belief. This way of reasoning from accepted propositions that we believe to be false (or at least we would hardly say we believe) in our philosophical practice seems to me both common and perfectly rational.

In both the exemplified cases of the Pascalian wager and the scientific practice, it seems perfectly rational for an agent to treat a proposition that is accepted but not believed as a reason for action. These cases constitute counterexamples to NEC and necessity versions of other doxastic norms.

From the case of scientific practice, we can also develop a counterexample to certain versions of SUFF. As I said in section 3.1, while some philosophers interpret epistemic norms as concerning a specifically epistemic sense of appropriateness, others conceive appropriateness in a more liberal sense (for example as substantive rational permissibility). The example I will consider below is problematic for all those endorsing the latter interpretation of SUFF (e.g. Hawthorne & Stanley 2008). A counterexample to this version of SUFF is one in which *S* knows that *p*, but given the setting of *S*'s situation, it is not rationally permissible for *S* to treat *p* as a reason for acting. Consider again the above scenario in which Mary knows the true and complex physical law of General Relativity necessary to calculate the precise amount of fuel needed for a lunar module to get to the moon and back. We can imagine a similar situation in which it is not rationally permissible for her to take that known proposition as a premise in her reasoning. Suppose again that Mary needs to calculate the amount of fuel only to a rough approximation, but it is particularly urgent that she does that in a very short time (e.g., she has only twenty seconds to enter an estimation of the amount of fuel into the control system of a machine).²² In such a situation, since adopting Newtonian laws would perfectly suffice for the purpose, it would be

²² Parameters including urgency that constitute deliberative contexts are discussed in Gerken (2011, 2017). One might worry here that though the agent in those cases seems to be blameless and fully excusable, he/she does violate some epistemic norm. I will address this possible worry in section 3.4.

unreasonable for her to use the complex law in her calculations. This is incompatible with SUFF. Note that Mary's epistemic position with respect to the proposition expressing the true physical law is also strong enough to satisfy the constraints required by other epistemic norms (justified belief that p , warranted belief that p to a degree that is adequate relative to deliberative context, or justified belief that S knows that p). And in the described situation it is inappropriate for her to use that proposition as a premise in her practical reasoning no matter how strong her epistemic position is with respect to that proposition. Therefore, all the sufficient versions of other epistemic norms, if interpreted in the liberal sense considered above, are confronted with a problem in dealing with this type of case as well.²³

A possible worry here could be that since acceptance doesn't involve rational commitments to the truth of the accepted proposition, acceptances are not liable to rational criticism and thus cannot serve as rationalizers of an action or a deliberation when used as premises in practical reasoning. However, I think that this worry is misplaced. As I said above, acceptances can be rationally assessed and criticised according to practical standards, and agents using acceptances as premises in practical reasoning are liable to rational criticism according to these standards. For example, if one concludes that accepting that God exists has the best expected utility, but then accepts that God doesn't exist and uses this acceptance as a premise in her reasoning, her acceptance can be assessed as unwarranted, and premising this proposition in her reasoning is liable to criticism. Similarly, consider the case of an engineer who accepts some law of Newton's theory in a context in which this theory doesn't provide sufficiently precise predictions for her specific practical purposes (e.g., for designing a particle collider machine), and she applies such laws as premises in her reasoning in that context (e.g., for making calculations whose results are necessary to design hadron accelerators). The

²³ Similar cases against SUFF can be made involving other attitudes. For example, suppose that Karen knows the axioms of number theory, and that Meera, who is reliable with mathematical knowledge, said that p is a theorem (although it is false). It becomes an urgent question for Karen whether p is a theorem. It is unreasonable for her to start reasoning from the axioms—she should instead rely on Meera's testimony.

reasoning of this engineer is liable to criticism and her acceptance cannot rationalize her action. This is because that acceptance is unreasonable given the practical purposes of the engineer in that context.

3.3 Counterexample two: rational action performed by the sceptic

Consider the following dialogue in which a sceptic (*K*) is trying to convince her friend (*F*) that she doesn't know that there is an external world.

K: "Do you know that you are not a brain in a vat?"

F: "No, I don't know."

K: "If you are a brain in a vat, then you cannot hold this cup of coffee in your hand, because you don't have hands at all. So, since you don't know that you are not a brain in a vat, you don't know that you have a cup of coffee in your hand."

After a brief thought, *F* concludes: "Yes, you are right, I don't know that." Thus she suspends her judgment. At the same time, *F* moves her cup to her lips and drinks the coffee.

I assume that: 1) *F* takes the conversation seriously and answers *K* sincerely, i.e., *F* does not give that answer to *K* due to any non-epistemic considerations, such as social graces and reluctance to displease her friend; rather, *F* is truly convinced by *K*'s reasoning and suspends her judgment as a consequence of that reasoning; 2) in moving the cup *F* genuinely exercises her agency; 3) the proposition that there is a cup of coffee in *F*'s hand (hereafter, *H*) is one of the reasons motivating her action; 4) The proposition that *F* treats as a reason for lifting her hand is *H*, not some complex proposition such as that, whether or not she is a BIV, seemingly lifting her hand will cause a pleasant taste and sensation of warmth.

Now, intuitively, it seems that *F*'s action cannot be criticised as irrational.²⁴ According to NEC, if it is appropriate for *F* to treat *H* as a

²⁴ Here I don't want to argue that before this conversation *F* didn't know that there

reason for acting, then *F* knows *H*. But *F* voluntarily suspended her belief that *H*. *F* doesn't believe *H* anymore, and consequently does not know that *H*. This is a counterexample to NEC.

It might be argued that in the above example there are other beliefs (plausibly amounting to knowledge) which *F* may be using as premises in her practical reasoning—in particular the belief that there *might* be a cup of coffee in front of *F*. Given the low cost of the action, this belief seems sufficient to rationalize it. This objection can be addressed by considering other analogous cases in which the sceptic doesn't merely suspend her judgment on the relevant proposition but believes that proposition to be false. Consider domain-relative forms of scepticism, such as scepticism about the existence of objects in the domains of mathematics, modality, ethics, etc. Such forms of scepticism are compatible with fictionalism with respect to each of these domains. For example, a fictionalist sceptic about mathematics holds that we should not believe in the existence of mathematical objects and we should regard sentences about mathematical objects, not as aiming at literal truth, but as telling part of a fictional story. For this sceptic, even if mathematical sentences are all false, engaging in a discourse about mathematics is rational because of its utility (Leng, 2015). Similarly, the fictionalist can rationally act as if those sentences were true. Lacking beliefs about mathematical propositions doesn't make the use of these propositions as premises in one's reasoning irrational or unreasonable. A fictionalist mathematician can perfectly well use sentences such as ' $7 + 5 = 12$ ' or 'There are no square prime numbers' as premises in her reasoning (both theoretical and practical); she can continue doing mathematics by adopting non-doxastic attitudes towards ordinary mathematical propositions.²⁵ In such cases, a sceptic about a certain area of discourse overtly believes that a proposition is false (and

was a cup of coffee in her hand. In fact *F* may have known that proposition before the conversation, but may have suspended her belief in that proposition as a consequence of the considerations proposed by *K*.

²⁵ It has been widely argued that people who refuse to accept the truths of mathematics can still continue doing mathematics by having an attitude towards mathematical objects sometimes referred as make-believe (also acceptance or exploitation) (see e.g. Daly (2008), van Fraassen (1980), and Yablo 2006)).

thus that it might not be true), but relies on it in her reasoning as if it were true. One can see the sceptic about the external world in my previous example as relying on an analogous sort of attitude in her practical reasoning.

One may object that *F* still involuntarily believes that *H* even though she would not be voluntarily willing to assent to it (Pritchard, 2000, p. 203). If this were the case, according to some externalist account of knowledge, *F* would know that *H* as long as some external condition is satisfied—for instance, if *F*'s belief that *H* was reliably formed. Furthermore, Williamson (2000) convincingly argues that knowledge is not a luminous mental state: one is not always in a position to reflectively know that she knows something. In the light of the above observations, one could say that what *F* lacks in the scenario is the higher-order knowledge that she knows that *p*, but not the first-order knowledge that *p*.

A possible way to defend the step from *F*'s claimed suspension of judgment about *H* to *F*'s ignorance of that proposition is by assuming that, even if in general knowledge may not be transparent to a subject, there are possible situations in which a subject in a case like that of *F* has an appropriate access to her own epistemic states. The possible failure of transparency in some cases does not entail that one is always wrong about her first-order mental state.²⁶ In particular, it is possible to conceive a scenario like the one described above such that, when *F* sincerely asserts that she doesn't know that *H*, she *has* a full epistemic access to the fact that she withholds her belief that *H*. In such a possible case, *F* would thereby not know that *H*. Even one single possible case like this is sufficient to provide a counterexample to NEC.

A more powerful objection is the following. There are possible ways of interpreting what's going on in *F*'s psychology in the above case, some of which are incompatible with the case being a counterexample to NEC. For example, Egan (2008) has recently argued for the view that the systems of belief that we in fact have are fragmented and could include subsets of beliefs which are possibly inconsistent. This view is opposed to idealized models of human cognition according to which our beliefs would be

²⁶ Neta (2009) argues that it is possible to know that *p* even if one believes that one does not know that *p*. But Neta's point doesn't conflict with what is suggested here.

part of a single coherent system. In the above case, one can interpret *F*'s cognitive system as fragmented, her scepticism not affecting her belief that *H*. If so, *F* can rely on the belief that *H* in her action. So described, the case wouldn't constitute a counterexample to NEC.

I concede that it's a much debated question in philosophy of psychology how to interpret similar cases, and that a 'fragmented mind' hypothesis seems a possible way of interpreting what's going on in *F*'s psychology in the above example. I admit that there are several possible interpretations of the case, some of which are incompatible with the case being a counterexample to NEC. But remember here that in order to make my point, all that I need is that there be at least one psychologically possible description of the example (or similar examples) under which *F* doesn't believe that *H*. The possibility of a single case in which *F* doesn't believe that *H* would be already a counterexample to NEC. This is perfectly compatible with there being other descriptions of this and similar cases according to which the subject believes the relevant proposition (as in the 'fragmented mind' interpretation). However, in the present context I cannot settle the issue of whether a description of the case that suits my purposes is psychologically possible. For this reason, I will set this issue aside and for the sake of argument I will simply assume that a similar description is indeed possible. My conclusion about the present case will thus be merely conditional: assuming an interpretation of this (or some similar) case as one in which the subject doesn't believe the target proposition, the case constitutes a counterexample to NEC.²⁷

²⁷ What could be the mental attitude that *F* has towards *H* if it is not belief? There are two possible interpretations here. One is acceptance: *F* decides to adopt the working hypothesis that she is not massively deceived in order to simplify her thinking, and uses this hypothesis as a premise in her reasoning. Alternatively, we can conceive the example as one in which *F* relies on *H* in an unreflective way, without first explicitly performing an act of acceptance. In this case, interpreting the attitude as acceptance would probably not be very accurate. A wide literature on Pyrrhonism suggests that the relevant attitude in such a case would be appearance. Sextus Empiricus (1994) describes an appearance as an involuntary affection (*pathos*) of the sceptic, something she passively undergoes. Unlike belief, an appearance makes no claims regarding the truth-value of *p*. Appearances, unlike beliefs, do not aim at truth, in the sense that they are not attitudes directed at correctly representing real states of affairs. Accordingly, they cannot be questioned and criticised with regard to their truth-dimension. Rather,

3.4 The excuse manoeuvre

One could defend the knowledge norm by arguing that though the agent in those cases seems to be blameless, he/she does violate the knowledge norm. Hawthorne and Stanley consider a situation in which someone in a situation of urgency is intuitively blameless in acting on mere partial belief. They claim that this kind of case doesn't ultimately threaten the knowledge norm; according to them, "the fact that we do not blame someone forced into a quick decision is no evidence at all against it." (Hawthorne and Stanley 2008, p. 587) In their view, the agent is blameless because the practical circumstances excuse her for violating the knowledge norm. Similarly, Hawthorne and Stanley could object to the above counterexamples by saying that the agent violates the norm but is excusable, for in these situations practical considerations render it excusable to act on less than knowledge (On the excuse manoeuvre, see also section 2.4).

Here are two replies. First, in the exemplified cases it is hard to see in what sense the agent needs to be excused. Our intuition suggests that the agent does not violate any epistemic constraint on practical rationality. For instance, when scientists have good reasons to use an out-of-date Newtonian law as a premise in their calculation, we neither judge them as acting inappropriately in any sense nor do we feel them in need of excuse for some wrongdoing. In these cases, there is no indication of the violation of some normative standard, either practical or epistemic: no criticisability, no blameability or excusability according to any normative assessment whatsoever. My reply here is similar to the one that I gave to a similar objection to my criticism of Williamson's approach in chapter 2.

Second, the manoeuvre of appealing to excuses is rather unpromising. Gerken (2011) points out that unless upholders of the knowledge norm can specify the notions of excuse and/or blamelessness, an appeal to excuses

they are appropriately assessed with regard to promoting a life free from turmoil and favouring the achievement of imperturbability. While beliefs involve a commitment to the truth of what is believed, appearances are attitudes supposed to represent the practical aim of acquiring peace of mind.

would be *ad hoc* and thus unconvincing. He then critically considers several possible principled accounts of excuses and argues that none of them is free from serious problems. It seems even more implausible to work out a viable account of excuse able to accommodate the types of cases considered above, not to mention that each of them involves completely different sorts of circumstances.²⁸

3.5 Concluding remarks

What conditions make it appropriate to treat *p* as a reason for action? In this chapter, we have seen that neither knowing that *p* nor believing that *p* is a necessary or sufficient condition for appropriately treating *p* as a reason for action. Notice however that the aim of this chapter was not to criticise RKP and other doxastic norms of practical reasoning on the ground that knowledge or warranted belief don't play any role in the rationalization of actions. Rather, it was to point out the limits of these principles by showing how rational actions may be based in some cases on other mental attitudes such as acceptance.²⁹

The conclusion of this chapter is that the normative connection between knowledge and practical reasoning is not tight enough to justify the knowledge norm of practical reasoning. If so, this undermines one main motivation for pragmatic encroachment and alleviates related objections or challenges to moderate invariantism (see section 1.3). Another important type of argument for pragmatic encroachment, based on intuitions about shifting patterns of knowledge ascriptions, still needs to be addressed. In chapter 1, we have critically considered pragmatics approaches to that challenge. In the next chapter, I shall explore another type of approach that I consider much more promising, the doxastic approach.

²⁸ Note that it seems to be even more implausible for upholders of the doxastic norms to use the excuse manoeuvre for defending their views. Those philosophers reject the excuse manoeuvre made by proponents of the knowledge norm. That would require developing a notion of excuse that could handle the cases of acceptance and appearance but maintain that there is no need of excuse in cases like the urgency situation considered by Hawthorne and Stanley (2008).

²⁹ A significant part of this chapter is adapted from Gao (2017).

Doxastic Accounts

In section 1.5, I have mentioned two prominent kinds of moderate invariantist accounts of practical factor effects on knowledge ascriptions: doxastic accounts and pragmatic accounts. There I provided a quite detailed explanation of the latter accounts and reviewed their most important problems. I only provided a sketchy outline of doxastic accounts, anticipating that they would be discussed in later chapters. It is now time to consider such accounts in more detail. In this chapter, I introduce the basic tenets of doxastic accounts and outline versions of such accounts discussed in the literature. In the following chapters, I shall defend a specific type of doxastic account.

According to doxastic accounts, the effects of practical factors on knowledge ascriptions discussed in section 1.1 can be explained by the causal-psychological influences exercised by these factors on our beliefs or degrees of confidence. For example, in the Bank case, the HS-subject fails to know that the bank will be open tomorrow because the perceived high stakes psychologically affect the HS-subject's belief or confidence, by causing her to reconsider her previous evidence and eventually doubt her previous convictions. Since the lack of knowledge is explained by a change in the doxastic condition necessary for knowledge, rather than by a change in the normative conditions necessary for knowledge such as justification and reliability, doxastic accounts are compatible with purism about knowledge.

Doxastic accounts of the Bank case and similar pairs of cases presuppose a particular view about belief, which I call *doxastic pragmatism*. Doxastic pragmatism holds that certain perceived practical factors generate systematic effects on normal subjects' outright belief. More specifically, this view predicts that other things being equal, a change in perceived practical factors tends to result in a shift of one's belief status

in a certain proposition. By ‘perceived practical factors’, I mean features of the environment or the circumstance that the subject perceives or is influenced by. It is possible to have cases in which certain perceived practical factors have misleading appearances; for example, when one falsely believes that one is in a high-stakes situation.

Doxastic pragmatism is not committed to any specific account of the mechanisms by means of which perceived practical factors generate effects on outright beliefs. The view itself only concerns the causal correlation between perceived practical factors and those effects. Nonetheless, doxastic pragmatists have so far provided different stories of the relevant mechanisms: in terms of dispositions constitutive of belief, psychological forces generated by the factors and affecting doxastic states’ regulation, and so on. Weatherson (2005) and Ganson (2008) appeal to functionalist accounts of belief to motivate the view (though they have different understandings of mental functionalism). Nagel (2008, 2010a) defends another version of doxastic pragmatism which focuses on psychological, rather than metaphysical, aspects of belief. Bach (2005, 2008, 2010) holds that high stakes affect confident, doubt-free belief (the type of belief required for knowing) turning it into a mere belief which is compatible with doubt.

In what follows, I will critically discuss these different versions of doxastic pragmatism in more detail. In sections 4.1 and 4.2, I consider Weatherson’s and Ganson’s accounts respectively. In section 4.3, I argue that Weatherson’s and Ganson’s functionalist/dispositionalist accounts of belief are affected by problems similar to those affecting simple conditional accounts of dispositions. In sections 4.4 and 4.5, I introduce Bach’s and Nagel’s versions of doxastic pragmatism. In section 4.6, I draw on Rose and Schaffer’s recent empirical findings to show that, contrary to what Bach and Nagel hold, knowledge doesn’t always entail confident belief or closed-mindedness. In section 4.7, I sum up the upshots of this chapter.

4.1 Weatherson’s doxastic pragmatism

Weatherson has been one of the first philosophers who endorsed doxastic

pragmatism. More recently, he has withdrawn his previous views on this matter, and he now argues for a pure form of pragmatic encroachment on knowledge (see Weatherson (2011, 2012)). I shall here refer to the views discussed in the paper “Can we do without pragmatic encroachment?” (2005), where Weatherson sets up his doxastic pragmatist account.¹

Weatherson develops his argument for doxastic pragmatism from the functionalist idea that to believe that p is to treat p as true for the purposes of practical reasoning. Weatherson spells out the functional role of belief in terms of a preference order of actions.² More specifically, believing p implies that your preferences make sense, by your own lights, in a world where p is true (p. 421). Suppose that I have a certain order of preferences over things that matter. If I believe that p , then conditionalising on p should not change the order of my preferences over those things; otherwise I do not believe that p , since that belief does not fit with my overall disposition to act. Weatherson provides a formal statement of this theory of belief as follows, where $Bel(p)$ means that the agent believes that p and A and B are two options (e.g., going to the sea or going to the mountain; checking whether the bank is open or going straight home...), $A _q B$ means that the agent thinks that A is at least as good as B given q :

$$(BEL) Bel(p) \longleftrightarrow \forall A \forall B \forall q (A _q B \longleftrightarrow A _{p \wedge q} B)$$

For Weatherson, the left-to-right direction seems trivial, and the right-to-left direction seems to be a plausible way to operationalise the idea that belief is a functional state, its function being that of providing grounds on which we can rely in our decisions.³

¹ In his paper “Knowledge, Bets and Interests” (2012), Weatherson abandons doxastic pragmatism given the difficulty of explaining some cases involving ignorance, or mistake, about the odds at which a bet is offered, and in the paper “Games, Beliefs and Credences” (2016), he gives up the simple Lockean reduction of belief to credence due to some cases derived from game theory in which having the maximal credence doesn’t lead to a full belief.

² Note that this is a quite unusual form of doxastic functionalism. I will come back to this point later in section 4.4.

³ As we will see in the next section, Ganson has a different opinion about this point.

According to Weatherson, this account should be further supplemented by additional constraints. These constraints are particularly important for our present discussion because they are responsible for the sensitivity of belief to practical factors. First, according to an unrestricted interpretation of (BEL), we barely believe anything about the truth of contingent facts. We can envisage that some bet has been offered to me about a contingent proposition p that when I wave at a statue in front of me, it doesn't wave back at me due to random quantum effects. If I take the bet and p is true, then I win a penny. If I take the bet and p is false, then I will go to hell and endure torture forever. Let's assume that q is a tautology (so that my preferences given q are my preferences *tout court*), A is that I decline the bet, B is that I take the bet. Given q , I prefer A to B , but given p and q , I prefer B to A .⁴ According to (BEL), I do not believe that p , which is very counter-intuitive (*ibid.*).

To avoid this problem, Weatherson suggests introducing some restrictions on the quantifiers in (BEL):

- 1) Only pragmatically *relevant* propositions play a role in influencing our belief status. A proposition is relevant if the agent is disposed to take seriously the question of whether it is true. This condition rules out wild hypotheses that the agent does not take at all seriously. Moreover, conditionalizing on a relevant proposition or its negation should bring about changes in some of the agent's unconditional preferences over live, salient options, such that if A and B are live salient options, $A \quad_q B \leftrightarrow A \quad B$ is false. A proposition is salient if the agent is currently considering whether it is true.
- 2) Proposition q should be restricted to *active* propositions relative to p . A proposition is active relative to p iff it is a

For one thing, Ganson doesn't consider the left-to-right direction of (BEL) trivial. On the contrary, she gives serious thought to whether we should provide a context-sensitive or context-insensitive reading of the condition. For another thing, she seems to take both directions as plausible ways to operationalise the functionalist idea of belief.

⁴ This case has been modified based on one considered by Weatherson (2005, p. 422).

(possibly degenerate) conjunction of propositions such that each conjunct is either relevant or salient, and such that the conjunction is consistent with p . This means that if q is not a conjunction, q has to be either relevant or salient. If q is a conjunction, each conjunct has to be either relevant or salient. 'Relevant' has been defined above. A proposition q is *salient* if the agent is currently considering whether it is true (p. 423).

3) Actions A and B should be *live* options in the sense that they are really possible for the agent to perform.⁵

4) A and B should be *salient* options in the sense that the agent takes those actions seriously in her deliberation.

Note that a live option may not be salient and vice versa. For example, even though it is open to me to gamble large amounts of money on internet gambling sites (live option), I never consider it relevant in making a decision (hence not a salient option), and properly so. By contrast, there are cases of salient options that are not live. For example, entertaining the thought of spending the winter in Hainan would please me, so it is a salient option to me at the moment. Still, it is not a live option given my other duty to attend events at my department. So, for Weatherson, *liveliness* and *salience* are two different constraints on quantifiers ranging over options of the agent.⁶

With these qualifications, Weatherson's view is that the agent believes that p iff conditioning on p changes none of her actual preferences over live and salient options, where the conditions are also active relative to p . Note that these restrictions of the quantifiers in (BEL) (in particular the liveliness and salience of options) depend on the subject's interests.

⁵ The relevant possibility here is stricter than the metaphysical modality. A live option is one that the agent can actually perform given all actual situational and psychological constraints of the agent.

⁶ Note that the quantifiers range over options for the agent who performs the action, not the person making the belief ascription. So his view is not a contextualist account of belief. In addition, what counts as a salient option doesn't involve any normative consideration about what the agent should take seriously. Weatherson is interested in providing a descriptive picture of what the agent does believe, not what they should believe (2005, p. 423).

Accordingly, (BEL) has the consequence that what an agent believes is in general sensitive to dispositions about practical matters, but insensitive to abstruse dispositions in situations far removed from actuality (Weatherson, 2012, p. 89). Weatherson takes this to exemplify the interest-sensitivity of belief. More specifically, this interest-sensitivity, according to Weatherson, consists of one being disposed to lose a belief if the circumstances changed to the extent that conditioning on the relevant proposition would alter some of one's order of preferences.

Now we can consider how Weatherson's account of belief explains paradigmatic cases used to motivate pragmatic encroachment. According to Weatherson, the LS-subject and the HS-subject are justified in adopting the very same degrees of belief. For example, they are equally justified in assigning a probability of around 0.9 to the target proposition p . Conditionalising on p doesn't change any of the preferences of the LS-subject over open and salient options. By contrast, conditional on p , the HS-subject prefers to choose risky opinions accompanied by possibly high costs (e.g., not taking the queue but coming back on Saturday). This choice is different from what HS-subject's actual order of preference which is queueing now instead of coming back on Saturday. Thus according to (BEL), the LS-subject believes that p , but the HS-subject does not. In this way, Weatherson explains the intuitive shift of knowledge ascriptions in bank cases in terms of the belief condition entailed by knowledge.

4.2 Ganson's doxastic pragmatism

Ganson (2008, 2019) also suggests that pragmatic considerations are relevant in determining the threshold for outright believing. In her view, two subjects who are equally justified in having an identical degree of credence in p can differ concerning whether they believe that p . The degree of credence might be high enough for one to be willing to act as if p but not for the other. Therefore, her view is also a *pragmatist threshold view* as Weatherson's: both views hold that whether one has enough confidence in p to believe it is sensitive to practical circumstances.

Let's look at Ganson's view in more detail. Ganson suggests the following constraint on belief:

(C) *S* outright believes that *p* only if *S* believes that *p* to a degree which is high enough to ensure that one is willing to act as if *p* is true—where one's being willing to act as if *p* means that what one is in fact willing to do is the same as what one would be willing to do, given *p*. (Ganson, 2008, p. 451)

As a necessary constraint on outright believing, (C) states that in order to count as outright believing that *p* in the circumstances, an agent must believe that *p* to a high enough degree such that she is willing to act as if *p* in the circumstances (*ibid.*, p. 453).⁷

How should we interpret this constraint? According to a context-insensitive interpretation, if *S* outright believes *p*, then *S* believes *p* to a degree which is high enough to ensure that one is willing to act as if *p* in *all circumstances*. But this condition is too demanding to be true. We would barely rely on any belief in whatever settings, in particular when the costs of acting as if *p* and being wrong are very high, and the benefits if we are right are very low.

What about a weaker context-insensitive reading? Ganson considers the following reading: a believer possesses a willingness to act as if *p* in *most (rather than all) circumstances* where the evidence for *p* is unchanged. However, there are some difficulties with such an interpretation. First, it is not clear how to understand 'most'. On the one hand, interpretations such as 'most relative to all logically possible scenarios' would be too demanding. On the other hand, with interpretations such as 'most normal circumstances' or 'most of those circumstances sufficiently like the believer's current circumstances in the relevant respects', it is difficult to

⁷ Ganson (2019, section 6) accuses Ross and Schroeder (2014) of uncharitably interpreting the pragmatist views of belief in Ganson (2008) and Fantl and McGrath (2009). Ross and Schroeder interpret Ganson and Fantl and McGrath's views to be about 'credence high enough to rationalize acting as if *p*'—a normative quality, whereas the original views concern strictly psychological features. In addition, Ross and Schroeder reconstruct the pragmatist view as a bi-conditional account of belief according to which having sufficiently high confidence in *p* is sufficient for you to be motivated by *p* itself, and to base your actions, emotions and reactions on *p* itself. By contrast, Ganson (2019) holds that outright belief is not just a matter of having confidence level reaching a point that triggers all the belief-relevant propensities.

specify what circumstances count as such. In addition, context-insensitive interpretations would involve further complications with cases in which the subject possesses a proclivity to act as if *p* in many circumstances except in the actual one (ibid., p. 452).

Hence, in order to take (C) as providing a necessary condition for outright belief, Ganson suggests adopting a context-sensitive interpretation:

In order to count as believing *p* in a range of circumstances, one must be willing to act as if *p* in those circumstances: one's degree of belief that *p* has to be high enough that one is willing to act as if *p* *under those circumstances*. (ibid., p. 452, italics added)

Ganson (2019) further develops her version of doxastic pragmatism and explicitly adopts a dispositionalist understanding of belief. She suggests the following account of belief that combines dispositionalism about belief with aspects of a threshold view:

Hybrid Doxastic Pragmatism*

Bp requires that you are confident enough that *p*

- for you to be able to have the spectrum of dispositions typical for outright belief that *p* under these conditions
- for *p* itself (rather than something more hedged proposition such as that *p* is likely) to be able to serve as a motivating reason for you to have or engage in the relevant sorts of activities, actions, reactions, feelings, habits typical for belief that *p* here. (ibid., p. 23)

While the above account still holds that having a high confidence in *p* is essential for belief, it does not reduce belief to confidence above a threshold. According to this account, the acquisition of various sorts of dispositions in reasoning, acting and feeling play a crucial role in marking the transition from mere degrees of belief to outright belief.

According to the both Ganson's and Weatherston's proposals, in typical high-stakes cases, the subject lacks sufficient confidence to be able to be motivated/willing to act on the basis of *p* and to use *p* itself as a

motivating reason in her choice about what to do. But having such degree of confidence is required for belief that p . Thus, the HS-subject in the bank case doesn't believe that the bank is open tomorrow. By contrast, the LS-subject is motivated/willing to act as if the target proposition is true; so she believes that proposition.

4.3 Objections to Weatherson and Ganson's accounts

Before discussing my objections to the views of Weatherson and Ganson, a preliminary question we could ask is whether their accounts truly count as versions of functionalism. Recall that Weatherson and Ganson deploy very different terms in explicating their respective functionalist accounts of belief. Ganson defines belief in terms of being "motivated/willing to act as if p is true", whereas Weatherson appeals to preferences and levels of confidence. Ralph Wedgwood (2012) questions whether Weatherson's account really vindicates the functionalist spirit. He says:

The functional role of a type of belief is a matter of the dispositions that characterize this type of belief, and the notion of a disposition is a causal or explanatory notion. According to Weatherson's account, the explanation of why an agent has the preferences that she has does not appeal to the agent's disposition to use the propositions that she has outright beliefs in as "premises in her practical reasoning". Instead, this explanation appeals to the functional role of preferences and levels of confidence, and to the fact that the agent has levels of confidence both in the propositions that she has outright beliefs in and in other incompatible propositions as well. (Wedgwood, 2012, p. 318)

As Wedgwood points out, Weatherson's account does not consider standard dispositions concerning the role of belief in practical reasoning. But one may further question why such dispositions should be reflected in ordinary practical reasoning rather than in consistent preferences. After all, aren't preferences given confidence the proper grounds for

deliberation, and so even more fundamental than the process of reasoning?

Wedgwood's worry doesn't apply to Ganson's account, given that Ganson explicitly adopts the standard dispositional formulation of the functionalist account of belief (especially in Ganson (2019)). My following objections to Weatherson's and Ganson's accounts do not hinge on Wedgwood's point. Here I remain neutral on the issue of whether dispositions in preferences of actions should be included in the relevant type of dispositions characterizing a functionalist account.

My objections are in the form of counterexamples, and they parallel other well-known objections in the debate about the metaphysics of dispositions. I will argue that Weatherson's and Ganson's context-sensitive accounts of belief are affected by similar problems.

The simple conditional account of disposition says:

(D) An object has a disposition D iff it would M if it were the case that C at time T . (cf. Choi & Fara, 2014)

According to this account, dispositions are defined in terms of the nature of certain characteristic manifestations M under some stimulus or activating condition C . Recall that according to Ganson's context-sensitive account of belief concerning the necessity condition for belief, in order to count as believing p in a range of circumstances, one must be willing to act as if p *in those circumstances*. This account can be formulated as follows:

(DG) One believes that p at time t only if one would act as if p is true if p is relevant for one's choice at time t .

The consequence of (DG) characterises the disposition involved in belief as suggested by Ganson's account. The manifestation condition at issue is to be willing to act as if p is true. The stimulus condition can be identified with circumstances in which p is relevant for one's choice at the present time, where 'relevant' means the same as in Weatherson's account, that the preference order conditional on p or its negation is different from the unconditional preferences over live and salient options.

The spirit of Weatherson's account (BEL) can be captured by the following claim:

(DW) One believes that p at time t *if and only if* one's order of preferences over live and salient options at time t would not differ from one's order of preferences conditional on p (where the conditions are active relative to p).

In (DW), the stimulus condition is the conditionalization of one's order of preferences on p . The manifestation condition is that one's order of preferences over live and salient options at time t does not change after so conditionalizing.

There are three very influential counterexamples to (D) in the literature on the analysis of dispositions. I first introduce the original counterexamples to (D) and then show how we can have parallel counterexamples to (DG) and (DW).

First, consider *maskers or antidotes* cases. The most discussed case is one in which a fragile glass has been carefully protected by packing material (see Johnston (1992, p. 232) and Bird (1998, p. 228)). The glass's disposition to break when struck is intact but it is masked by the protecting layer.⁸ Another often cited case is the disposition of methanol to be oxidized to formaldehyde if accompanied by the enzyme alcohol dehydrogenase. This disposition can be masked by adding ethanol. The latter prevents the oxidation of methanol by consuming the enzyme in the production of a different chemical complex (acetaldehyde) (Ekins, 1985, pp. 337–340; cf. Proctor et al. 1998).

Here is a parallel case for belief in which the disposition to take p as a premise in practical reasoning or acting as if p is true is masked. Consider a scientist, Mary, who must deliberate about which specific act of computation she should perform in order to calculate the amount of fuel needed to get to the moon and back in a lunar module. Mary needs to calculate the amount of fuel quickly (e.g., she has only twenty seconds to enter an estimation of the amount of fuel into the control system of a machine). She doesn't have time to use General Relativity, which (let's say) she actually believes to be the true theory. So she calculates the amount

⁸ An exception is Choi (2008) whose view about dispositions implies that when the packing material is so pervasive that it is ruled out from the ordinary conditions for fragility, the glass is not fragile anymore.

of fuel more quickly by using Newton's laws, which she believes to be false but a good approximation to the truth for her present purpose. In this case, the proposition about the General Relativity theory is relevant to the current circumstances, but Mary doesn't use that proposition as a premise in her practical reasoning. The disposition to use that proposition as a premise in practical reasoning is not removed but masked by other practical concerns (in this case, by urgency). Again, according to Ganson's version of doxastic pragmatism, Mary doesn't believe in General Relativity. But this is plainly false.⁹

Let us consider a second type of counterexample to conditional accounts of dispositions: *finks and reverse-finks*. Suppose that an electrical wire is live just in case, if the wire were stroked by a conductor at t , the electrical current would flow from it to the conductor. Suppose that the device is operating on a reverse cycle, attaching to a naturally live wire but removing its property of being live if ever it is touched by a conductor. For example, the wire is connected to a safety switch (or a reverse electro-fink). The switch turns off the electrical current whenever the activating conditions of the disposition occur. In this case, the left-hand side of (D) would be true (an object would have disposition (D) at t) whereas its right-hand side false (the manifestation condition would not occur despite the stimulus condition being the case). As such cases illustrate, some dispositions are 'finkish' in the sense that the conditions for an object's acquiring or losing the disposition might be the same as the stimulus conditions (Martin, 1994).

We can construct a similar counterexample to (DG) where the disposition of acting on a certain belief is 'finked'. Here is one case. Imagine that an evil scientist has implanted a chip in your brain. The chip has the effect that every time the belief that $2 + 2 = 4$ becomes relevant to the current choices, it will inhibit you from deploying your belief that $2 + 2 = 4$ as a premise in practical reasoning, or act as if it is true. In such a case, the left-hand side of (DG) is true whereas its right-hand side is false. A similar counterexample also applies to (DW). In this case, we can imagine that the evil scientist always generates a change in the order of preferences over live and salient options conditional on p . Again, the left-hand side of (DW)

⁹ For more discussion of this case, see also section 3.2.

would be true whereas the right-hand side is false.

Weatherson's (BEL) account is also affected by a further problem. This is due to the non-monotonicity of reasons aggregation. In short, it is almost universally held that for some reasons p and q , p can be a reason to F , q can also be a reason to F , but p and q together constitute a reason not to F (e.g., Dancy, 2004, chapter 2; Horty, 2012). Suppose that p is the proposition that Mary will go to the party and q is the proposition that Katy will go to the party. You really like both Mary and Katy. The presence of each of them is a reason for you to go to the party. But Mary and Katy hate each other, and if both are present at the party it will be an awful night for all those present at the party. Now, you believe that p . However, where action G is 'go to the party' and action S is 'stay at home', it is obviously false that $G \text{ }_q S \leftrightarrow G \text{ }_{p \wedge q} S$. According to (BEL), you don't believe that p , which contradicts our initial assumption.¹⁰

The above counterexamples pose serious problems to (DG) and (DW). In my view, they constitute good reasons to reject these accounts. Notice, however, that the objections considered in this section are not against every possible dispositional account of belief. Rather, these objections constitute serious challenges only for accounts sharing the context sensitive manifestation conditions proposed by Ganson and Weatherson. Certain context-insensitive accounts may avoid the above considered problems by restricting the manifestation conditions of the dispositions relevant for belief to specific 'normal' circumstances in which there are no finks, reverse-finks and maskers preventing the manifestation of the relevant dispositions.

This leads us to consider whether there are any suitable more sophisticated versions of (DG) and (DW) which could avoid those difficulties. One typical move in the literature on dispositions used to avoid the above types of counterexamples is to hold that any specification of the stimulus condition at issue must include covert reference to a *ceteris paribus* clause or 'all else being equal'. After all, behavioural dispositions, or dispositions with a behavioural element, seem particularly defeasible since they can be overturned by other practical considerations. In this way, with respect to (DG), for example, it follows that one believes that

¹⁰ Thanks to Davide Fassio for pointing out this case to me.

p iff one would use p as a premise in one's practical reasoning if it were the case that p is relevant for one's choices *ceteris paribus*, under certain standard conditions. In this way, by classifying cases in each of the alleged counterexamples to (DG) and (DW) as nonstandard conditions, it delivers the right verdicts to each case.

Here there is a further question about how to spell the *ceteris paribus* clause. Serious doubts have been raised that the only way to spell it out is to render the proposed analysis vacuous (see Bird (1998), Fara (2005), Hauska (2008), Martin (1994), Mumford (1998, 2001)). Yet it has also been argued that the *ceteris paribus* laws and generalisations can still be useful and productive in most cases (see e.g., Choi (2008), Mellor (2000) and Steinberg (2010)). What is agreed is that if we can spell out the *ceteris paribus* clause in a way that does not render the resulting conditionals vacuous, then we should go for it. I will consider this question a gain in section 6.2, where I will propose a context-insensitive account of dispositional belief.¹¹

4.4 Bach's doxastic pragmatism

Bach (2005, 2008, 2010) provides a moderate invariantist account of the intuitive judgments about the bank and airport cases. As upholders of other doxastic accounts, Bach claims that the HS-subject in the bank case does not know that p because the subject does not meet the doxastic condition for knowing. However, rather than holding that the HS-subject simply does not believe the target proposition because the subject is not yet ready to act as if the proposition is true, Bach suggests that the subject does not have confident, doubt-free belief required for knowing (Bach 2008, pp. 83–84, 2010, p. 118).

According to Bach, knowing requires *confident, doubt-free* belief rather

¹¹ It is also worth mentioning that while a restriction of (DG) and (DW) to *ceteris paribus* circumstances would avoid the considered problems, the resulting accounts would be inadequate to express a general necessary condition for outright belief. The reason is that there would still be beliefs in *non-ceteris paribus* circumstances not considered by the necessary conditions of these accounts, such as the belief that $2+2=4$ in the evil scientist case above.

than a mere belief that is compatible with doubt or doubtfulness.¹² Now, in *High Stakes* the higher stakes increase the range of possibilities one should guard against. And it is clear that the HS-subject cannot exclude those further possibilities of error without acquiring further evidence. It is thus reasonable for the HS-subject to remain partially doubtful about the truth of the target proposition (say, p). Hence, the subject does not confidently, doubt-free believe that p .

Bach then explains the lack of knowledge of the HS-subject in terms of a lack of enough confidence in p . He says:

[Y]our practical interest may lead you to want to make sure that p before you act on the supposition that it holds true. As a result, you don't yet believe that p , at least not with confidence, and wish to guard against certain possibilities of error. This means that you don't yet know that p . The reason of this is not that you have an insufficient epistemic reason for believing that p but that you don't meet the doxastic condition on knowing. (Bach 2010, p. 118)

As for the airport cases, in which the stakes concern the attributor but not the subject, according to Bach we should bear in mind two things about knowledge ascription:

[T]hat attributing to someone knowledge that p involves (confidently) believing that p yourself; and that denying knowledge of someone who has the same evidence you have involves being at least somewhat doubtful about p . (Bach, 2005, p. 76)

Accordingly, Bach suggests that since the attributor is not confident that p and deems that she still needs to verify the truth of p , she cannot *coherently* ascribe knowledge to the subject in question, even though the subject confidently believes and knows that p . The attributor has to deny

¹² Note that the difference between confident belief and mere belief is not normative but psychological.

knowledge to the subject, since she thinks that the truth of p is still to be confirmed and the subject has no more relevant evidence than the attributor herself (see also Bach (2008, p. 84, 2010, p. 115)).

As Weatherston and Ganson, Bach commits to a mobile threshold for confident, doubt-free belief, which rises in relevant high-stakes cases. He says:

[T]he higher stakes raise the threshold of confident, doubt-free belief. [...] 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. (Bach, 2008, p. 83)

So, I am suggesting, willingness to attribute knowledge does not track the standards on the truth of a [knowledge] attribution; rather, it tracks one's *threshold of doxastic confidence*. In the so-called high-standards cases, the attributor's doxastic threshold goes up to the point that without additional evidence she implicitly, but mistakenly, thinks she is not in a position to know. (Bach, 2005, p. 78, italics added)

According to Bach, the HS-subject does not meet the doxastic condition for knowing due to a rise in the doxastic threshold of confident and doubt-free belief. A similar account also applies to knowledge ascription cases. If an attributor's confidence doesn't meet the relevant threshold and thus fails to meet the doxastic condition for knowledge, the attributor wouldn't attribute knowledge to herself and to others sharing the same amount of evidence. The above accounts together explain our intuitive judgments about relevant high-stakes cases.

4.5 Nagel's psychological account

While the three doxastic pragmatist approaches we discussed so far all focus on the metaphysics of outright belief, it is undeniable that the psychology of a (rational) agent is also highly pertinent to the present

debate. After all, it is an empirical claim that the threshold of credence for belief is sensitive to practical considerations. So, in principle, the claim can be confirmed or falsified by empirical evidence. In this respect, Nagel's work is particularly important. Nagel (2008, 2010a) brings into the discussion relevant psychological studies, and uses these data to work out a more informed doxastic pragmatist account.

Nagel identifies two related psychological forces important for regulating humans' allocation of cognitive efforts in various practical situations. One is *epistemic anxiety*, which serves to regulate our cognitive activities directed toward ascertaining the truth of practically relevant propositions (Nagel, 2010a, p. 408). This is a type of emotive response that normally affects cognitive efforts, such as our readiness to collect new evidence and our choice of an evidence-weighting strategy leading to the formation of a given opinion. Such a force enables us to use our cognitive resources in proportion to the expected costs and benefits. For example, by perceiving high stakes, the level of epistemic anxiety of the subject normally rises, which in turn results in a tendency to allocate more cognitive efforts and adopt a more accurate and complex evidence-weighting strategy. This natural expectation has been repeatedly confirmed by a number of psychological studies. As Nagel sums up:

In general, high-stakes subjects think more systematically and less heuristically, relying more on deliberate and controlled cognition and less on first impressions and automatic responses [Kunda 1990; Lerner and Tetlock 1999]. Many cognitive biases—a recent survey article on accountability counts sixteen—are known to be attenuated when subjects take themselves to be shifted into a higher-stakes condition [Lerner and Tetlock 1999]. (Nagel, 2008, p. 282)

Together with an elevated level of epistemic anxiety, according to Nagel, other things being equal, high stakes also induce a lowered *need-for-closure*. The current notion of closure is introduced by psychologist Arie Kruglanski and his colleague to refer to the transition from the hesitant conjecture to a subjectively firm and settled belief. Nagel summarises:

Achieving closure or judgemental commitment on a question puts an end to the experience of ambiguity and delivers the sense of having a firm answer. The opposite of closure is openness or judgemental non-commitment, in which we are able to continue juggling alternative possibilities, perhaps lingering in ambiguity or confusion. (Nagel, 2008, p. 286)

Although the two forces are closely related, epistemic anxiety and need-for-closure cannot be completely inter-defined. Epistemic anxiety only refers to the inclination for increased cognitive activity. This inclination can be overshadowed by factors such as haste and distraction (Nagel, 2010a, p. 414). By contrast, need-for-closure is the final consequence of an interaction of multiple practical and psychological factors—including but not restricted to epistemic anxiety (Nagel, 2008, pp. 287–288). This makes it possible to have cases where epistemic anxiety is high, whereas the need-for-closure could be either neutral or high, depending on the presence of other factors neutralising the effects of epistemic anxiety.

In paradigmatic high-stakes cases, time pressure and other factors that could neutralize the inclination to lower one's need-for-closure are absent, so high-stakes subjects experience a low need-for-closure as well as high epistemic anxiety. By contrast, in low-stakes cases, there is no particular practical or psychological factor that affects the level of need-for-closure. The low-stakes subject then tends to have a low level of epistemic anxiety and a neutral need-for-closure.

In drawing a connection between making up one's mind (closure) and the belief status required for knowledge, Nagel appeals to the above three doxastic pragmatist approaches in different places. For example, in an earlier paper (Nagel, 2008), where she uses the notion of need-for-closure to analyse the psychology of HS-subjects, she appeals to Bach's view:

If our intuitions about the shift in Stanley's cases are driven by the contrast between low and neutral need-for-closure, then something like the Bach objection will be right: we ascribe knowledge in Low Stakes and deny it in High because we naturally attribute higher and lower confidence belief to the contrasted subjects, or confident belief and a state of evidence

assessment that precedes fixed belief. (Nagel, 2008, p. 289)

In her later paper (Nagel, 2010a), instead of talking of ‘forming a firm belief’, Nagel sticks to a talk of ‘outright belief’ and uses the views of Weatherston (2005) and Ganson (2008) to support her view.

Summing up, Nagel defends the view that stakes affect our knowledge ascriptions in high-stakes cases by affecting our perception of whether in such cases the subject has met the doxastic condition for knowledge, i.e., whether she made up her mind on the relevant proposition or formed an outright belief.

The diagnosis of Nagel discussed so far doesn’t exclude the possibility that a subject could stick to an outright belief or a confident belief in a high-stakes situation (a belief which, by stipulation, is required for knowing). We could imagine a case in which the subject’s psychology is insensitive to the high cost of error and thus the subject continues believing the relevant proposition in high stakes. However, this wouldn’t alter our intuitive judgment that the HS-subject still doesn’t know the relevant proposition. This intuition seems to speak against the approach of Nagel and other doxastic pragmatists, and in favour of the view that practical factors affect the knowledge-level-justification rather than psychological factors such as belief.

In response, Nagel argues that a mere stipulation that the HS-subject has the same knowledge as his/her low-stakes counterpart is plausible only if the reliability of the HS-subject’s belief formation is perceived as not compromised by factors such as thinking hastily or being biased by wishful thinking (Nagel, 2008, p. 292). The point is that if the high-stakes subject forms an outright belief without further collecting any evidence in spite of high epistemic anxiety, it would be natural for us to question the reliability of her belief formation. She says:

To come across as knowing, subjects who are settling a given question—say, whether the bank will be open—need to collect a level of evidence that would satisfy the epistemic anxiety we naturally ascribe to them. Failure to do so would make it natural for us to see the subject as suffering from a reliability-compromising condition such as haste, distraction or wishful

thinking. Subjects come across as passing the reliability condition if they are seen as forming evidence-based beliefs, where more or less evidence may need to be collected in order to seem to satisfy that condition. (Nagel, 2010a, p. 419)

According to Nagel, since it is stipulated that the HS-subject forms a confident belief that the bank will be open on Saturday as her low-stakes counterpart, it is natural to perceive the evidence-gathering strategy at issue as defective, hence preventing one from knowing. Similar arguments also apply to the formation of the HS-subject's degrees of belief. According to Nagel, we could stipulate that the HS-subject has the same confidence level as the LS-subject based on the same moderate evidence. But the reader would naturally perceive that high-level confidence as unfounded and the accuracy of the HS-subject's judgment as compromised, given that subjective confidence ordinarily co-varies with the stakes (Nagel, 2008, pp. 291–292, 2010a, p. 422).

Based on Nagel's arguments, we can construct a dilemma for pragmatic encroachment. Either it is natural to perceive the HS-subject as less confident than the LS-subject, and hence as not believing the relevant proposition, or it is stipulated that HS-subject has the belief, but it is then natural to perceive her belief as inappropriately formed and the reliability of her belief formation as compromised. Either way, the lack of knowledge of the HS-subject is explained within a purist, moderate invariantist framework.

Nagel's strategy has been criticised by Sripada and Stanley (2012). According to Sripada and Stanley's interpretation of Nagel's argument, this argument is weak, since it doesn't apply to those cases where it is stipulated that the evidence-gathering strategies used by the subject are exactly the same in low and high stakes cases (*ibid.*, p. 22). By using vignettes with such settings, it is implausible for readers to perceive any difference between the evidence-gathering strategies deployed by the subjects in high- and low-stakes. According to a moderate invariantist, if a belief is formed on the basis of strategies that are good enough to know in a low-stakes situation, then the same belief formed on the same bases is also equally well supported in a high-stakes situation, and the HS-subject is in a position to know the believed proposition. So, from a moderate

invariantist perspective, if a LS-subject is in a position to know p based on evidence-gathering strategy S , a HS-subject must also be in a position to know p based on the same strategy.

In reply, Nagel could argue that the reliability condition doesn't merely consist of an evidence-gathering strategy. It also includes the evidence-weighting process that is subject to the interference of reliability-compromising factors. Indeed, the reliability-compromising factors suggested by Nagel, i.e. haste, distraction or wishful thinking, concern the evidence-processing or evidence-weighting procedure rather than the evidence-gathering strategy. Thus, merely holding fixed the evidence-gathering strategy across cases but leaving unspecified other important conditions relevant for the reliability of belief formation allows for a different intuitive verdict in low- and high-stakes cases, as suggested by Nagel.

Summing up, Nagel's doxastic account of the practical factor effects on knowledge ascriptions relies on psychological features of the HS-subject. In particular, she attributes specific psychological properties to outright belief and its regulation processes. A belief requires closure and the absence of epistemic anxiety. These features are absent in other versions of doxastic pragmatism we discussed previously. Although Nagel appeals to a core tenet of doxastic pragmatism, i.e., the pragmatist threshold view, to support her diagnosis, this doesn't constitute an essential part of her account. Her account only requires establishing a correlation between high perceived epistemic anxiety and perceived difficulty in attaining outright belief.¹³ In addition, Nagel sometimes questions the stipulation that the confidence level of the HS-subject remains as high as that of the LS-subject (e.g., Nagel, 2008, p. 286, 2010a, p. 421). Though she just briefly

¹³ Nagel says, "The larger goal here is to show how variations in perceived epistemic anxiety can explain certain patterns in our attribution of knowledge, in particular, patterns involving changing stakes. If we need to see subjects as having made up their minds on the basis of evidence in order to see them as knowing, then changes in stakes could make a difference to whether a subject seems to know simply by making a difference to whether they seem to have made up their minds...[F]or our purposes here, it would be enough to establish that someone with high perceived epistemic anxiety would need to be seen as having more evidence than his low-anxiety counterpart in order to be seen as having a normal evidence based outright belief." (2010a, p. 419, italics added)

suggests such a confidence variability, this constitutes another distinctive feature of her account.

4.6 Objections to Bach and Nagel's accounts

In this section, I will draw on Rose and Schaffer's recent empirical findings to show that, contrary to what Bach and Nagel hold, knowledge doesn't always entail confident belief or closed-mindedness (by which I mean that the question of whether p is closed for the subject). Recall that Bach's doxastic pragmatist account of the bank cases relies on the assumption that knowledge requires a confident belief that is incompatible with any serious doubt—in contrast with mere (possibly unconfident) belief. Similarly, Nagel holds that confident belief or outright belief requires closed-mindedness on the relevant issue. Against these views, I am going to present some cases in which the subject knows but doesn't confidently believe the target proposition.

In a series of empirical studies, David Rose and Jonathan Schaffer show that people tend to ascribe knowledge and dispositional belief to subjects lacking a confident belief. They consider the following case (modified from Radford (1966)):

Unconfident Examinee. Kate is taking a history test. She had studied carefully and has been doing well on all the questions so far. She has now reached the final question, which reads "What year did Queen Elizabeth die?" As Kate reads this question she feels relief, since she had expected this question and memorized the answer. But before Kate can pause to recall the date, the teacher interrupts and announces that there is only one minute left. Now Kate panics. Her grip tightens around her pen. Her mind goes blank, and nothing comes to her. She feels that she can only guess. So, feeling shaken and dejected, she writes '1603'—which is of course exactly the right answer.

This type of case has been used in many places to argue that knowledge does not entail belief (Annis, 1977; Black, 1971; Margolis, 1973; Harker,

1980; Lewis, 1996; Myers-Schulz & Schwitzgebel, 2013; Ring, 1977; Schwitzgebel, 2010; Shope, 2002; Williams, 1973; Woozley, 1952). It has been generally assumed that Kate knows the answer since she answers correctly, but she lacks belief, since she feels no confidence in her answer. However, in their recent studies, Rose and Schaffer (2013) found that the majority of the participants tend to ascribe both knowledge and a certain type of belief to Kate.

One prominent explanation of the result is to distinguish two notions of belief, an *occurrent* and a *dispositional* notion. It is the dispositional notion of belief that people attribute to Kate. The *occurrent* notion denotes a thought consciously and confidently endorsed. The *dispositional* notion refers to the mere availability of information in one's mind. Arguably, *occurrent* belief is associated with some characteristic phenomenology, whereas *dispositional* belief merely involves a disposition to manifest such phenomenology. Concerning the Unconfident Examinee case, Kate lacks *occurrent* belief since she is unable to consciously endorse the thought. But she still has *dispositional* belief, for she has the information 'stored' in her mind, learned from her studies, and from that stored information she is able to unconsciously draw a right 'guess'. This verdict also fits with what I have argued in section 4.4, that dispositions can be masked by practical and psychological factors. In this case, Kate's disposition to recall her memory at a conscious level is masked by her temporary panic.

In addition to these experimental results, there are also strong theoretical reasons to hold that *dispositional* belief, rather than *occurrent* belief, is the type of belief entailed by knowledge. As Rose and Schaffer say:

If *occurrent* belief were required for knowledge, one would know far too little. A normal human adult knows all sorts of mundane propositions about her own personal life, current events, basic arithmetic, and various other topics at any given time (even when her mind goes blank; even when she is asleep). For instance, a normal human adult will have known that $7 + 5 = 12$ since early childhood without interruption, while only *occurrently* believing this on a handful of scattered occasions. Indeed a normal human adult will know many basic

arithmetic truths that she has not explicitly considered even once. (Rose & Schaffer, 2013, p. S23)

The above observation provides good grounds for the claim that the sense of belief in the entailment thesis “if *S* knows that *p*, then *S* believes that *p*” cannot be belief consciously endorsed by the subject at the present moment. Since dispositional belief accommodates the verdict that we could retain a huge amount of knowledge at any moment even though we do not entertain any explicit thought, it should be the kind of belief relevant to the entailment thesis.

These considerations constitute potential threats to Bach and Nagel’s views about what kind of belief is entailed by knowledge. In the Unconfident Examinee case, Kate loses confidence in the answer she gives due to panic and anxiety. Given that, the question’s answer is not closed for Kate at that moment. Kate is genuinely doubtful about the answer. Bach and Nagel’s views predict that Kate simply does not believe that Queen Elizabeth died in 1603 and does not know that proposition. However, as mentioned above, empirical studies indicate that Kate only lacks the occurrent belief in the target proposition but maintains the corresponding dispositional belief. Since mere dispositional belief (rather than occurrent belief) is required for knowledge, and no truth-relevant factor concerning the target proposition is affected in the case, Kate knows that Queen Elizabeth died in 1603.

Note also that, presumably, Kate will recover her memory at a later moment when her anxiety lowers. This indicates that she never stopped holding that belief, even though her ability to entertain it consciously was temporarily impaired.

The problem with Bach and Nagel’s view is that their accounts of belief seem to be plausible for occurrent belief, but not for dispositional belief, which is the only type of belief necessary for knowledge. Their doxastic accounts of the practical factor effects on knowledge ascription could work only under the overly demanding and intuitively implausible assumption that knowledge entails occurrent belief.

There are also other theoretical reasons to question the views of Bach and Nagel. It has been convincingly argued by Williamson (2000) that the KK principle, according to which if one knows that *p*, then one knows

that one knows, is wrong. Like many other mental attitudes, knowing is non-luminous, in the sense that one is not always in a position to know that one knows that p (Williamson, 2000, chapter 4). Consider a case in which one knows that p but is not in a position to know that she knows that p . For some reason, the subject engages in reasoning about whether she knows that p . It is possible that the subject finds out that she is not in a position to know that she knows that p . It would be reasonable for the subject to have some doubt and not be completely confident that p is true. However, the first-order knowledge that p is not *necessarily* destroyed for the simple fact that one realises that she doesn't have second-order knowledge that p .¹⁴ This constitutes another type of case in which one knows that p but is not highly confident in p .

4.7 Concluding remarks

In this chapter, we have critically examined some influential doxastic accounts of the practical factor effects on knowledge ascription available in the literature. These are all versions of doxastic pragmatism. However, we have seen that these versions face specific problems. Should we conclude that doxastic accounts are implausible? I don't think so. In the following chapters I develop a new type of doxastic account. This new account is based on a different sort of doxastic pragmatism, one on which practical factors affect credence rather than the threshold for outright belief.

¹⁴ Once someone realises that she doesn't have the second-order knowledge, this might provide good reasons for the subject not to rely on p as a premise in practical reasoning in a situation in which the costs of being wrong about p are high. See Williamson (2005), Gao (2019) and chapter 2 for further discussions of this issue.

PART II

Pars Construens

Credal Pragmatism

In the previous chapter, I critically considered some prominent doxastic accounts of the practical factor effects on knowledge ascriptions. According to such accounts, the practical effects of knowledge ascriptions can be explained in terms of the influence of practical matters on doxastic attitudes such as belief. These views are all varieties of doxastic pragmatism, according to which certain perceived practical factors generate systematic effects on normal subjects' outright belief. Other things being equal, a mere change in perceived practical factors tends to result in a shift of one's belief status in a certain proposition.

In this chapter, I am going to introduce a distinction between two types of doxastic pragmatism: threshold pragmatism and credal pragmatism. The types of doxastic pragmatism considered in the previous chapter are all versions of the former view. According to *threshold pragmatism*, an outright belief requires a quite high degree of confidence, above a certain threshold t (for those unfamiliar with the distinction between confidence and outright belief, I will explain it in more detail in section 5.2). Moreover, this view holds that the sensitivity of belief to the relevant practical factors is due to a corresponding sensitivity of threshold t to these factors. Thus, assuming that evidence remains constant across low- and high-stakes contexts, the threshold will be affected by stakes changes, while credence will stay the same. *Credal pragmatism* also holds that an outright belief requires a degree of confidence above threshold t . However, the relevant practical factors do not affect the threshold but the degree of confidence. Assuming that evidence remains constant across low and high-stakes contexts, confidence will be affected by stakes changes, while the threshold will stay the same.

Although credal pragmatism seems an obvious alternative, it has

received surprisingly little attention in the literature.¹ The primary aim of this chapter is to defend a version of this view. In section 5.1, after a short review of doxastic pragmatism, I introduce a standard classification of the relevant practical factors into two types: caution-oriented and rush-oriented. This distinction plays an important role in the characterization of the version of credal pragmatism I shall defend later in the chapter. In section 5.2, I provide a more accurate statement of the two types of doxastic pragmatism mentioned above: threshold and credal pragmatism. In section 5.3, I argue that certain intuitive and empirical data are better explained by credal pragmatism than by threshold pragmatism. In section 5.4, I consider and reply to possible worries about my arguments and more in general about credal pragmatism. A central issue of this chapter is how practical factors affect the formation of doxastic attitudes. A further issue is whether this influence is epistemically rational. In section 5.5, I consider the issue of whether the sensitivity of our doxastic attitudes to practical factors can be considered rational, and if yes, in what sense. In section 5.6, I sum up the main upshots of discussions in this chapter.

Before moving forward, let me add an important remark. My goal in this chapter is restricted to the comparative question of which amongst doxastic pragmatist accounts is better. Hence the present discussion will assume at the outset that some form of doxastic pragmatism is right. I've already discussed and criticised other views in the previous chapters (chapters 1–3). Readers who are not convinced by my considerations in those chapters are free to read my conclusions as merely conditional: amongst versions of doxastic pragmatism, the one providing the best explanation of the available intuitive and empirical data is credal pragmatism.

¹ To my knowledge, in addition to Nagel, other philosophers who have contemplated this idea have been Jonathan Kvanvig and Jason Stanley (2005, p. 6). Rubin (2015) considers and criticises a corresponding normative view about rational credence, which will become relevant in section 5.4.

5.1 Two types of practical factors

As I have already anticipated in section 4.1, according to doxastic pragmatism, certain perceived practical factors generate systematic effects on normal subjects' outright belief. More specifically, this view holds that other things being equal, a mere change in perceived practical factors tends to cause a shift of one's belief status in a certain proposition. This view is descriptive rather than normative. Normal subjects might fall short of being ideal rational subjects in many respects, but their mental states and behaviours present regularities that are largely predictable by psychology and tend to approximate certain standards of rationality, at least in most cases.² Doxastic pragmatists claim that one of such regularities obtains between perceived practical factors such as stakes-variations and doxastic attitudes such as belief, disbelief and withholding judgment.

By 'perceived practical factors', I mean features of the environment or the circumstance that the subject perceives or is influenced by. It is possible to have cases in which certain perceived practical factors have misleading appearances; for example, when one falsely believes that one is in a high-stakes situation. Furthermore, doxastic pragmatism is not committed to any specific account of the mechanisms by means of which perceived practical factors generate effects on outright beliefs. The view itself only concerns the causal correlation between perceived practical factors and those effects. Doxastic pragmatists have provided different stories of the relevant mechanisms, some of which were reviewed in the previous chapter: in terms of dispositions constitutive of belief,

² The standard of rationality I have in mind is not ideal or unbounded rationality. The latter kind of rationality (e.g., conformity to Bayesian standards) encompasses decision-making strategies that have little or no regard for the constraints of time, knowledge, and computational capacities that real humans face. It is widely acknowledged that real humans often go astray from exhibiting ideal rationality given their heavy reliance on fast-and-frugal heuristics in decision makings. Following Simon (1956) and Gigerenzer et al. (1999), I take that the type of rationality that applies to real humans is bounded rationality. In section 5.4, I will draw a more detailed distinction between these two types of rationality.

psychological forces generated by the factors and affecting doxastic states' regulation, and so on.

Relevant practical factors that could causally affect doxastic attitudes are diverse. They could be classified in various ways. Let me introduce a classification of such practical factors, drawing upon a well-recognised distinction between two intellectual duties or epistemic goals.³ On the one hand, there are practical factors related to the avoidance of error. These have the effect of leading one to assign more weight to the goal of not believing p if p is false. Examples are the high costs of being wrong about p , the significant benefits of achieving accuracy in judgment, and the availability of further evidence. On the other hand, there are factors related to truth acquisition. They have the effect of leading one to assign more weight to the goal of believing p when p is true. Examples include considerable benefits of having a settled opinion or remarkable costs of not having it and difficulty of acquiring further evidence. Let's name the former group *caution-oriented factors* and the latter group *rush-oriented factors*. Notice that this division is only a convenient tool in order to frame them in a simple and intuitive classification.

Using the above distinction between two groups of practical factors, we can characterise the systematic effects of perceived practical factors on beliefs, i.e., the sensitivity to practical factors (hereafter, *practical sensitivity*) of belief suggested by doxastic pragmatism, as follows: in the presence of rush-oriented factors, people tend to believe on the basis of less evidence and fewer cognitive efforts; by contrast, in the presence of caution-oriented factors, people tend to form beliefs only on the basis of more evidence and more cognitive efforts. For example, it is often the case that the higher the stakes, the more evidence one collects and/or the more cognitive efforts one allocates before forming a belief. On the contrary, the more urgent it is to form a settled opinion, the less evidence one gathers and/or the less cognitive effort one spends to reach a settled opinion. One of the consequences of the practical sensitivity of belief is that normally although one believes p in low-stakes circumstances, if the stakes on whether p had been higher one would lose the belief that p , even

³ This distinction was first introduced by James (1956). See Ganson (2008), Kelly (2014) and Wedgwood (2012, p. 325) for recent endorsements.

if the evidence for p remains the same.

5.2 Threshold pragmatism and credal pragmatism

This section introduces the distinction between two types of doxastic pragmatism: threshold pragmatism and credal pragmatism. Both these views rely on a distinction between two types of doxastic attitudes: belief and credence. *Credence* characterises the degree of confidence in a proposition. For example, among things that one believes, it is natural to think that one is more confident that $2+2=4$ than that Ulaanbaatar is the capital of Mongolia. Similarly, among things that one fails to believe, one could take one thing to be more unlikely than another. *Belief*, on the other hand, is a binary doxastic state. One either believes a proposition or does not believe it.

A hotly debated issue in epistemology and philosophy of mind concerns how these two types of doxastic attitudes relate to each other.⁴ According to a familiar view, so-called *threshold view*, according to which an outright belief is reducible to or requires a degree of credence above a certain threshold t .

Threshold view

For any subject S , proposition p , S believes that p if and only if S 's degree of credence is above threshold t ($Cr(p) > t$).

As anticipated in the introduction, both threshold pragmatism and credal pragmatism are kinds of threshold view. They both hold that an outright belief requires a certain degree of confidence above a certain threshold t . The difference between the two views concerns how the relevant practical factors affect belief. According to threshold pragmatism, such factors directly affect the threshold, leaving unaffected the degree of credence. More precisely, the threshold is practically sensitive in the following way: holding fixed the amount of available evidence, rush-

⁴ For a recent overview, see Jackson (2020). For a discussion on the relation between rational credence and rational belief, see Fassio & Gao (2020).

oriented factors tend to lower the threshold, while caution-oriented factors tend to raise the threshold. By contrast, according to credal pragmatism, they affect the degree of credence, leaving unaffected the threshold. More precisely, credence is practically sensitive in the following way: holding fixed the amount of available evidence, rush-oriented factors tend to raise the subject's degree of credence, while caution-oriented factors tend to lower the subject's degree of credence.

Here is a possible definition of these views:

Threshold pragmatism

Where a belief *B* is sensitive to relevant practical factors *p*, this is due to the sensitivity of threshold *t* to *p*. *p* affects *t* in the following way:

$$(t|\text{caution-oriented factors}) > (t|\text{no practical factors}) > (t|\text{rush-oriented factors}).$$

Credal pragmatism

Where a belief *B* is sensitive to relevant practical factors *p*, this is due to the sensitivity of the subject's degree of credence to *p*. *p* affects the degree of credence in the following way:

$$(Cr(p)|\text{rush-oriented factors}) > (Cr(p)|\text{no practical factors}) > (Cr(p)|\text{caution-oriented factors}).$$

So, for example, suppose that in a variant of the bank case, the subject passes from a low-stakes situation to a high-stakes one. According to doxastic pragmatism, this stakes-change affects the belief of the subject. In particular, when stakes go high, this has a causal effect on the subject's doxastic states so that she stops believing that the bank will be open on Saturday (*p*). According to threshold pragmatism, the high stakes 'move upward' the threshold *t* marking the degree of credence sufficient to believe *p*. The subject has the same confidence as in low stakes that the bank will be open, but now that stakes are high, that degree of confidence is not sufficient anymore to outright believe that *p* given the more demanding threshold. By contrast, according to credal pragmatism, the

high stakes leave the threshold t as I was in low stakes, but they ‘move downward’ the subject’s credence to a degree below t . The threshold is always the same in low- and high-stakes situations, but when the stakes are high the (lowered) degree of confidence is insufficient to reach it.

Upholders of doxastic pragmatism considered in the previous chapter have all endorsed some particular version of threshold pragmatism. We can easily identify endorsements of this view in the works of doxastic pragmatists mentioned above. Weatherston endorses a pragmatist threshold view in his account of belief’s practical sensitivity. This view holds that the subject’s perceived practical interests partly determine the threshold for outright belief. Weatherston writes:

In cases like this, interests matter not because they affect the degree of confidence that an agent can reasonably have in a proposition’s truth. (That is, not because they matter to epistemology.) *Rather, interests matter because they affect whether those reasonable degrees of confidence amount to belief.* (That is, because they matter to philosophy of mind.) There is no reason here to let pragmatic concerns into epistemology. (Weatherston 2005, pp. 435–436, italics added)

As Weatherston puts it, the responsible for why the subject lacks beliefs in cases involving abnormal practical interests is the practical sensitivity of the threshold for belief rather than the practical sensitivity of degrees of rational confidence.⁵

⁵ From Weatherston’s view about justification, we can arguably derive a pragmatic encroachment on justification. Weatherston suggests that by adding a normative operator to both sides of the thesis that S believes that p iff S prefers as if p , we have the claim that S is justified in believing that p only if S is justified to prefer as if p (2005, pp. 417–418). Consider cases where a subject irrationally maintains the same order of preferences conditionalising on p as in a low-stakes situation. But actually, the subject is in a perceived high-stakes situation. In such cases, according to Weatherston’s account of belief, the subject believes that p . But given that the order of preferences of the subject in high-stakes is irrational, hence not justified, according to Weatherston’s view about justification, the subject’s belief in p is not justified. Thus, it seems we can derive pragmatic encroachment on justification from his view about belief and justification.

Ganson builds threshold pragmatism directly into her account of belief:

In order to count as believing p in a range of circumstances, one must be willing to act as if p in those circumstances: one's degree of belief that p has to be high enough that one is willing to act as if p under those circumstances. (Ganson 2008, p. 452, italics added)

Here is another passage that explicitly illustrates her view:

Practical features, such as the cost if p should turn out to be false, [...] are relevant to determining the threshold for outright belief. Higher thresholds will demand stronger, more conclusive, or more plentiful evidence than lower thresholds. (ibid., p. 454)

Bach commits to a mobile threshold for confident, doubt-free belief. For example, he writes:

[T]he higher stakes raise the threshold of confident, doubt-free belief. [...] 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. (Bach 2008, p. 83)

Nagel also expresses her sympathy for threshold pragmatism. In particular, she finds that Weatherson and Ganson's idea, according to which the threshold for outright belief co-varies with the willingness to treat p as true for the purpose of practical reasoning, 'fits neatly' with uses of expressions such as 'desired confidence level' in psychologists' descriptions of certain cases (Nagel 2010a, p. 417). She observes:

Certain psychological accounts of epistemic anxiety even use terminology similar to Weatherson's: Daniel Hausmann and Damian Läge, for example, describe the variable evidence

In the rest of my discussion, I will ignore this potentially problematic consequence of his view and focus on his account of belief.

threshold as marking a “desired confidence level”, which rises and falls in step with stakes, among other factors (Hausmann and Läge, 2008). Alice Eagly and Shelly Chaiken describe transitions between heuristic and systematic ways of settling questions as motivated by discrepancies between the subject’s actual and desired confidence levels, where higher-stakes problems set higher desired confidence levels (Eagly and Chaiken, 1993). (Nagel 2010a, p. 417)

Nagel also identifies some data as evidence in support of a variable threshold view of outright belief. She writes:

In Ofra Maysseless and Arie Kruglanski’s work, subjects called off their search for evidence and reached a final decision on each digit at noticeably different average levels of reported confidence when they were in low-stakes as opposed to high-stakes conditions — 68.42 vs. 91.46 on the 0–100 scale (Maysseless and Kruglanski, 1987). (Nagel 2010a, pp. 417–418)

However, Nagel also suggests that, as a consequence of an elevated level of epistemic anxiety and lowered level of need-for-closure, it is natural for the reader to feel odd that the HS-subject remains as confident about the truth of the target proposition as she was before, or as his/her low stakes counterpart was. She writes:

DeRose himself takes the psychological attitude of the paired subjects to be held fixed by a stipulation that HIGH remains as confident as he was before that the bank would be open. It is arguably somewhat difficult to register this stipulation, given the other content in the scenario: in announcing a decision to “go in and make sure”, *HIGH certainly seems to be displaying lowered confidence*, in some sense of that term...Independently, if HIGH is seen to have just the same subjective confidence as his counterpart, this could be a sign of old-fashioned trouble in HIGH’s epistemic predicament: ordinarily, as we actively consider a broader range of hypotheses consistent with our

evidence, *confidence in any one of these hypotheses should fall* (Kelley et al., 1972). (Nagel 2010a, pp. 421–422, italics added)

In sum, almost all proponents of doxastic accounts considered so far endorse a threshold pragmatist view according to which believing that p is to have one's degrees of belief reaching a certain threshold which varies with one's perceived practical interests. With the exception of Nagel, all other doxastic pragmatists take the sensitivity of the threshold to practical factors to be part of an inference to the best explanation of the practical sensitivity of belief. However, they put forward the hypothesis with neither substantial argument nor evidence. They simply stipulate it.

Credal pragmatism has been far less popular than threshold pragmatism in the contemporary literature. According to credal pragmatism, *the practical sensitivity of credence* (i.e., the credence's variation in accordance with perceived practical factors), rather than the practical sensitivity of the threshold, is the cause of the practical sensitivity of belief. More specifically, according to the practical sensitivity of credence, holding fixed the amount of available evidence, rush-oriented factors tend to raise the degree of credence (i.e., subjective confidence), while caution-oriented factors tend to lower the degree of credence. According to credal pragmatism, it is credence that changes with practical factors, while the threshold on credence necessary for outright belief is unaffected by the relevant practical factors.⁶

Credal pragmatism has been a less popular view because practical sensitivity of credence implies violation of widely accepted evidentialist norms. In particular, according to the standard evidentialist view of rational credence, credence in a proposition p should be proportional to the degree of evidential support that one has in favour of p . In this picture, degrees of credence should be probabilistically coherent and should be exclusively updated on new evidence by following conditionalization rules. It is thus natural for philosophers, including most existing doxastic

⁶ Or at least the threshold is not affected in the way prescribed by threshold pragmatism. Nonetheless, it's worth noting that credal pragmatism is compatible with some moderate flexibility of the threshold generated by other mechanisms, as some of the data considered in section 5.3 seem to suggest.

pragmatists, to take credence to be insensitive to practical factors. However, such an assumption seems to be a mistake, as I will argue below.

In the next section, I will argue that credal pragmatism is the more plausible version of doxastic pragmatism. In particular, I will show that this view not only can accommodate data supporting threshold pragmatism but also fits well with other empirical and intuitive data that cannot be easily explained by threshold pragmatism.

5.3 Credal pragmatism vs. threshold pragmatism

In this section, I argue that credal pragmatism accommodates a wide range of intuitive and empirical data better than threshold pragmatism. First, let's consider intuitive data that are in favour of credal pragmatism. Consider a subject who experiences a transition from a low-stakes situation to a high-stakes one. For example, imagine a variant of the initial LOW-HIGH case in which Hannah makes many sandwiches, several of each kind (tuna, peanut butter and almond butter), arranged in the fridge in the same order as in the original case. Suppose Sarah receives visits on the same day from her friend Almira, who does not have any allergy, and her nephew Algernon, who has a severe peanut allergy. Almira arrives first. Sarah gives her an almond butter sandwich, picking it from the right side of the refrigerator. After Almira has left, Algernon arrives. As in HIGH, Sarah could not distinguish the peanut sandwiches from the almond butter sandwiches, so she served the tuna sandwich to Algernon.

According to threshold pragmatism, although Sarah doesn't believe r' —that the sandwiches on the right are almond butter—when she chooses the sandwich for Algernon, her confidence with respect to r' remains the same as before. However, it sounds odd for Sarah to say things like “I have exactly the same confidence in r' as before when I gave the sandwich to Almira, but I don't believe r' now.” On the contrary, it seems very reasonable for Sarah to mutter things such as “Well, I am not as confident in r' as before. This is why I don't know r' .” This indicates that, in accordance with credal pragmatism, the high stakes in the new situation affect not only Sarah's belief but also her confidence.

In addition, we can conceive a further variation of the case in which

another guest of Sarah, Bob, who knows all about the guests' preferences, Algernon's allergy and the disposition of sandwiches in the fridge, observes the whole story and asks Sarah why she didn't give Algernon the sandwich on the right. According to threshold pragmatism, the following conversation should sound perfectly fine:

Bob: "Why didn't you give an almond butter sandwich to Algernon?"

Sarah: "Well, I am not sure anymore that r '."

Bob: "But are you equally confident?"

Sarah: "Yes of course! After all, my evidence for r ' is the same as before. But even though I'm exactly as confident as before that r ', it is unreasonable for me to believe that now."

Compare Sarah's last response with another answer,

Sarah*: "Well, I don't feel as confident as before. I'd rather remain agnostic about r now."

At least to my ears, Sarah's response in the first conversation sounds very odd while the latter response sounds very natural. Once again, this intuition cannot be easily accommodated by threshold pragmatism, but it fits perfectly well with credal pragmatism.⁷

If you share the same intuitions in the cases considered above, we can

⁷ Someone may suggest that a rewording of the dialogue would not necessarily favor credal pragmatism over threshold pragmatism. In particular, she may observe that if we replace 'confident' with 'likely for you/me', the dialogue would not sound particularly odd. She may also observe that it would sound odd for Sarah to say: "Even though my evidence hasn't changed, given how much is at stake r is less likely for me now." I think that the above considerations touch an important point about the relation between confidence and likelihood. An explanation of why one may not find counterintuitive the modified dialogue is that we tend to attribute different meanings to expressions such as 'being likely for someone'. This expression may sometimes refer to subjective confidence, but more frequently it is used to refer to epistemic probability (probability of a proposition given one's evidence). The notion of probability relevant for epistemic chance and epistemic modals is almost universally considered to be epistemic—not subjective—probability. With an epistemic reading in place, it is clear that if in the

conclude from this that our intuitive judgments in cases involving a shift from low stakes to high stakes provide an at least *prima facie* support for credal pragmatism and speak against threshold pragmatism. Moreover, our intuitive reactions seem to reveal a general principle relating belief and confidence in line with credal pragmatism—that is, don't revise a belief if you don't lose confidence.

In addition to the above intuitive data, empirical studies on the phenomenon of need-for-closure provide further support to credal pragmatism. A series of psychological studies have investigated how motivational forces affect attainment and avoidance of what psychologists call 'closure'. 'Closure' is a notion first introduced by Arie Kruglanski referring to the phenomenological transition from hesitant conjecture to a subjectively firm and settled belief.⁸ The relevant motivational forces include 'need-for-closure' (NFC) and 'need-to-avoid-closure' (NTAC). NFC is a concept referring to some form of desire or other tendency to possess some knowledge on a given topic, any definite knowledge as opposed to confusion and ambiguity. NTAC is a concept referring to the opposite desire or tendency to avoid acquiring a definite answer to a question. Importantly, NFC and NTAC are conditions that can be triggered through manipulating circumstantial factors. Typically, rush-oriented factors bring about NFC, and caution-oriented factors give rise to NTAC. Nagel (2008, 2010a) appeals to some of the existing empirical researches on NFC and NTAC in order to bring evidential support to doxastic pragmatism. It turns out that some of those researches also provide data relevant for a comparative assessment of threshold

above case Sarah's evidence remains the same across the contexts, also the likelihood of *r* on her evidence remains the same (by definition). With such a reading, it is not surprising that the modified dialogue doesn't sound odd and Sarah's claim sounds inconsistent. However, if one moves to alternative readings of likelihood not indicative of subjective confidence, the intuitive judgments are not relevant to testing credal and threshold pragmatism.

⁸ In his words, closure is "the juncture at which a belief crystallizes and turns from hesitant conjecture to a subjectively firm 'fact'" (Kruglanski & Webster, 1996, p. 266). Given this definition, closure implies the self-transparency of one's belief; hence closure entails belief. The opposite is not always true: there can be non-transparent beliefs, and in such cases, belief doesn't involve closure.

pragmatism and credal pragmatism.

Let us start with evidence concerning the practical sensitivity of credence and credal pragmatism. A series of studies provides both direct and indirect positive evidence for the practical sensitivity of credence. Direct evidence can be found in a study by Mayseless and Kruglanski (1987, Study 2). This study features a measurement of the shift of confidence corresponding to every single increase of evidence in the process of reaching a settled opinion. In their study, all participants were asked to identify a digit very briefly shown on a tachistoscope. All of them were allowed to control the tachistoscope and were able to repeat the stimulus presentation as many times as they wished. When a participant formed a settled opinion about one digit, the next digit would start being presented, and so on until 12 digits had been named. Participants were evenly divided into three groups measuring specific conditions: NFC, NTAC and Neutral. One group of participants was told that there is a close connection between forming unambiguous, clear-cut opinions and high intelligence. This was supposed to heighten their NFC. Participants in another group were promised very attractive rewards only if all 12 digits were correctly identified. This was supposed to strengthen a tendency to maximum accuracy, thereby heightening NTAC. Finally, a third group of participants in a 'neutral' group were not given any special instructions or rewards.

During the test, participants were asked to rate their confidence in the guessed result following each presentation of the digit on a 0-100 scale, where 0 represents being not at all confident and 100 represents being confident beyond a shadow of doubt.⁹ This resulted in two kinds of data concerning the shift of confidence with each new piece of evidence. One is the *initial confidence*, which measures the initial confidence rating.

⁹ In the studies of Mayseless and Kruglanski (1987), the data on subjective confidence comes from participants' self-evaluation. What is measured, more precisely, is a higher-order evaluation about one's own confidence. If we accept Williamson's claim about the non-luminosity of mental states, the subjective confidence might not be always transparent to the subject herself. However, since paradigmatic examples of non-luminosity concern only marginal cases, this should not create a substantial problem for taking the empirical data at face value. That said, it might be good to take these empirical data with some reservation.

This reflects the influence of early evidence on one's confidence. The other is the *direct confidence shift*. It measures the difference between two adjacent ratings of confidence, and it represents the contribution of each new datum to the participants' confidence level. The study reports the following significant results: the averaged initial confidence is highest in the NFC condition (50.84), intermediate in the Neutral condition (34.78), and lowest in the NTAC condition (27.43). In addition, the averaged direct confidence shift in the NTAC condition is the smallest (13.46) and in the NFC condition is the highest (38.49), with the Neutral condition (20.44) falling in the middle.¹⁰

The above results vindicate the practical sensitivity of credence prescribed by credal pragmatism, i.e., in responding to the same kind of evidence, certain pragmatic conditions generate systematic variations in the subject's credence. Other things being equal, under the influence of rush-oriented factors, people tend to rely heavily on evidence and give more credibility to each piece of evidence than in normal circumstances in which those factors are absent. In contrast, with exposure to caution-oriented factors, people tend to assign less importance to early information and give less credibility to each piece of evidence than in normal circumstances.

In addition to this direct evidence for the practical sensitivity of credence, the same study also provides indirect evidence for the practical sensitivity of credence on the basis of a cluster of other data. First, Mayseless and Kruglanski also measured the number of presentations that each participant chose to observe. Moreover, they recorded the confidence of each participant in the formation of closure, which they call *final confidence*. The difference between one's initial confidence and final confidence constitutes the total confidence shift. The result of the total confidence shift divided by the number of presentations indicates

¹⁰ These data measure the confidence change score of all presentations. There is another group of data measuring confidence change of all presentations that excludes null presentations, where null presentations mean participant reports seeing nothing on the screen. Again, in this group of data, there is a significant difference in magnitude of confidence shift between the NTAC condition (18.07), the Neutral condition (32.64) and the NFC condition (42.6).

the average confidence shift given each piece of evidence. Call it *derived confidence shift*. Table 1 illustrates how the derived confidence shift in the three conditions is calculated based on original experimental data.¹¹ According to the results, although the values of the derived confidence shift are different from the corresponding direct confidence shifts, the derived confidence shift in the three conditions exhibits the same ranking pattern as the direct confidence shift, i.e., NTAC < Neutral < NFC. This result, again, vindicates the phenomenon that people tend to respond differently to the same kind of evidence under different pragmatic conditions, as prescribed by credal pragmatism.

	NFC	Neutral	NTAC
Initial confidence (IC)	50.84	34.78	27.43
Final confidence (FC)	78.03	57.98	90.8
Total confidence shift (TC = FC – IC)	27.19	23.2	63.37
Number of presentation (N)	3.24	5.14	18.28
Derived confidence shift (= TC/N)	8.39	4.51	3.47

Table 1. Calculation of derived confidence shifts under three conditions. Shaded lines are derived results based on original experimental data of Maysless and Kruglanski (1987, Study 2).

Further indirect evidence for the practical sensitivity of credence comes from an important phenomenon called ‘unfounded confidence paradox’. According to this phenomenon, heightened NFC generates higher confidence in less accurate judgments. This is already detectable in Maysless and Kruglanski (1987, Study 2), where participants in the NFC condition group tend to finish with higher confidence than those in the Neutral condition even if the evidential basis (in terms of the number of presentations) held by the former group is weaker than the latter. Two

¹¹ Note that the number of presentations clearly demonstrates the practical sensitivity of belief—that is, people tend to collect less evidence under NFC conditions and more under NTAC conditions in order to form settled beliefs.

series of studies add further confirmation to this phenomenon. One series of studies confirms part of the above results by reporting elevated subjective confidence under heightened NFC. The studies manipulate NFC in different ways—for example, by putting subjects in environments with distracting noises (Kruglanski & Webster, 1991; Kruglanski et al., 1993), making the task appear to be dull for subjects (Webster, 1993), and adding time pressures to decision-making (Kruglanski & Webster, 1991). According to another series of studies, heightened NFC increases the subject's tendency to use quick but less accurate, even biased, evidence-processing strategies (e.g. relying on stereotypes and various heuristics) in reaching judgments, as opposed to more extensive and accurate evidence-weighting and evidence-integrating methods (Freund, 1983; Freund et al., 1985; Gigerenzer & Todd, 1999; Heaton & Kruglanski, 1991; Jamieson & Zanna, 1989). Now, if confidence or credence were not sensitive to rush-oriented factors, it would be difficult to see how we can even make sense of the 'unfounded confidence paradox'. After all, if subjects' credence only depends on accuracy or truth-related factors, we should expect that in cases where less accurate judgments are reached, subjects end up with lower confidence than what it is actually reported. This supports the practical sensitivity of credence, providing further grist to the mill of credal pragmatism.

So far, we have considered evidence supporting credal pragmatism. Let us now consider evidence relevant for assessing threshold pragmatism. If we think that believing p is having credence in p that passes some threshold, then the formation of a settled belief implies that one's credence has passed that threshold. Recall that final confidence is the degree of confidence necessary for the formation of closure for each participant. Since closure is the point at which a subject reaches a final verdict, thereby forming a firm belief about the matter, we can also assume that that point is also indicative of the threshold level in that situation. Therefore, the final confidence is indicative of the confidence necessary for reaching the threshold in a circumstance. Should there be a flexible threshold affected by practical factors, it is reasonable to expect that the averaged final confidence under different conditions would somewhat reflect the (alleged) shift of the threshold. Assuming the correlation between the threshold level and the final confidence, according to

threshold pragmatism, the averaged final confidence in NTAC condition should be the highest and the one in NFC condition should be the lowest.

However, according to the results of Mayseless and Kruglanski (1987, Study 2), the averaged final confidences do not differ significantly between the NFC condition (78.03) and the NTAC condition (90.8), both of which are higher than the Neutral condition (57.98). Threshold pragmatism can explain why the final confidence in the NTAC condition is higher than the one in the Neutral condition, but it cannot explain why the final confidence in the NFC condition is higher than in the Neutral condition. Similar results featuring higher final confidence in NFC conditions than in Neutral conditions have also been reported in Kruglanski and Webster (1991), Kruglanski et al. (1993) and Webster (1993). These results indicate that the threshold for belief on credence is lower in normal cases than in cases involving rush-oriented factors—contrary to what we should expect if threshold pragmatism were true. In addition, Webster (1993) also finds that the final confidence in NTAC conditions is significantly lower than that in both the Neutral conditions and the NFC conditions.¹² Contrary to what threshold pragmatism predicts, this result indicates that the threshold for belief on credence in cases involving caution-oriented factors is lower than that in those involving rush-oriented factors. All these results constitute counterevidence against what threshold pragmatism predicts about the relation between the various conditions.

These data pose serious explanatory challenges for threshold pragmatism, but are perfectly compatible with credal pragmatism. According to credal pragmatism, the high final confidence under NFC conditions reported in most of the above-mentioned studies can be explained in terms of the fact that under those conditions one tends to be overconfident in one's judgment despite the paucity of evidence. As for the experiments highlighting a high final confidence under NTAC conditions, credal pragmatists can emphasize the robustness of the

¹² Webster detects this result in all three experiments. Here is the result of the experiments: i) Experiment 1 – 10.50 for NFC condition, 8.44 for neutral condition, 5.66 for NTAC condition; ii) Experiment 2 – 10.61 for NFC condition, 8.42 for neutral condition, 5.76 for NTAC condition; iii) Experiment 3 – 9.91 for NFC condition, 7.53 for neutral condition, 4.62 for NTAC condition.

epistemic position that subjects with elevated NTAC tend to acquire in the process of reaching a settled judgment. For instance, the high averaged final confidence in the NTAC conditions (90.8) in Mayseless and Kruglanski (1987, Study 2) can be seen as a natural consequence of the numerous repetitions of the digit presentation on average (18.28 times) and thus of the more robust evidential support. Finally, for what concerns the low final confidence under NTAC conditions reported in Webster (1993), credal pragmatists can explain it in terms of one's inclination to be extra cautious in these circumstances, possibly to a degree that overbalances the excellence of one's relevant epistemic position.

In summary, empirical research provides evidence for the practical sensitivity of credence in accordance with what credal pragmatism predicts and evidence against the practical sensitivity of the threshold as required by threshold pragmatism. Thus, empirical data provide support for credal pragmatism over threshold pragmatism.¹³

5.4 Objections and replies

Some may wonder whether the psychological studies discussed in section 5.1 are genuinely relevant and helpful to the present philosophical debate. In particular, one may wonder whether it is appropriate to draw conclusions about one's degree of confidence on the basis of studies using self-reports of how one would rate one's own confidence. Admittedly, phenomenological feelings of sureness and self-reporting are not the typical dispositions philosophers have looked at for measuring subjective confidence.

In response, I agree that these studies focus only on a specific disposition manifesting subjective confidence. In this respect, my

¹³ Here I focus on the standard version of threshold theory of belief where the threshold is smaller than 1 and construe credal pragmatism as committing to this version of threshold theory. Another way to construe credal pragmatism is to combine it with the credence-one view, which says that belief requires credence one or maximal confidence (Clarke, 2013; Dodd, 2017; Greco, 2015). In Gao (2024), I compared the threshold view version of credal pragmatism and the credence-one view version, and argued that the former one has several advantages over the latter.

arguments are supposed to provide only partial, defeasible, *prima facie* evidence for the superiority of credal pragmatism over other forms of doxastic pragmatism. However, I also think that introspection is one of the best means to access one's own mental attitudes (though a fallible one) and that sincere self-reports of confidence can provide important indications of our degrees of confidence. This is evident if we consider how these dispositions interact with other dispositions typically related to doxastic attitudes. For example, if someone has behavioural dispositions indicating that she is not confident that *p*, but at the same time sincerely reports that she is confident that *p*, we take her as having incompatible dispositions, indicating some incoherence in her doxastic attitudes.

This is not to say that such dispositions constitute infallible guides to our degree of confidence. Undoubtedly there are cases in which we lack transparent access to our confidence. Phenomenological dispositions are indeed fallible. However, cases in which these dispositions fail to manifest one's confidence are abnormal, often related to pathological conditions, and relatively uncommon. The studies concern a large number of subjects participating in several and varied studies in controlled conditions. A significant and systematic divergence of self-reports from actual confidence in all or most subjects would be surprising and in need of explanation. For this reason, I think that, in general, self-reports of normal and presumably reasonable subjects provide at least defeasible evidence of their degree of confidence.

A second worry about the studies is that at least some of the factors triggering NFC and NTAC conditions are not obviously like those typically featured in the philosophical literature on knowledge ascriptions. For example, there seems not to be much in common between factors such as being in high or low stakes and being told there is a close connection between forming unambiguous, clear-cut opinions and high intelligence (Mayseless & Kruglanski, 1987, Study 2).

My response to this kind of worry is twofold. First, I would like to observe that, as stressed in the introduction, the debate on doxastic pragmatism is not limited to explaining traditional cases relevant for adjudicating specific debates on knowledge ascriptions. The debate concerns more broadly the general way in which various kinds of practical factors affect the regulation of our doxastic attitudes. However,

second, I also think that there is a closer connection than one may initially think between the factors involved in psychological studies and those traditionally considered by philosophers. In particular, doxastic pragmatists are keen to include amongst practical factors affecting doxastic attitudes, not only stakes but also, for example, the significant benefits of achieving accuracy in judgment, urgency in forming a settled opinion, available cheap means of double-checking, or the difficulty of acquiring further evidence. The factors considered in the psychological studies triggering NFC and NTAC seem to be reducible, more or less directly, to one or the other of these factors. For instance, in the study of Mayseless and Kruglanski (1987), the promise of a reward for identifying the digits correctly is supposed to indirectly influence one's dispositions to accuracy by stressing the potential gain in getting the digits right, and thus the costs of being wrong about them. Similarly, being told there is a close connection between forming unambiguous, clear-cut opinions and high intelligence is supposed to create a tendency to reach conclusions faster, thus influencing urgency to form a settled opinion on the matter. Noises (for example Kruglanski & Webster (1991), Kruglanski et al. (1993)) and unattractiveness of the task (for example Webster (1993)) are supposed to increase the difficulty of acquiring further evidence.¹⁴

It is not my intention here to draw a full analogy between the practical factors involved in traditional cases and those in the considered psychological studies. However, I think that it is fair to stress important similarities amongst them. It is not casual that the very same practical factors relevant for credence shift in the experimental studies have been considered relevant by other doxastic pragmatists. Notably, Nagel (2008) used the same studies to argue for her view, and has interpreted the data in the same way in which I interpreted them—with the only difference being that I have stressed the special support that these data provide to a specific type of doxastic pragmatism. Some of these studies are also discussed by Mikkel Gerken (2017, chapter 12) in connection with the debates on

¹⁴ It is important to stress that these interpretations are not merely my conjectures. The authors of the studies designed them for tracking features such as significant benefits or costs of achieving accuracy in judgment, benefits or costs in forming a settled opinion, and difficulties of information processing, and interpreted their results in this way.

knowledge ascription and shifty epistemology.¹⁵ Since the paper's main focus is on the comparative evaluation of different forms of doxastic pragmatism, the fact that other doxastic pragmatists consider such studies as relevant for their views makes the present use not unmotivated, at least from a dialectical perspective.

Another worry is that most of these studies are a few decades old, limited in scope and number, and not designed to test or fit the specific alternatives at issue. One may have doubts about their significance for the present debate. I admit that the considered body of studies may be insufficient to draw definitive conclusions. However, I would like to observe that the set of studies is sufficiently complete and significant to provide at least *prima facie* evidence for my conclusions. It is important to stress that while these studies are a bit dated, they are still considered paradigmatic in the growing literature on the NFC, and their results have never been contested. Recent overviews in this literature show that the considered studies are neither partial, nor marginal or considered unreliable by the scientific community. On the contrary, they have been very well-received and continuously referenced in the growing number of studies on the theory of the NFC and its applications.¹⁶ Moreover, some recent studies in descriptive decision theory seem to provide further confirmation of the results of the studies discussed in the previous section. In particular, a series of studies shows that lay people's probability estimates for negative events are systematically influenced by features of perceived outcomes and risks, such as the severity of the consequences, affective richness, salience and vividness of the scenario (Harris et al., 2009; Loewenstein et al., 2001; Rottenstreich & Hsee, 2001; Sunstein, 2002). These studies partially confirm the effects of practical factors on credence predicted by the studies discussed in section 5.3.¹⁷

¹⁵ Weisberg (2020) uses the same cluster of studies to discuss related issues about the metaphysics of doxastic attitudes.

¹⁶ See e.g. Roet et al. (2015) for a recent review.

¹⁷ In addition, Tetlock and Kim (1987), Kassir et al. (1991) and Lerner and Tetlock (1999) report that participants who are accountable for their judgments by expecting to have to justify them to an audience are less confident in their judgments than those who are not accountable.

I conclude this section by considering a quite different kind of worry. One may wonder whether credal and threshold pragmatism are really incompatible alternative, conflicting views. In particular, one may conceive a more nuanced version of threshold pragmatism, allowing exceptions to the threshold shift in some specific circumstances, such as those considered in the above studies (which would be accounted for by credence shifts). According to this view—which we may call *mixed* pragmatism—the practical sensitivity of belief would be explained by the practical sensitivity of both credence and the threshold.

Although we cannot utterly deny the viability of this view, I think that credal pragmatism has at least two advantages over it. First, considerations of simplicity and ontological parsimony favour credal pragmatism. In order to explain the available data, a mixed pragmatism must presuppose two separate mechanisms of doxastic attitudes' formation and regulation: one mechanism regulating the practical sensitivity of the threshold in specific circumstances and another regulating the practical sensitivity of credence. In contrast, credal pragmatism can explain the same data assuming a single mechanism of credence regulation. Thus, credal pragmatism provides a simpler, more parsimonious explanation of the data than the mixed view. Hence, assuming a threshold sensitive to practical factors seems unnecessary and would require more burdensome ontological and psychological commitments. Second, as argued above, the idea that the threshold is sensitive to practical factors in the way suggested by threshold pragmatism seems incompatible with some of the available data. In particular, some of the previous studies indicate that the threshold is not sensitive to the relevant practical factors in any of the ways threshold pragmatists suggest. On the contrary, in the considered studies, the threshold sometimes seems to shift in exactly the opposite of the predicted direction, lowering in cases involving caution-oriented factors and rising in cases involving rush-oriented factors. Absent an account of why the threshold would be affected in completely different ways in different contexts, the mixed view would sound *ad hoc*, less principled than credal pragmatism.

5.5 Three dimensions of epistemic rationality

The practical sensitivity of credence provides a picture incongruous with the common normative presumption about doxastic attitudes, according to which credences should be exclusively sensitive to truth-relevant considerations. Does this suggest that the systems regulating the formation of doxastic attitudes are fundamentally defective (at least in the circumstances involving the relevant practical factors), and therefore that belief and subjective confidence influenced by practical factors are epistemically irrational? The answer, I suggest, is a qualified no. Building on recent empirical research in psychology, I will distinguish between three types of epistemic rationality: *ideal rationality*, *purist rationality* and *bounded rationality*. For agents who share the cognitive limitations typical of human beings, it is boundedly rational (even if not purist or ideally rational) that the amount of evidence necessary for forming outright belief is affected by practical factors in certain specific circumstances.

Other doxastic accounts do not distinguish between different kinds of epistemic rationality. Rather, proponents of these accounts seem to hold that it is simply epistemically rational to require more evidence for belief in high-stakes situations. In addition, as we saw in the previous chapter, according to doxastic pragmatism, i) the degree of rational credence is exclusively determined by the amount of available evidence, and ii) the degree of credence necessary for outright belief is variable depending on practical factors—where these accounts do not distinguish between occurrent and dispositional belief, and thus the claim should apply to both. In other terms, according to these views, the threshold that rational credence should reach in order to rationalize an outright belief is not stable and is sensitive to practical factors such as stakes. When the amount of evidence necessary for rational outright belief goes up (e.g., in high-stakes cases), the degree of rational credence necessary for belief also goes up. By contrast, according to the present version of credal pragmatism, given a certain fixed amount of evidence, it is boundedly (though not purist or ideally) rational for the degree of credence to vary in different circumstances depending on practical factors, while the threshold on

the degree of credence necessary for outright belief remains fixed across contexts. For example, in perceived high-stakes situations a (boundedly) rational agent would need more evidence in order to reach the same level of subjective confidence as her low-stakes counterpart.

In the literature, we can already find some arguments for the irrationality of credence sensitivity to practical factors. In particular, Rubin (2015) considers a pragmatic encroachment on credence according to which maximally rational credence is interest-relative or sensitive to practical factors. She shows that an agent whose credence shifts merely with a change in stakes is vulnerable to a kind of diachronic Dutch Book. Such an agent would accept a certain series of bets, provided respectively before and after a particular high-stakes situation turns up, that collectively leads to a sure loss. According to classical Dutch Book arguments, if a subject is vulnerable to a Dutch Book, then her credences cannot be (ideally or maximally) rational. In particular, one's credences cannot be rational due to violations of Conditionalization rules.¹⁸ Accordingly, as Rubin illustrates, having credence which is sensitive to practical factors cannot be rational, at least if we assume the standard Bayesian picture of rationality.

While admitting this point, we should keep in mind that the vulnerability to diachronic Dutch Book arguments can only work as an objection to normative theories of credence and credence update (for example, to a theory claiming that maximally rational credence *should be* sensitive to practical factors). Nonetheless, the practical sensitivity of credence and credal pragmatism are descriptive theses, concerning how belief formation, retention and revision work for normal human beings in cases involving perceived practical factors. Thus, the point made by Rubin doesn't constitute a direct threat to these theses.

¹⁸ There are a variety of different conditionalization rules. The most prominent ones are Bayesian Conditionalization ($P_{\text{new}}(X) = P_{\text{initial}}(X|E)$ (provided $P_{\text{initial}}(E) > 0$) and Jeffrey Conditionalization ($P_{\text{new}}(X) = P_{\text{initial}}(X|E_1) \cdot P_{\text{new}}(E_1) + P_{\text{initial}}(X|E_2) \cdot P_{\text{new}}(E_2) + \dots + P_{\text{initial}}(X|E_n) \cdot P_{\text{new}}(E_n)$). Conditionalization rules are considered as the exclusively correct methods for credence change. According to standard views (e.g., orthodox Bayesian theory), it is rational to modify one's credence on the basis of these methods, *and only on their basis*. This remark is important because, according to credal pragmatism, practical factors also affect credence in absence of new evidence.

Furthermore, it is important to acknowledge that many philosophers and psychologists have radically different opinions about which attitudes or methods should count as rational. Many of them do not buy the idea that a subject is rational only when her credences are exclusively sensitive to truth-relevant factors and are updated using standard formal methods. Rather, they suggest that human beings can still be considered rational in some important sense even when they do not obey strict Bayesian standards. They oppose standards of *bounded rationality* to standards of *unbounded* or *ideal rationality*.¹⁹

On the one hand, we have unbounded, ideal rationality that takes truth and accuracy as its only rationality standards, abstracting away from limitations of cognitive abilities. An ideally rational subject's doxastic attitudes are completely isolated from influences of non-truth-relevant factors, such as influences of psychological, emotional, practical or environmental factors. Ideally rational subjects' credences should obey principles based on rules of logic, probability theory and so forth and can measure the truth-relevant factors to the highest accuracy.

On the other hand, human beings are not ideal rational agents. In reality, our cognitive performances are bounded by serious physical, ecological and temporal limits. Even though the human mind commits to certain patterns of cognition (such as various types of heuristics, biases and fallacies) that are not recommended by the standards of ideal rationality, humans might not be irrational. *Bounded rationality* characterizes the type of rationality relative to subjects with constraints due to limitations of mental and environmental resources.²⁰ A theory of bounded rationality focuses both on the structure of the environments and on the adaptation of the capacities of cognitive systems to the environments, for example, through evolution and development.²¹

¹⁹ For overviews of discussions and relevant literature on different types of rationality see e.g. Samuels et al. (2004) and Hertwig and Pedersen (2016).

²⁰ According to a famous analogy suggested by Hebert Simon, "Human rational behaviour...is shaped by scissors whose two blades are the structure of the task environments and the computational capabilities of the actor." (Simon, 1990, p. 7)

²¹ One important line of research following this conception of bounded rationality goes under the label of 'ecological rationality'. This captures the importance of the

In the latter picture, it is boundedly rational that doxastic attitudes' regulation (formation, possession and retention) is adapted to environmental and practical factors. Our limitations due to cognitive resources (time, energy and etc.) do not allow us to carry on perpetual information seeking procedures or allocate infinite energy toward reaching a given judgment or opinion. The search and deliberation must be ended at some point. But we cannot terminate the search arbitrarily or consciously leave some relevant evidence out of consideration either; otherwise, we would end up forming a shaky and highly uncertain base of judgments for actions and decisions. Then how can we form a solid judgmental base without exhausting ourselves? From an evolutionary point of view, the natural thing to do is to allocate our energy in accordance with the demands of each task and reach closure at the point in which the accuracy of judgment is sufficient for a given purpose. As Kruglanski observes:

It seems that Mother Nature (probably via the evolutionary process) came to our rescue with a simple solution: the capacity to occasionally shut our minds, that is, develop the sense of secure knowledge that obviates our felt need for further agonizing deliberation. Is the solution adequate? Does it always work? Does it invariably yield the intended results? The answer is a threefold no (whoever claimed that Mother Nature was a paragon of perfection?), yet our capacity for closed mindedness allows us to get on with our lives, rather than remain in an indefinite cognitive limbo, perennially buried in thought, as it were. Besides, our mental shutdown is hardly irrevocable. When its potentially adverse consequences become salient, we often seem capable of reopening the internal debate and appropriately adjusting our opinions. (Kruglanski, 2004, p. 2)

Accordingly, bounded rationality allows the amount of evidence necessary

environment in constraining and enabling decision making. See Todd and Brighton (2016) for a recent development of the theory of ecological rationality and relevant references.

for forming and maintaining an outright belief to be affected by practical factors in certain specific circumstances. In particular, given practical features of a particular situation (such as the importance of being right, the availability of further evidence and time, etc.), sometimes it may be boundedly rational to gather more evidence before settling one's mind on a question, while in other situations it may be boundedly rational to stop inquiring and form a belief. Similarly, sometimes a certain amount and quality of evidence may be sufficient to rationally reach closed-mindedness on an issue, while other times this may be irrational. For example, the subject in perceived high-stakes cases will be boundedly rational to feel epistemic anxiety and suspend her belief on whether the bank is open, and the subject in conditions of urgency will be boundedly rational to inquire less about the matter at stake.²² Hence, although the practical sensitivity of belief is obviously irrational from the point of view of ideal rationality, it manifests bounded rationality. As argued above, the practical sensitivity of belief depends on the practical sensitivity of credence, so the latter is boundedly rational as well.

Credal pragmatism fits extremely well with the demands of bounded rationality. An agent whose credences vary with practical factors in the way specified by credal pragmatism would give more weight to error possibilities and reduce confidence in the relevant proposition in NTAC cases. For example, in high-stakes cases, circumstances engender cognitive pressures on the subject, requiring one to be cautious and avoid risks of getting things wrong on a certain matter. Lower degrees of confidence and underestimation of evidence in such cases can be seen as functional to the goal of postponing closure until one's epistemic position has been sufficiently strengthened. In other NTAC cases, environments can be cognitively cooperative, providing cheap means to strengthen one's epistemic position—such as the easy availability of further evidence. In those cases, diminished confidence allows for delaying the time of closure in order to take advantage of these means. In NFC cases, overconfidence helps in reaching opinions in contexts in which rapid formation of opinion and decision is more important than accuracy, or where

²² As I will discuss later in more details, this pattern of cognition falls into what is labeled as bounded cognition.

unfriendly environments involve obstacles to the enhancement of one's epistemic position. Thus, far from compromising the agent's ability to solve decision problems effectively and efficiently, credence's sensitivity to practical factors in different environments helps the achievement of one's specific goals.²³

In addition to ideal and bounded rationality, we can conceive a third type of rationality, which I call *purist rationality*. This type of rationality is still bounded by most of the limitations typical of human beings. However, purist rationality is insensitive to specific non-truth-relevant factors typical of abnormal cases, such as anxiety typical of high need-to-avoid-closure cases or haste typical of high need-for-closure cases. A purist rational (but not ideally rational) agent still needs to resort to various types of heuristics in her cognition and defeasible reasoning as opposed to probabilistic reasoning. However, this agent will react in every circumstance (including abnormal ones) as a bounded rational subject would react in normal circumstances (i.e., circumstances involving neutral need-for-closure). For example, in high need-to-avoid-closure cases (a type of abnormal circumstance), a purist rational agent will maintain the degree of credence and belief that she would have in normal (neutral NFC) circumstances, being fully insensitive to practical factors such as stakes and derived psychological ones such as epistemic anxiety.²⁴ While this type of rationality may not correspond to any ordinary use of the term 'rational', its role will become apparent when I discuss the rationality of dispositional belief in the next chapter. In my view, this is the kind of rationality occupying a central place in the discussions of many traditional epistemologists about knowledge and epistemic justification, whereas ideal epistemic rationality is more often discussed by Bayesians or formal epistemologists.

As I conceive of them, all three types of rationality introduced so far fall into the category of epistemic rationality broadly construed, for

²³ In Gao (2021), I also show that bounded rationality is also applicable to certain self-deception cases that satisfy what pragmatic encroachment considers sufficient conditions for epistemic rationality.

²⁴ With the obvious exceptions of attitudes about the specific circumstances and eventually higher-order attitudes about these attitudes.

all of them concern the achievement of epistemic goals, i.e., truth and accuracy.²⁵ Each type of rationality concerns a specific type of epistemic agent and differs with respect to the specific limitations of these agents and their environments and the ways in which they can achieve the truth-goal. Ideal rationality concerns ideal agents who are capable to reach the highest degree of accuracy and are not bounded by contingent limitations of cognition, time and environment. Purist rationality concerns epistemic agents subject to limitations of cognition, time and environment typical of human beings but whose doxastic attitudes and their regulation are not influenced by non-truth-relevant concerns and factors such as pressure and anxiety. Bounded rationality concerns epistemic agents like us with limited cognitive resources and confined by environmental boundaries, whose doxastic attitudes should be regulated by holding into account practical influences in order to maximize the achievement of epistemic goals. Therefore, each type of epistemic rationality assessment is incommensurable, since bounded, ideal and purist rationality reflects very different concerns. However, the three types of rationality could eventually coincide in normal circumstances when actual cognitive limits do not prevent the subject from reaching accuracy standards typical of an ideally rational agent. The key difference between the latter two types of rationality is that practical factors such as stakes and urgency are only relevant for bounded rationality but not for purist rationality. We would commit a serious mistake if we considered either type of rationality as superior or more important than the other. Thus, I suggest that we can only have a qualified answer to the question of whether it is epistemically rational for our doxastic attitudes and knowledge status to be affected by practical factors in an indirect way: it is rational in a sense (boundedly rational), but not in another (ideally and purist rational).

A legitimate worry here concerns whether bounded rationality has to be classified as a brand of epistemic rationality to the extent that it is affected by practical factors related to the circumstances and by the subject's non-truth-related concerns. In my view, as long as a type of

²⁵ I take this to be the mark of epistemic rationality as opposed to other types of rationality. However, if one conceives the difference between types of rationality in different terms, I am open to alternative ways of shaping the distinction.

rationality concerns the achievement of epistemic goals, i.e., seeking truth and accuracy and avoiding falsity and error, it should be classified as epistemic.²⁶ A boundedly rational subject in abnormal cases (those involving high NTC or high NTAC) will adopt heuristics and methods sensitive to cost-effectiveness in problem-solving in those circumstances. This does not prevent those heuristics and methods from being directed at forming true beliefs or avoiding false ones and be thereby sensitive to evidential and accuracy-conducive considerations. This is true even in high NFC cases in which, under the effects of practical factors, boundedly rational subjects don't allocate as much cognitive efforts to achieve true beliefs as they would in normal circumstances and base their judgments on mediocre evidence. Nonetheless, boundedly rational cognitive methods are still epistemically rational as long as they are conducive to true belief or avoid false ones and are sensitive to evidential considerations.²⁷

One could insist that a boundedly rational agent does not seem to maximize her epistemic goals in certain circumstances. For example, in cases in which someone takes into consideration the importance of a question in her cognitive efforts to investigate whether p or in forming a belief given relatively weak evidence when being wrong is unimportant. However, it must be noted that in such cases, the practical factors do not compromise her epistemic goals for practical reasons. They do not interact with the subject's doxastic deliberation as practical reasons, but rather work by indirectly influencing the strength of certain evidential considerations or by weighing more a certain epistemic goal than another (e.g., avoiding errors rather than getting the truth, or vice versa).²⁸ In this

²⁶ I take this to be the mark of epistemic rationality as opposed to other types of rationality. However, if one conceives the difference between types of rationality in different terms, I am open to alternative ways of shaping the distinction.

²⁷ This notion of epistemic rationality has been discussed in recent works in epistemic utility theory. Decision theory admits an influence of practical factors in the determination of utilities, such as for example psychological effects of risk aversion (Buchak, 2014). Some argued that these considerations apply also to epistemic utilities (Campbell-Moore & Salow, 2020). The relevant influence of practical factors does not make rationality less epistemic, for it is still directed at maximizing accuracy (or truth).

²⁸ On the indirect ways in which practical factors influence evidential considerations

respect, by taking into consideration the importance of a question in her investigation, an agent seeks to minimize the risk of error (eventually to the detriment of truth maximisation). Similarly, by forming a belief given relatively weak evidence when being wrong is unimportant, an agent strives to maximize the quantity of true belief (to the detriment of risk minimization).

One may still be concerned about the compatibility of credal pragmatism and Bayesianism. One may worry that if credences are subject to easy and systematic variance with simple changes in practical factors, the powerful Bayesian model can no longer be usefully applied to credences.²⁹ It is worth stressing that standard Bayesianism is commonly considered a normative theory of credence, not a descriptive one. This is confirmed by a series of studies showing that, as a matter of fact, in most circumstances people do not regulate their doxastic attitudes according to Bayesian standards and do not update credence using Bayesian methods.³⁰ Thus, if the worry is that credal pragmatism would threaten Bayesianism as a descriptive theory, there are independent reasons to think that Bayesianism is not a good model for how people actually regulate and update their credences. Nonetheless, one could wonder whether a systematic divergence from Bayesian standards at the descriptive level would constitute a genuine cost for credal pragmatism.

In response, I don't think we should be too pessimistic about the possibility of applying a Bayesian framework to credal pragmatism. The standard Bayesian model can still be usefully applied to one's credences in normal circumstances where no practical factors affect the subject's confidence. Furthermore, credal pragmatism is compatible with non-standard Bayesian models. In particular, Bayesianism has two main components: first, rational credences should satisfy the probability calculus (probabilism); and second, that rational credences should be

and epistemic rationality in the relevant cases, see e.g., Grimm (2011) and Wedgwood (2012, p. 325).

²⁹ Thanks to a reviewer of *Philosophical Studies* for pressing me to address this important worry.

³⁰ See Phillips and Edwards (1966), Robinson and Hastie (1985), Zhao et al. (2012), Douven and Schubach (2015).

updated by conditionalizing on new evidence. Agents whose credences are sensitive to practical factors can still satisfy probabilism. Even though affected by practical factors, a system of credences can perfectly cohere in ways that satisfy Kolmogorov's probability axioms.

It seems that if credal pragmatism is true, one would systematically violate Conditionalization principles. This is because credence's updates do not depend uniquely on new evidence, but also partially on practical factors. However, credal pragmatism is compatible with a modified version of conditionalization featuring *weighted evidence* rather than *raw evidence*. In this model, before factoring the evidential support into conditionalization rules, this support is weighted differently in different practical circumstances. More precisely, a lower probabilistic weight is given to evidential support in caution-oriented circumstances and a higher weight in rush-oriented circumstances. This mechanism allows for preserving a version of the Conditionalization principle as a useful tool for updating credence on new evidence also within a credal pragmatist framework.³¹

In sum, I do not think that credal pragmatism conflicts with a Bayesian model. On the contrary, we can conceive a Bayesian-friendly version of credal pragmatism preserving the core tenets of Bayesianism—probabilism and specific versions of the Conditionalization rule.

While I find compelling the above considerations in favour of classifying bounded rationality as a genuine kind of epistemic rationality, one may persist in disagreeing with the specific characterization of epistemic rationality I have provided or with the fact that bounded rationality is truth-conducive in the way that I described. In particular, many epistemologists conceive epistemic rationality in a narrower sense, as the type of rationality relevant for epistemic justification and knowledge-level belief (rather than for the achievement of truth and accuracy broadly

³¹ The present discussion may be a bit too sketchy and abstract. Unfortunately, the limited space doesn't allow a detailed discussion. Clarke (2013) shows that credence sensitivity to contextual factors is compatible with a Bayesian framework. Though I disagree with his framework, I think that it provides another illustration of how credence sensitivity to contexts can be compatible with a Bayesian framework.

conceived).³² I agree that bounded rationality is not a type of epistemic rationality in this narrower sense. As I will argue in the next chapter, the notion of rationality relevant for epistemic justification and knowledge-level belief is rather the purist one. With this precision in mind, I will use ‘epistemic rationality’ in the wider sense considered above, and thus I will assume that bounded rationality is epistemic. However, importantly, nothing in what follows depends on this assumption. If one prefers using ‘epistemic rationality’ in a stricter sense, not including bounded rationality, one is free to do so. In the next chapter, for the purpose of my argument, I only need a more modest assumption: that purist rationality is a genuine form of epistemic rationality (both in the wide and narrow sense considered above). On the basis of this weaker assumption, I will show that my view implies a form of moderate purist invariantism about knowledge.

Having clarified the three types of epistemic rationality, we are in a position to consider the rationality/irrationality of the practical sensitivity of doxastic attitudes at issue here in more detail. The practical sensitivity of belief and credence is obviously irrational from the point of view of the ideal rationality and purist rationality standards, but it manifests bounded rationality. The practical sensitivity of belief exemplifies so-called *bounded cognition*. As it has been widely recognised and studied, bounded cognition characterises how we human beings as finite creatures allocate our limited cognitive resources given the nature of a task: the higher anticipated rewards in making accurate judgments or anticipated costs in being inaccurate, the more energy one would allocate to a given task. Most times, allocating more energy would result in a more accurate judgment. In one of the earliest studies on bounded cognition, McAllister et al. (1979) found that MBA students were more willing to give positive assessments and select more complex and accurate hypothetical business decisions in high-stakes than in low-stakes situations. Studies mentioned in section 5.3 on how the evidence-collecting behavior is affected by the level of need-for-closure are also illustrations of bounded cognition.³³

³² See, for example, Burge (2003, 2010), Gerken (2013b), Goldman (1979) and Graham (2012).

³³ Different models for understanding bounded cognition have been proposed.

5.6 Concluding remarks

In this chapter, I have argued for a new picture of the ontological relations between doxastic attitudes, credal pragmatism. According to credal pragmatism, given a certain fixed amount of evidence, the degree of credence of a boundedly (but not epistemically ideal or purist) rational agent varies in different circumstances depending on practical factors, while the threshold on the degree of credence necessary for outright belief remains fixed across contexts. Credal pragmatism provides a wide-ranging picture of the nature and interaction of different doxastic attitudes, the role of non-truth-relevant factors in their rational regulations, and knowledge. Credal pragmatism fares better than threshold pragmatism in accommodating a range of intuitive and empirical data. The available data seem to vindicate the practical sensitivity of credence predicted by credal pragmatism. Furthermore, the data indicate that if there is a threshold on credence necessary for outright belief, this is not sensitive to practical factors as threshold pragmatism predicts. In response to the challenge that credal pragmatism doesn't fit well with the common normative presumption about doxastic attitudes, there are two sets of considerations. First, the practical sensitivity of credence and credal pragmatism are descriptive, not prescriptive theories. Second, the contemporary debate on the nature of epistemic rationality makes room for assessing practically sensitive credences as rational, at least in a qualified sense.

In order to clarify the type of epistemic rationality manifested by the practical sensitivity of credence, I distinguish three dimensions of

According to the 'bounded toolbox' approach, bounded cognition is manifested in cognitive strategy selection: depending on the practical demands, either an automatic and heuristic strategy or a more controlled and taxing strategy will be deployed in information process (Gigerenzer & Todd, 1999; Gigerenzer, 2008). Traits such as the need-for-closure (and its counterpart need-to-avoid-closure) work as a metacognitive determiner of the deployed process strategy (Stanovich, 2011). According to the alternative 'evidence accrual' model, the evidence threshold is determined by the decision context (Bröder & Newell, 2008; Lee & Cummins, 2004; Newell, 2005). I am indebted to Nagel (2010a, pp. 411–412) for the discussion of the relevant psychological works.

epistemic rationality: i) ideal rationality—one relative to procedures leading to highly accurate epistemic achievements in ideal circumstances, ii) purist rationality—one relative to epistemic achievements accessible to a human cognition exclusively affected by truth-relevant factors, and iii) bounded rationality—one relative to the cognitive regulation for non-ideal agents with limited cognitive abilities in real-life environments and possibly affected by non-epistemic factors. For doxastic attitudes, sensitivity to non-truth-relevant factors could be boundedly rational as long as it leads cognitively limited subjects to achieve epistemic goals in abnormal circumstances. More specifically, it is part of the requirement of bounded rationality to proportionate one's cognitive efforts and the strength of one's epistemic position to the practical significance of the relevant beliefs.³⁴ In the next two chapters, I will show how credal pragmatism accounts for other intuitions concerning concessive knowledge attributions and ordinary epistemic assessments supporting the knowledge norm of practical reasoning.³⁵

³⁴ There is a further question of whether practical sensitivity of credence is all-things-considered rational. In Gao (2023), I argue that it is.

³⁵ A significant part of this chapter is adapted from Gao (2019b).

Credal Pragmatism and the Distinction Between Dispositional Belief and Occurrent Belief

In the previous chapter, I defended credal pragmatism. According to this view, credence is sensitive to practical factors. For example, if the stakes on being right about whether p become very high, or we have easily available evidence relevant to whether p and no time constraints, we will tend to have a lower degree of credence, possibly insufficient for an outright belief. By contrast, if the stakes on being right about whether p are very low, or it is urgent to settle on a given opinion, we will tend to have a higher degree of credence, possibly sufficient for believing outright that p .

This chapter explores the implications of this view for the distinction between two varieties of full belief: *dispositional belief* (which is the type of belief typically considered necessary for knowledge) and *occurrent belief*. This chapter also considers how these two varieties of beliefs are affected by practical factors and explores the implications for our intuitions about high-stakes cases. Section 6.1 introduces the distinction between dispositional and occurrent belief. Section 6.2 clarifies practical factor effects on dispositional and occurrent beliefs. Section 6.3 proposes a new doxastic account of practical factor effects on knowledge ascriptions. Section 6.4 sums up the main upshots of this chapter.

6.1 Dispositional belief and occurrent belief

The distinction between occurrent belief and dispositional belief is widely

recognised among philosophers.¹ A similar distinction is also drawn in folk psychology. We believe many things, such as that the first letter of my given name is the 10th letter in the English alphabet, and that Turin is the capital of Piedmont. These are things we believe but don't think much about and don't actually come to the forefront of our minds without some special reason. But this is information held in our memory and could be manifested in certain circumstances. Beliefs like these are often referred to as *dispositional beliefs*. By contrast, at any given time (except when one's mind goes blank or when one is asleep), some thoughts are actively brought to mind whose contents vary depending on circumstances. Such actively endorsed thoughts are often referred to as *occurrent beliefs*.^{2,3}

¹ The distinction between the dispositional and occurrent conception of belief traces back at least to Ryle's (1949, p. 135) and Campbell (1967). See Schwitzgebel (2021, section 2.1) and Rose and Schaffer (2013, section 1.2) for a useful overview.

² Philosophers give different orders of priority to one form of belief over the other. According to Price (1969), given shared interests in the phenomenology of belief, early modern philosophers focused on an 'occurrence analysis' of belief in terms of an introspectible mental act. Occurrent belief occupied a central stage in this project and dispositional belief was mostly neglected. A famous example is Hume (1740)'s account of belief that treats beliefs principally as occurrences. Twentieth-century behaviourists, by contrast, switched to a 'dispositional analysis' in terms of overt behaviour as a result of shifted interests to the role of belief in the explanation of action. Dispositional rather than occurrent belief became the focus. Nowadays there are still philosophers privileging an 'occurrence analysis', such as Campbell (1967) who argues that the occurrent (by which he called 'episodic') conception is more fundamental than the dispositional conception, since the relevant disposition includes: "the tendency to react to some of the relevant situations with episodic belief." (p. 206) For an account of belief which combines both phenomenological and action-based criteria, see Braithwaite (1932–3).

³ I would like to remain neutral with respect to the contentious claim that occurrent beliefs can be actively formed in the sense discussed by Frankish (2004), according to which we have the power to decide what attitude to take towards a proposition. Frankish's idea is that we can consider a proposition, reflect upon evidence for and against it, and then decide whether or not to accept it as an object of belief (maybe under pressure of making up our minds on the issue). In my opinion, the attitudes described by Frankish are more like acceptance than belief (see section 3.2 for the distinction between belief and acceptance). By weighing inconclusive evidence, one's rational subjective confidence could be around 0.5. Making up one's mind on p with

Plausibly, when we say that knowledge entails belief, we are talking about dispositional belief rather than occurrent belief. For one would not have much knowledge, at least not as much as we think we have, if occurrent belief is the type of belief entailed by knowledge (Rose & Schaffer, 2013, S23).

According to the traditional picture, occurrent belief and dispositional belief are different aspects of the same state. For example, representationalism identifies dispositional beliefs with stored representations, and occurrent beliefs with activations of these representations, preparatory to their employment in reasoning and decision-making (see e.g. Fodor 1987). Likewise, in other frameworks about belief, such as dispositionalism, dispositional belief is conceived as a dispositional state of the subject, and occurrent belief is understood as the manifestation of this disposition. In addition to activated dispositional beliefs, occurrent beliefs also include those that are not drawn from memory but are just formed based on freshly collected evidence, or as a result of an explicit judgment on the truth of some matter about which one doesn't have a settled opinion beforehand. According to the traditional picture, if a subject holds an occurrent belief that *p*, she also holds a dispositional belief that *p*. Once a belief is occurrently formed for the first time, it is automatically added to 'the belief box' and ready to be recalled into explicit thought when needed.⁴

Arguably the distinction between dispositional and occurrent belief does not fit neatly with the distinction between conscious and non-conscious belief. While dispositional but not occurrently endorsed beliefs are non-conscious beliefs, occurrent beliefs can be either consciously or non-consciously endorsed, for some beliefs can be activated and influence our reasoning and behaviour non-consciously. For example, when I am

such relatively low degree of subjective confidence in *p* naturally involves a voluntary endorsement of *p* for non-epistemic reasons. This has usually been taken as a sign of acceptance rather than belief.

⁴ For this traditional picture of occurrent belief, see Schwitzgebel (2015, section 2.1) and literature quoted therein. See also Frankish (2004, section 2.1) for an overview. I use this analogy only for explanatory purposes and remain open to whether it is legitimate and whether it can be applied to my present account.

driving a car, my behaviour is guided by non-conscious beliefs about street regulations. It is plausible that those beliefs must be in an occurrent form to exert influence on my behaviour. Similarly, if we accept the idea that animals act on their beliefs, we do not commit to the idea that animals have conscious occurrent thoughts (Frankish, 2004, p. 16).⁵

Here I would like to construe the distinction between dispositional belief and occurrent belief in a way similar to how folks commonly conceive it. According to recent psychological studies, the mental condition of ‘closure’ recognized by Kruglanski and his colleagues seems to be necessary and essential for occurrent belief but not for dispositional belief. Recall that Kruglanski introduces the notion of ‘closure’ to refer to the phenomenological transition from a hesitant conjecture to a subjectively firm and settled belief. In his words, closure is “the juncture at which a belief crystallizes and turns from hesitant conjecture to a subjectively firm ‘fact’” (Kruglanski & Webster, 1996, p. 266). As Nagel summarises:

Achieving closure or judgmental commitment on a question puts an end to the experience of ambiguity and delivers the sense of having a firm answer. The opposite of closure is openness or judgmental non-commitment, in which we are able to continue juggling alternative possibilities, perhaps lingering in ambiguity or confusion. (Nagel, 2008, p. 286)

Rose and Schaffer (2013) hold that occurrent belief is something like explicit judgment, involving the conscious endorsement of the content.

⁵ It is worth stressing that the issue of the relation between occurrent and conscious belief is much more complex than how I have introduced it here, and it is dependent on specific accounts of belief. For example, according to certain dispositionalist views, the only relevant dispositions to the activation of an occurrent belief are dispositions to sincerely assert the believed proposition (e.g., Braithwaite, 1932–1933; Marcus, 1990), and arguably one can sincerely assert only what one consciously takes to believe. Similarly, certain forms of representationalism identify occurrent beliefs with beliefs consciously recalled to mind to be employed in reasoning. Furthermore, as pointed out by Frankish (2004, p. 17), occurrent belief, as it is commonly conceived by folks, may be unique to the conscious mind.

Similarly, according to Frankish (2004, p. 17), occurrent belief, as commonly conceived by ordinary people, may be unique to the conscious mind. While one may think that these accounts are too narrow in restricting occurrent belief to conscious episodes, the closure requirement is coherent with Frankish's and Rose and Schaffer's views, since conscious endorsement entails firm and settled opinion. Furthermore, it is reasonable to think that non-conscious occurrent belief also requires closure or full acceptance.⁶ If someone didn't settle her mind on whether *p*, it is highly dubious that she can rely on *p* as a background premise even at an unconscious level, as it happens when one recalls the way home without consciously considering the directions.⁷

Things are different for dispositional belief. Although it is true that if one has a dispositional belief that *p*, one often also has the disposition of closing the question of whether *p* and form a settled judgment, closed-mindedness is not necessary for holding a dispositional belief. This is easy to see when we consider special contexts in which the disposition to actualize a belief is prevented by contingent factors, such as a momentary difficulty in remembering a piece of information. In such cases, the subject has a dispositional belief, but due to contingent factors, she is unable to settle her opinion on the matter.

Empirical studies conducted by Myers-Schulz and Schwitzgebel (2013) and Rose and Schaffer (2013) on belief ascriptions helpfully illustrate the folk conception of the distinction between occurrent and dispositional belief. Recall the unconfident examinee case in §4.7. In that case, Kate has memorized the year in which Queen Elizabeth died before the test, but she doesn't feel very confident about the (correct) answer she gives. This case (along with some other more controversial cases) has been used to elicit folks' judgment about belief in the two studies mentioned above. However, the two studies provide very different results of certain patterns

⁶ While I find the latter thought very plausible, I recognize that it might be controversial. However, following Rose and Schaffer, I am tempted to classify all cases of belief not involving closure as not fully activated, and thus as not occurrent.

⁷ Notice that this claim should be qualified and restricted to doxastic attitudes. As I argued in chapter 3, one can also rely on non-doxastic attitudes in reasoning, such as acceptance.

of belief ascription.⁸

In the experiment of Myers-Schulz and Schwitzgebel (2013), the participants were simply asked, “Did Kate believe that Queen Elizabeth died in 1603?” According to the result of the experiment, the majority of the participants (63%) were inclined to deny belief to the protagonist. Later, Rose and Schaffer (2013) repeated the experiments with varying designs of the belief probe to elicit the dispositional reading of belief in the vignettes. For example, in one design, they modify the belief probe to include a parenthetical clarification of the sense at issue. Instead of simply asking whether Kate believes that Queen Elizabeth died in 1603, they ask: “Did Kate still believe (in the sense that she still held the information in her mind even if she could not access it) that Queen Elizabeth died in 1603?” According to their results, the majority of participants were willing to ascribe belief to Kate with an elicited dispositional reading: in three designs, the percentages of positive belief ascription are 74%, 58% and 71% respectively.

Rose and Schaffer (2013) argued that there are two reasons for thinking that Kate has a dispositional belief, but not an occurrent one. First, Kate seems still to retain that information in her memory. Presumably the memory trace is not eliminated by the momentary panic. Second, Kate does guess correctly, and presumably her memory guides her action in the background in some unconscious and indirect way. Presumably, the majority of participants in the experiments of Rose and Schaffer ascribe the target dispositional belief to Kate for the above reasons.

Back to the results of Myers-Schulz and Schwitzgebel’s study, if the majority of participants agreed that Kate has the target dispositional

⁸ The debate that the two papers are about is whether intuitive judgments about certain cases constitute evidence against the claim that knowledge entails belief. Using five vignettes, Myers-Schulz and Schwitzgebel report that the majority of participants tend to deny belief while ascribing knowledge to the protagonists in the case of the unconfident examinee. Rose and Schaffer re-ran the experiments introducing significant revisions to the vignettes. According to them, Myers-Schulz and Schwitzgebel have failed to distinguish occurrent belief and dispositional belief. By using probes where a reading of dispositional belief is elicited, Rose and Schaffer report an inclination to ascribe belief and knowledge to the protagonists in cases including the unconfident examinee.

belief, they would deny the target belief to Kate in the occurrent sense. After all, Kate is described as feeling unsure when she writes down the answer, which suggests that she is not confident at all about her answer (as it is also suggested by the title of the case).⁹ Apparently, Kate does not have a settled opinion when she writes down her answer. The result of Myers-Schulz and Schwitzgebel's experiment supports the closure requirement for occurrent belief.

While I think that the above-considered empirical studies and interpretations provide a sound basis for the distinction between occurrent and dispositional belief, someone may disagree with how these philosophers and I interpret the results of the experiments and draw the distinction. For the sake of argument, I will simply assume here that occurrent belief that *p* necessarily involves: i) closed-mindedness about whether *p* (be it explicit and conscious or not);¹⁰ and ii) (defeasible) dispositions to rely on *p* as a premise in reasoning and assert that *p* in the actual circumstance—e.g., if asked whether *p*, the subject would answer affirmatively. Dispositional belief doesn't necessarily involve these conditions, as Kate's case shows. Those who think that this distinction doesn't correspond to any ordinary intuitive one can read the present characterization as stipulative for now. As I will argue in this and the next chapter, this distinction has the advantage of settling several problems about belief, knowledge and knowledge ascription. This counts as a further argument for the claim that the present distinction is a real and

⁹ This is also supported by empirical researches on memory that suggests that confidence in memory-based beliefs appears to be constructed at the time of recall, rather than stored. In particular, Kelley and Lindsay (1993) and Koriati et al. (2006, 2008) suggest that the confidence in a recalled memory is influenced by the ease with which the agent recalls it and the amount of related information that comes to mind.

¹⁰ This requirement excludes another diagnosis of the data according to which the subjects occurrently believe that *p* but do not feel that they have enough evidence to act on it. For otherwise the subjects would occurrently believe that *p* without being closed-minded about *p*. In addition, as some empirical studies that I have discussed in section 5.3, under high need-for-closure, agents tend to be more confident about their final judgment than in normal circumstances in which the need-for-closure is neutral. It is very implausible that one could be highly confident in a proposition while thinking that one does not have enough evidence for that proposition.

substantive one, even if it failed to correspond to the distinction drawn in folk psychology.

With the above understanding of occurrent belief, it is easy to see how occurrent belief formation is affected by psychological factors like anxiety and panic. It should be out of doubt that in the unconfident examinee case, Kate dispositionally, but not occurrently, believes that Queen Elizabeth died in 1603. The moral of the case is that dispositional beliefs do not always automatically turn into their occurrent forms under the relevant triggering conditions. In particular, psychological factors such as panic and anxiety could block the proper activation of dispositional belief into occurrent forms.

Moreover, it seems that whether one has a dispositional belief should be separated from one's actual level of subjective confidence. Subjective confidence characterises one's feeling about how likely to be true a proposition is. In the cases considered above, the subjects' subjective confidence in a relevant proposition seems to be substantially lowered by influences of interfering psychological factors. For example, while in normal cases Kate would take p to be true, during the exam her confidence is shaken and lowered to the point that she takes p to be merely likely (or even not that). As a consequence, the prospect of having a uniform threshold view for the two varieties of belief seems to be undermined. At least, the threshold view introduced in the previous chapter (section 5.2), according to which belief requires a degree of credence above a certain threshold t , seems to apply only to occurrent belief, not to dispositional belief. However, this doesn't mean that a threshold view for dispositional belief is not possible. I will discuss the threshold view for both types of belief in the next section.

6.2 Dispositional belief, occurrent belief and practical factor effects

In the previous chapter (section 5.3), I introduced the notion of closure that characterises the phenomenological transition from hesitant conjecture to a subjectively firm and settled belief. Practical factors that could influence belief formation and revision can be separated into two

groups, given their influence on one's desire or tendency to reach closure. Rush-oriented factors bring about need-for-closure (NFC), and caution-oriented factors give rise to need-to-avoid-closure (NTAC). Psychological studies have shown that our belief formation and revision are sensitive to the influences of these practical factors. Given the distinction between dispositional and occurrent belief considered in the last section, one may wonder whether these two varieties of belief are subject to the effects of practical factors to the same extent and what exactly those effects are. This section aims to explore these issues.

According to the rough accounts of dispositional and occurrent belief discussed in the previous sections, only occurrent belief requires closure; dispositional belief doesn't. Thus, the effects of NFC and NATC only apply to occurrent beliefs. Accordingly, the empirical findings discussed in the previous chapter concerning need-for-closure effects constitute evidence for the practical sensitivity of occurrent belief only, not for the practical sensitivity of dispositional belief.

Note that NFC can be distinguished *in need for non-specific closure* and *need for specific closure*. The former kind triggers desires for any answer whatsoever as long as this is definite and does not bring about bias for specific judgments. A typical example is the desire to possess some knowledge on a given topic, any definite knowledge, as opposed to confusion and ambiguity (Mayesless & Kruglanski, 1987, p. 164). By contrast, when one's wishes and desires are only compatible with some judgmental contents and incompatible with others, one has a need for specific closure. This type of NFC is supposed to sway the judgmental process towards a particular direction that is antecedently considered desirable, leading one to possible wishful thinking (*ibid.*, p. 165). I will focus here on the need for non-specific closure (I will be using 'closure' to denote this type of closure).

People's NFC varies both intrapersonally (i.e., variations within one person, depending on the circumstances) and interpersonally (i.e., variations of baseline need-for-closure between different people).¹¹ High

¹¹ See Webster and Kruglanski (1994) for interpersonal variations; and see Kruglanski and Webster (1991, 1996), Webster (1993), Kruglanski et al. (1993) and Webster et al. (1996) for intrapersonal variations.

NFC can be triggered in the presence of rush-oriented factors that feature expected benefits of closure, such as situations when attaining closure brings others' approval. Other factors having similar effects include increased costs of continuing in ambiguity under conditions such as time pressures, and conditions rendering information processing more difficult and laborious, such as tiredness, distracting background noise and making the task seem dull. Conversely, high NTAC can be triggered in the presence of caution-oriented factors, including expected benefits of openness or costs of closure, like antecedent emphases on accuracy of the judgment, high costs of being wrong with the judgment, others' esteem and appreciation for accuracy, making the task seem enjoyable and interesting, etc. When there is no situational factor amongst those mentioned above, one's NFC is neutral. Neutral NFC constitutes a baseline NFC that differs from one subject to another.

As section 5.3 has shown, a series of studies found that the level of NFC or NTAC has significant effects on the amount of evidence necessary for one to reach closure. Given the relation between closure and occurrent belief, those studies also provide evidence concerning the effects of practical factors on occurrent belief formation. Recall that in Mayseless and Kruglanski (1987, Study 2), it was found that participants in the NTAC condition group repeated presentations many more times than the other groups before answering which digit was flashed: the NTAC condition's average was 18.28 times, while 5.14 times on average for the Neutral condition and 3.24 times for the NTC condition. Similar studies show that when a firm judgment is not yet formed, high NFC could make individuals feel uneasy with an absence of settled opinion and lead them to be considerably cognitively, although maybe not practically, impatient and hasty in processing information. Individuals under such conditions tend to seize on whatever information or hypothesis is offered and settle on inconclusive evidence or information accessed in the early stage (Mayseless & Kruglanski, 1987; Kruglanski & Webster, 1991, 1996; Kruglanski et al., 1993). In such cases, one would rely on less information than what it is normally required in forming a corresponding occurrent belief.

Given the inadequate evidential basis on which 'closure' is reached under high NFC conditions, one might argue that we should classify

closure as a sort of acceptance rather than belief. However, there are good reasons to classify it as belief rather than acceptance. Recall that in section 3.2, I drew a distinction between belief and acceptance along three dimensions. Closure fulfils two conditions for belief: commitment to truth and absence of voluntary control. Let me explain these points in more detail.

First, in cases in which closure is reached under NFC conditions, from the subject's point of view, a final judgment is reached exclusively on the basis of evidential considerations. Practical factors do not affect the formation of a judgment by directly influencing one's weighing of reasons to believe: we do not treat these practical factors as reasons to be weighed on a par with evidential reasons (Kelly, 2003; Owens, 2003). Rather, they exert their effects on how much one is willing to trust each piece of evidence by defining the working environment for cognition or the practical payback for accuracy.

Second, the adjustments made in cognition, such as the alteration of the amount of evidence required for an outright belief and the selection of information process strategy, are typically adapted to the practical demands automatically (Alter et al., 2007; Gigerenzer & Todd, 1999; Newell, 2005; Rieskamp & Otto, 2006). Thus, the achievement of closure is done involuntarily.

Empirical studies report that people tend to cherish uncertainty with heightened NFC and be cognitively cautious in belief formation. For example, compared with individuals with high NFC, individuals with high NTAC tend to be less influenced by early information and be reluctant to commit to a definite opinion (Maysseless & Kruglanski, 1987; Webster et al., 1996). The data suggest that one would need to collect more evidence than one would do with neutral NFC to form a settled opinion on which one is willing to rely for further reasoning.

High NFC and NTAC also have effects on the preservation of opinion after a question is closed. A high NFC would involve a tendency to 'freeze' one's opinion about a question. For example, one would be more resistant to reopening the question when there is new evidence or hypotheses that might threaten the opinion coming up (Kruglanski & Webster, 1991; Kruglanski et al., 1993. See also Kruglanski & Webster (1996) and Kruglansk (2004, chapter 5) for summaries of several other results). On

the contrary, conditions for high NTAC make one more prone to reopen the question and less resistant to new counter-evidence or alternative hypotheses (Mayseless & Kruglanski, 1987; Kruglanski & Freund, 1983; Freund et al., 1985).

Where does the above discussion leave us for what concerns the relation between occurrent belief and dispositional belief? Sensitivity to rush-oriented factors allows us to form an occurrent belief based on less evidence. Moreover, the resultant belief tends to be more stable given the ‘freezing’ effect of NFC. This suggests that occurrent beliefs formed with high NFC will be maintained in normal circumstances afterwards. However, the story is not so simple with impacts of high NTAC on occurrent belief given that effects of caution-oriented factors are the opposite to those of rush-oriented factors.

Suppose that a subject originally has a dispositional belief that p based on moderate evidence for p . In the type of case in which there are caution-oriented factors, heightening NTAC (such as high-stakes scenarios), when a judgment on the truth of p is called for, a dispositional belief that p might not automatically turn into its occurrent form. For example, according to the relation between hypothesis generation and NFC, when certain practical forces heighten NTAC, one tends to generate new alternative hypotheses that cannot be eliminated by current evidence. Given the presence of alternative hypotheses, one cannot reach closure, for closure requires that no alternative hypotheses consistent with the evidence come to mind. The presence of alternative hypotheses also tends to lower one’s subjective confidence in p —at least when there are no other countervailing relevant practical factors defeating the lowering effect. From the perspective of the threshold view, once one’s subjective confidence drops to a point below the threshold, one does not occurrently believe the target proposition. Similar considerations apply to other possible mechanisms influencing the regulation of occurrent beliefs triggered by practical factors. For example, mechanisms that regulate occurrent belief formation by directly affecting the level of credence and one’s psychological anxiety.

A question about such type of cases is whether the subject still maintains the dispositional belief that p once these mechanisms affecting credence have been triggered (for example, when alternative hypotheses

and anxiety have been raised). Here it is helpful to compare the effects of such practical factors to those of psychological factors. It seems that pure psychological factors, such as panic and anxiety, do not typically remove the dispositional belief.¹² These psychological factors only temporarily lower one's subjective confidence in the target proposition, not permanently. Once disturbing psychological factors are removed, the dispositional belief would become accessible to the subject again. For example, we can imagine that in the case of the unconfident examinee, when Kate calms down and recovers from the mental confusion after the test, she should be able to recall her memory fluently and correctly. Her belief about the death year of Queen Elizabeth becomes fully accessible again. Do practical factors, like psychological factors, only have temporary effects on credence and do not destroy dispositional belief at all? The issue is a bit more complicated. While psychological factors (e.g., agitation, distress, etc.) directly block access to one's stored information, the relevant practical factors (e.g., high stakes) that trigger high NTAC do not have immediate effects on the access to stored information unless they trigger psychological reactions that could affect the retrieval of one's memory.

Let's focus on simplified cases in which practical factors motivating high NTAC do not generate psychological disturbance in one's mind, but simply lower one's occurrent credence in the target proposition. We can consider, for example, the specific mechanism of generating new alternative hypotheses discussed above. Depending on the epistemic characters of different agents, in some cases, it seems plausible that the corresponding dispositional belief is maintained, while in other cases, it is destroyed. As mentioned in section 3, the baseline NTC (or NTAC), i.e., the level of NFC (or NTAC) that is not affected by any environmental manipulations, varies among individuals (Webster & Kruglanski 1994). Individuals with high baseline NTAC are cautious with one's beliefs in general, regardless of the practical relevance of a specific belief. It is reasonable to think that, for such types of agents, the presence of alternative hypotheses accompanied by a rise of NTAC does not simply go away by removing practical pressures. The subjective confidence of those

¹² Exceptions include specific circumstances in which traumatic experiences affect one's dispositional belief.

agents with respect to the target proposition would not automatically shift back to the original value when the high NTAC situation goes away. Once doubts appear in their minds, those doubts will linger unless evidence is explored thoroughly and the doubts or generated alternative hypotheses are properly eliminated by new evidence. Given that, it is probable that the agents with high baseline NTAC also lose the corresponding dispositional belief once the occurrent credence in the target proposition goes lower than the threshold.¹³ Thus, although dispositional belief is not sensitive to practical factors as occurrent belief, in certain circumstances, dispositional belief can be destroyed by high NTAC. In such cases, if the dispositional belief amounted to knowledge, also knowledge is doxastically defeated.

By contrast, agents with high baseline NFC have the tendency to arrive at a fixed opinion rather quickly on relatively shaky evidential grounds. It is reasonable to think that once the alarm of high stakes is removed, it doesn't take long for them to regain their confidence in the old opinion and stick to that opinion. For such agents, high NTAC only inhibits the generation of occurrent belief but probably not the corresponding dispositional belief.

We are now in a position to consider the plausibility of the threshold view as an account of the two types of belief. The standard version of the threshold view can apply to occurrent belief. As it has been shown, both occurrent belief and subjective confidence are sensitive to psychological and practical factors in a related way. The presence of occurrent belief requires closure and closure depends on the actual degree of credence in the circumstance. Hence, occurrent belief depends on the actual degree of credence.

However, the standard version of the threshold view doesn't apply to dispositional belief. For one can maintain a dispositional belief even though the actual degree of credence is rather low and the subject lacks

¹³ The phenomenon is coherent with what is identified as easy epistemic ascent—difficult epistemic descent by some philosophers. Against epistemic contextualism, it has been argued that one cannot properly adopt a low epistemic standard in a context in which the epistemic standard has shifted from high to low (Pritchard, 2001; see McKenna (2011) for discussion).

occurrent belief. As argued above, dispositional belief is more resistant to the effects of psychological and practical factors than subjective confidence. But we can still establish a reasonable connection between dispositional belief and subjective confidence. In particular, there can still be a tight relation between dispositional belief and the degree of subjective confidence one would have in normal circumstances. By *normal circumstances*, I mean circumstances in which no psychological factors could temporarily block access to one's possessed information or opinion and in which one's NFC is neutral. A modified threshold view for dispositional belief would be the following: there is a threshold such that an agent has a dispositional belief that p just in case she would have a degree of confidence in p greater than (or equal to) that threshold in normal circumstances.

Since occurrent belief is strictly related to the actual degree of credence and it is boundedly rational for credences to be sensitive to practical factors, it follows that bounded rationality governs credence and occurrent belief, the two doxastic attitudes sensitive to practical factors. By contrast, rational regulation (formation, revision and retention) of dispositional belief follows a purist rationality standard. This is because dispositional belief is defined in terms of subjective confidence passing the threshold in normal circumstances, where there is no non-truth-relevant factors interfering with the regulation of doxastic attitudes. By definition, purist rationality and bounded rationality coincide in normal circumstances. In such circumstances, there is no impediment for a dispositional belief that p to turn it into the corresponding occurrent belief that p (and vice versa).

One may wonder here whether purist rationality is the only one governing dispositional belief or whether this attitude can also be evaluated according to bounded rationality. This is not possible, given the definition of dispositional belief and bounded rationality. Bounded rationality is an assessment sensitive to circumstances in which doxastic attitudes and their regulations are affected by practical factors, but dispositional belief is here defined as the occurrent belief one would have in circumstances in which the degree of credence is not affected by influences of psychological and practical factors. Although bounded rationality standards do not generally apply to dispositional belief, bounded rationality can be negatively relevant to dispositional belief

in specific contexts. These are high-stakes contexts where the agent is boundedly rational (permitted) to abandon, not only occurrent belief, but also dispositional belief. So, while bounded rationality assessments do not apply to dispositional belief in general, they apply to some of its revision processes.

While occurrent belief involves *closed-mindedness* (by which I mean that the question of whether p is closed for the subject), dispositional belief does not. Since achieving closed-mindedness is influenced both by psychological factors (e.g., panic and anxiety) and practical factors (e.g., time pressure, the importance of forming an accurate judgment), occurrent belief formation is sensitive to these factors, whereas dispositional belief is typically insensitive to the influence of these factors (except in special circumstances considered above).

In spite of the differences mentioned above, the two varieties of belief seem to be clearly related. My hypothesis is that their relation is roughly the following. Recall that credal pragmatism is a variant of the threshold view, according to which an outright belief is reducible to or requires a certain degree of credence above a threshold t . The relevant threshold for occurrent and dispositional belief is the same. However, for occurrent belief, we should look at whether the *actual degree of credence* meets this threshold. For dispositional belief, we have to look at whether the *degree of credence one would have in normal circumstances* meets that same threshold. With *normal circumstances* here I mean those circumstances in which the degree of credence is not affected by influences of psychological and practical factors—when stakes are normal, there is no urgency to settle your mind on whether p , no epistemic anxiety, and so on. In other words, a subject S 's dispositional belief that p is the occurrent belief that S would have if circumstances were normal. Accordingly, assuming that we already have an intelligible notion of occurrent belief at hand, we could provide the following counterfactual definition of dispositional belief:

Dispositional belief

For any subject S , proposition p , S dispositionally believes that p if and only if, if S were in normal circumstances (with respect to whether p), then S would occurrently believe that p .

The following schema summarizes the overall picture of credal pragmatism about the doxastic attitudes and their respective rationality conditions:

Doxastic attitudes	Main relevant rationality standard applying to the attitude	Threshold	Relevant for knowledge?
<i>Credence</i>	Bounded rationality (possibly influenced by practical factors)	Not-applicable	Qualified 'yes' (credence constituting dispositional belief)
<i>Occurrent belief</i>	Bounded rationality (possibly influenced by practical factors)	Fixed and dependent on the actual degree of credence	no
<i>Dispositional belief</i>	Purist rationality (not influenced by practical factors)	Fixed and dependent on credence in normal circumstances	yes

Table 2. A summary of features of the three types of doxastic attitudes

6.3 Diagnoses of high-stakes cases

On the basis of the distinction between dispositional belief and occurrent belief drawn above, we can now provide new diagnoses of the high-stakes cases commonly used for motivating pragmatic encroachment. Besides providing a response to pragmatic encroachers' arguments against moderate invariantism, these diagnoses together constitute a new doxastic account of practical factor effects on knowledge ascriptions. The account is different from existing doxastic accounts, given that it is essentially based on the distinction between dispositional and occurrent belief, whereas the other accounts are not.

First, let's consider whether the HS-subject knows that p . Again, the belief relevant for knowledge is dispositional belief, not occurrent belief. Knowledge can be conceived as dispositional belief plus other traditional

conditions such as reliability, justification, etc. In ignorant high-stake cases, since neither the belief condition nor truth-conducive conditions for knowledge are affected, the subject knows that *p*. In cases where the subject is aware of the high stakes, the diagnosis is that the subject may know or not know that *p* depending on whether she preserves a dispositional belief in such circumstances.

Concerning whether the high-stakes subject would lose the relevant belief, almost all doxastic pragmatists that have been discussed in chapter 4 are committed to saying that the subject loses both occurrent and dispositional belief. This is because those doxastic pragmatists aim to explain the loss of the high-stakes subject's knowledge in terms of the belief condition for knowledge, and it is dispositional belief that is entailed by knowledge. But according to the account I provided in the previous section, the type of belief that is subject to direct influences of psychological and practical factors is occurrent belief. Dispositional belief is more stable than occurrent belief. Only in certain circumstances with heightened NTAC, dispositional belief can be undermined together with occurrent belief. Thus, whether the subject still has dispositional belief that *p* is an open question.

In cases where the lowered credence that *p* remains stable and will not go over the threshold for belief unless there is new evidence for *p*, the subject loses the dispositional belief that *p* together with the occurrent belief that *p*. In such cases, the subject loses knowledge as well. In other cases, the subject maintains the dispositional belief. In these cases, the high-stakes subject will readily regain high confidence in *p* (and occurrent belief) when stakes lower again without needing to acquire any new evidence about *p*. For example, you might lose confidence in the proposition that Julius Caesar was born in 100 BC at a psychological study in which you will receive an extremely painful electric shock if you judge that proposition incorrectly.¹⁴ Nonetheless, we can imagine that immediately after the study, you regain confidence in that proposition, given the absence of the stress of a pending punishment.

I have argued that it is sometimes boundedly rational for the high-

¹⁴ The original case is from Reed (2010, pp. 228–229). See also Fantl and McGrath (2009a, pp. 192–193).

stakes subject to lose dispositional belief and hence knowledge. However, although bounded rationality is a kind of epistemic rationality, it does not follow straightforwardly that it is *epistemically rational* for the high-stakes subject to lose dispositional belief and knowledge in such cases. Epistemic rationality, as I conceived it in the last chapter, encompasses ideal rationality, purist rationality and bounded rationality. Under perceived high stakes, an ideally rational or a purist but not ideally rational agent's credence, occurrent belief, dispositional belief and knowledge would remain untouched from any impact of practical factors. Thus, it is only boundedly rational, but not ideally or purist rational for the high-stakes subject to lose dispositional belief and knowledge, in such cases. And if the subject doesn't lose her dispositional belief in high-stakes cases, she knows as long as she satisfies other conditions for knowledge.

Since the doxastic attitude relevant for knowledge is dispositional belief and the rationality of dispositional belief is insensitive to circumstances involving abnormal psychological and practical factors, knowledge inherits such stability and insensitivity to contingent practical changes. Given that dispositional beliefs are governed by purist rationality standards, if other conditions necessary for knowledge, such as evidence and reliability, are also practically insensitive, the resulting account of knowledge can be a purist one. *Credal pragmatism is thus compatible with a moderate invariantist account of knowledge.*¹⁵

One may object that occurrent belief is not irrelevant for knowledge. Suppose that I have an occurrent belief that not-*p* but a dispositional belief that *p* (for example, in a high-stakes case). Even if we consider dispositional belief central for knowledge, one may say that an occurrent belief that contradicts a dispositional one should at least be regarded

¹⁵ Consider cases where one under high NFC conditions forms occurrent belief that *p* with high credence based on relatively low evidence for *p*. It is plausible that the formed occurrent belief is automatically added to the belief box and hence becomes a dispositional belief. Again, it is boundedly rational but not ideally rational and purist rational to do so. But is the agent in a position to acquire knowledge that *p* should *p* be true? The answer is no. For the belief that *p* is supported by inadequate evidence. Thus, practical factors can only undermine knowledge by sabotaging belief required for knowledge, but cannot generate knowledge by prompting the formation of the belief required for knowledge.

as a defeater of knowledge, and thus be relevant for knowledge. My reply here is that if the content of an occurrent belief is not- p and that of a dispositional belief is p , the former simply cannot be a defeater of knowledge that p . This divergent occurrent belief is the product of specific practical factors of the abnormal circumstances and it does not constitute or is based on any additional evidence for or against the relevant proposition or about the source of evidence of that proposition (as would do a genuine defeater). Thus it cannot provide a sufficient reason to abandon the dispositional belief.¹⁶

An example could be helpful for making this point clear. Suppose Mary knows her name at t_1 . At t_2 , she is presented with a bet on what her name is. If she wins, she will receive candy, but if she loses, she will have to give up everything she has. Assume that Mary loses occurrent belief in her name and eventually forms an occurrent belief that she doesn't know her name. However, this doesn't seem in any way to prevent Mary from knowing her name, and as a matter of fact, at a time t_3 , when the bet is not in place anymore, she can easily regain her occurrent belief in her name. I see no reasons to think that a temporary loss of occurrent belief and/or the formation of an occurrent belief in the contrary proposition generated by very high stakes should count as a genuine epistemic defeater of one's knowledge, functioning as counterevidence of the relevant proposition.

On the basis of credal pragmatism, we can construe an alternative doxastic account of the practical factor effects on knowledge ascriptions. Intuitive judgments about the high-stakes cases can be explained by appealing to the occurrent belief status that readers are inclined to attribute to the high-stakes subjects and the relation between the occurrent belief status and knowledge ascriptions. This account differs with respect to those of other doxastic accounts for what concerns i) the specific variety of full belief at stake and ii) the mechanisms regulating the

¹⁶ Note that the claim that occurrent belief is irrelevant for knowledge (rather than merely unnecessary) is an important, non-negotiable one for my view. If rational occurrent belief were sufficient to defeat knowledge, the resulting view would be incompatible with epistemological purism. Since rational occurrent belief is sensitive to practical factors, there would be knowledge defeaters dependent on practical factors such as stakes. The resulting view would be a form of pragmatic encroachment similar to that defended by Fantl and McGrath (2009a).

effects of practical factors on the relevant doxastic attitudes. This kind of doxastic account is more plausible than other doxastic accounts because it is backed up by a refined and well-developed theory of doxastic attitudes that is supported by independent evidence (see section 6.2).

Let's consider bank cases where the HS-subject is aware of the high stakes and denies knowledge that p to herself. Since this is a case featuring high NATC, and by default, the HS-subject is boundedly rational, it would be natural for the HS-subject to lose confidence in p and fail to form an occurrent belief that p . When one does not occurrently believe that p , it is natural for one to deem herself as not knowing that p and admit this. Likewise, an appropriate self-ascription of knowledge that p —and in general taking oneself to know—would require that one occurrently believes that p . The disposition to self-ascribe knowledge is one of the typical dispositions specific to occurrent belief. If we as readers take that the HS-subject doesn't occurrently believe that p , it would be unsurprising for us to find it felicitous for the HS-subject to deny knowledge that p to herself. It is worth noting that while the knowledge denial in such cases is appropriate to the extent that it is a natural consequence of the absence of occurrent belief, the content of the knowledge denial may be false. This happens when the subject still preserves dispositional belief that p and hence also knows that p .

Other high-stakes cases, such as the airport and boat cases,¹⁷ feature the speaker who is in a perceived high-stakes situation about whether p and another person who is presumably in a low-stakes situation. From the bounded rationality perspective, the speaker does not have enough evidence to support high credence in p . So, it is natural to perceive her as not having the occurrent belief that p .¹⁸ But a third-person knowledge

¹⁷ See section 1.1 for discussions of those cases.

¹⁸ This is the case although she can maintain a dispositional belief and knowledge (see section 6.1). It is also worth mentioning that there may be cases in which NATC in the situation is not so high to make boundedly rational a loss of occurrent belief, but still it may be irrational for the subject to act on the believed proposition for completely different reasons. For example, Lucy may refuse a very high-stakes bet about which is her name while maintaining full occurrent belief that her name is Lucy. In such case, her refusal to bet is not due to a lack in her epistemic position, but to other non-epistemic reasons such as the judgment that it is always morally wrong to accept such

ascription, like a first-person knowledge ascription, requires that the speaker occurrently believes that *p*. Thus, it is infelicitous for the speaker to ascribe knowledge to the subject.

As Nagel (2010a, p. 425) points out, when we evaluate the mental states of others who are less informed than us, we have the tendency to judge them as if they share the same information and our concerns. This tendency is often labeled *egocentric bias* (see e.g., Baron & Hershey (1988), Nickerson (1999), Royzman et al. (2004), Birch (2004)). Under the influence of egocentric bias, we as readers and the speaker in the case would project the concern of high stakes and the inadequate evidence possessed by the speaker on the subject of the knowledge ascription. This explains why we find it felicitous for the speaker to deny knowledge to the subject, for we cannot refrain from projecting the feelings we would have under high stakes and the lack of mental attitude (i.e., a lack of occurrent belief and self-denial of knowledge) on the subject. Again, it is worth remembering that this account appeals to occurrent beliefs and their boundedly rational dispositions in abnormal contexts. In such cases, purist rationality of the dispositional belief is not compromised. If the subject maintains the dispositional belief, she also knows.

The egocentric bias explanation delineated above also deals well with some controversial cases, such as the ignorant high stakes cases. Following Nagel (2010a) and Gerken (2017, chapter 12), when we read an ignorant high stakes case, we could be expected to feel the force of high NATC with respect to the target proposition and then project the need for more evidence to the ignorant subject at issue. Likewise, a similar account also works for the non-linguistic cases, since what determines our intuitive judgment of all those high-stakes cases, at the bottom, is how we would feel in that situation.¹⁹

In most of the perceived high-stakes cases involving first-person knowledge ascriptions, it is stipulated that the high-stakes subject retains belief that *p* or a high degree of subjective confidence in *p* as her low-stakes counterpart. However, such a stipulation is subject to many

type of bets.

¹⁹ For further discussions of egocentric bias accounts, see Sripada and Stanley (2012, section 4), Dimmock (2019) and Dinges (2021).

problems. It either i) undermines the intuition that the high-stakes subject doesn't know that p , ii) it undermines the perceived reliability condition for knowledge, or iii) readers do not register this stipulation given the effects of egocentric bias.

For what concerns i), as diagnosed above, it is plausible for the high-stakes subject to maintain a dispositional belief that p and it would be epistemically (both purist and boundedly) rational to do so. We should not exclude the possibility that readers perceive the stipulation of outright belief held by the high-stakes subject in the sense of dispositional belief. However, if readers understand the stipulation at issue in terms of dispositional belief, it is doubtful that the intuition that the high-stakes subject does not know that p or the felicity of self-denying knowledge that p still holds. After all, dispositional belief is the type of belief required for knowledge and other epistemic conditions for knowledge are held fixed across low- and high-stakes cases.

However, many readers may not be aware of the distinction between dispositional and occurrent belief. In this case, it is likely that by registering the stipulation that the subject retains the belief that p or a high degree of subjective confidence in p , the readers perceive that the reliability condition for knowledge is undermined. If the high-stakes subject closes the question of whether p in spite of the perceived high stakes, it is likely that the subject has been in the grip of wishful thinking or hastened to form her belief, which would undermine the reliability of her belief formation. Thus, the subject would lack knowledge because she would fail to meet the reliability conditions.²⁰

²⁰ The current explanation is similar to Nagel (2008, pp.219–220, 2010a, pp. 419). Sripada and Stanley (2012, p. 22) argue that the explanation only apply to cases where we are uncertain about the evidence-gathering strategy deployed by the subject and cannot apply to high-stakes cases where evidence-gather strategy deployed by the subject is specified. However, according to Nagel, in addition to evidence-gathering strategy, evidence-processing strategy also matters for the reliability of belief formation. Proponents for Nagel's explanation may say that in high-stakes cases where only evidence-gather strategy deployed by the subject is specified, whereas the evidence-processing strategy is not, the reader would tend to perceive the evidence-processing strategy to be defective and its reliability is insufficient even to qualify a person for knowledge in low stakes situations.

Of course, it is also possible that the readers do not register the stipulation at all, i.e., iii). This is because, given the presence of egocentric bias, we are inclined to project a lack of occurrent belief to the subject in spite of such stipulation. If the reader perceives the subject as not closing her mind on the relevant question, then it would also be natural for her to deny knowledge that p to the high-stakes subject. After all, as suggested above, self-ascription of knowledge that p normally requires that one occurrently believes that p .

6.4 Concluding remarks

This chapter completes another piece of the puzzle for credal pragmatism. Practical factors affect both dispositional and occurrent belief, but they do so very differently: while occurrent belief is dependent on the actual degree of credence and is highly sensitive to influences of practical factors, dispositional belief is only sensitive in a much more restricted way. The types of epistemic rationality regulating these two kinds of belief are also different: occurrent belief is governed by bounded rationality, whereas dispositional belief is governed by purist rationality. If we accept that dispositional belief is the type of belief entailed by knowledge (rather than occurrent belief), we reach the conclusion that credal pragmatism is compatible with a moderate invariantist account of knowledge. In this chapter, I have argued for a new doxastic account capable of explaining away intuitive judgments about high-stakes cases used for motivating pragmatic encroachment. Based on the distinction between the two varieties of belief discussed in this chapter, in the next and last chapter of this book I will introduce a new account of the intuitive relations between knowledge and practical reasoning that are commonly used to motivate the knowledge norm of practical reasoning.

Fallibilism and the Knowledge Norm of Practical Reasoning

In chapter 3, I argued that there are no general epistemic norms for practical reasoning. With ‘general epistemic norm’, I mean a norm applying to every instance of practical reasoning. I did that by showing that there are certain cases of practical reasoning to which doxastic norms (hence also the knowledge norm) do not apply. This argument and other objections to the knowledge norm of practical reasoning reviewed in chapter 1 provide some *prima facie* reasons to reject the knowledge norm of practical reasoning. However, proponents of this norm often motivate their view by appealing to the fact that ‘knowledge’ and its cognates play a prominent role in ordinary epistemic assessments of rational actions. In order to fully rebut these arguments, moderate invariantists need to fulfil a further task: to explain the prominent role of ‘knowledge’ and its cognates in ordinary epistemic assessments of practical reasoning and action. In this chapter, I will consider how a fallibilist moderate invariantist can meet this task.

Here is the plan of this chapter. Section 7.1 introduces fallibilism and reviews various contemporary formulations of this view. While I do not endorse a specific version of the view, I think that a probabilistic conception (discussed in section 7.2) better fits the version of credal pragmatism discussed in the rest of the book. Section 7.3 explains why fallibilism doesn’t square well with the knowledge norm of practical reasoning. In particular, as anticipated above, fallibilist views cannot fully explain why in folk epistemological practices, knowledge is taken to provide a sufficient epistemic ground for relying on a proposition in practical reasoning (i.e., the sufficiency version of the knowledge norm, SUFF; see section 3.1). I propose a new account of the intuitiveness

of knowledge assessments of rational action and practical reasoning. According to this account, such folk epistemological practices are closely related to an infallibilist intuition: if *S* knows that *p*, then there is no possibility for *S* that not-*p*. But this infallibilist intuition is false, given that it is inconsistent with fallibilism.

Fallibilists should explain away the infallibilist intuition. There are pragmatic accounts of the infallibilist intuition (e.g., Reed (2013), Gerken (2015, 2017)). However, they are not free of problems. My own proposal is a psychological-pragmatic account. This account explains the linguistic data used to motivate the infallibilist intuition, such as concessive knowledge attributions and mental versions of those data. The account is based on credal pragmatism and how this view conceives the nature of occurrent belief. In addition, I will show how the same account could explain the intuitiveness of the infallibilist intuition and ordinary epistemic assessments of action and practical reasoning commonly used to support the knowledge norm of practical reasoning.

7.1. Fallibilism and the knowledge norm of practical reasoning

According to an ordinary notion of fallibility, when we say that we are fallible knowers, what we mean is that human beings are prone to make mistakes. Two thoughts further motivate this claim. First, our cognitive faculties are limited and imperfect. Our knowledge is fallible in the sense that it is produced by fallible cognition. Second, we are often incapable of providing infallible reasons for our knowledge when we are asked how we know something. This happens especially when we don't remember the way in which the belief was acquired (Reed, 2012, p. 585). In such cases, we might admit we could be wrong upon reflection.

The epistemological doctrine of fallibilism is related but not identical to the above two thoughts. Fallibilism in epistemology is neither about cognitive faculties nor about providing articulable reasons for knowledge. It is about the character of the evidence supporting one's knowledge.¹

¹ For the sake of presentation, here I assume an evidentialist framework. The same

The basic idea is that we can know something on evidence that doesn't preclude the possibility of error. For example, you might know that your child went to school today because you left him at the entrance of the school. But there have been people in similar situations to your own whose children didn't go to school that day. They had reasons that apparently were just as good as yours to believe so, but they were wrong. Does the possibility of such 'bad cases' make us lose knowledge in the 'good cases'?² Contrary to familiar sceptics, fallibilists say no: we can and do have fallible knowledge. However, there are different ways to characterise what exactly it is to have fallible knowledge.

According to Reed (2002), "fallibilism is endorsed by virtually all contemporary epistemologists. Despite this near unanimity, or perhaps because of it, there has been some confusion as to how fallibilism should be best analysed." (p. 143) Given the variety of possible forms of fallibilism, it is imperative to first clarify the form of fallibilism that I endorse and compare it to other forms, some of which I deem less plausible. In the following subsections, I will introduce and examine some candidate conceptions of fallibilism. I will defend a version of probability fallibilism (which is also a version of epistemic modal fallibilism). This view relies on the threshold view discussed in previous chapters (chapters 5–6).

7.1.1 The logical conception

If there can be situations in which the evidence in the good case is the very same as it is in the bad case, then evidence that is sometimes sufficient for knowing that p must be logically consistent with the truth of not- p . The latter claim expresses what is sometimes called a *logical conception of fallibilism*. One way to formulate this conception is as follows:

point can be restated substituting evidence with other properties supposed to be justifiers.

² There are important varieties of 'bad cases'. While I stick to this popular terminology of 'good' and 'bad' cases, I observe here that 'bad cases' shouldn't be read as involving a strong negative evaluative connotation. Rather, 'bad case' merely expresses the thought that in such cases the subject is misled, though fully rational.

LF

S fallibly knows that p iff i) S knows that p on the basis of evidence e , and ii) e doesn't logically entail that p .³

Another way to express the logical conception of fallibilism is to put it in terms of a *logical modal condition*:

LMF

S fallibly knows that p iff i) S knows that p on the basis of evidence e , and yet ii) S 's belief that p on the basis of e could have been false.⁴

When the possibility in LMF is understood as logical rather than epistemic, the modal formulation is equivalent to the logical one.⁵ Notice that some philosophers use 'justification' instead of 'evidence' in the above formulations. Assuming that propositional justification only supervenes on evidence, this alternative use does not make any substantial difference.

In spite of being a popular formulation of fallibilism, logical fallibilism faces two familiar problems. First, logical entailment is not an epistemic relation. Satisfaction of logical entailment does not necessarily reflect the epistemic support of evidence to the relevant proposition (Harman, 1973; Reed, 2012). Second, this conception entails that knowledge of necessary truths is infallible no matter the evidence one has for these truths (Fantl & McGrath, 2009c, p. 57; Fumerton, 2006, p. 60; Hetherington, 1999, p. 565; Lehrer, 1974, pp. 82–83; Merricks, 1995; Reed, 2002, 2012, p. 586). Since a necessary truth p is logically entailed by everything, S 's evidence e will entail that p . And if it is necessary that p , then S 's belief that p could not have been false. However, presumably not all knowledge of necessary truths is infallible. For example, knowledge of some mathematical truths on the basis of testimonial evidence is hardly infallible. Given these

³ See Audi (1998), Cohen (1988), Fogelin (1994), Jeshion (2000) and Stanley (2005).

⁴ See Alston (1992), Ayer (1956), Bonjour (1985, 1998), Hetherington (1999), Lehrer (1974, 1990) and Reed (2002).

⁵ I assume here a notion of logical entailment as implying logical necessitation, i.e., if A logically entails B , then it is logically impossible that A and not B .

difficulties, neither LF nor LMF seems to be good enough to capture the familiar intuitive conception of fallible knowledge outlined at the beginning of this section.

7.1.2 The probability conception and the epistemic modal conception

According to another more prominent account of fallibilism, one can know something even though one's evidence or justification only makes the belief probable. This conception can be stated as follows:

PF

S fallibly knows that *p* iff i) *S* knows that *p*, and ii) it is epistemically probable that *p*.⁶

Philosophers have not reached a consensus on how to best understand epistemic probability. Some have taken epistemic probability to be grounded in frequencies in actual or relevant counterfactual situations.⁷ Others have taken the relevant sort of probability to be an internal relation holding between propositions, perhaps knowable *a priori*.⁸

According to another conception of fallibilism, to fallibly know is to know despite the fact that there is a non-zero epistemic chance for you

⁶ See e.g., Alston (1988), Bonjour (2010), Chisholm (1957, p. 28), Derksen (1978), Fumerton (1995, pp. 18–19), Goldman (2011, section 16.7), Lewis (1996, p. 551), Swinburne (2001, chapter 3) and (2011), Conee and Feldman (2004, fn. 32), Moser (1988), Plantinga (1993, chapter 9), Pryor (2004, pp. 350–351, 2005, p. 181), Reed (2010, 2013) and Russell (1948, chapter 5).

⁷ Goldman's reliabilism is one prominent example of a view that takes justification to be grounded in probability as a measure of either actual or counterfactual frequencies; see Goldman (1979, p. 96). Plantinga (1982, 1993) and van Inwagen (1996) hold counterfactual analyses of epistemic probability. See also Otte (2006) for criticism of both these views.

⁸ This is an idea rooted in Keynes (1921). Contemporary proponents include Fumerton (2004), Kyburg (1971) and (2003), Chisholm (1989a, pp. 54–56, pp. 63–64) and (1989b). See Russell (1948, part 5) and Mellor (2005) for more on these interpretations of probability.

that not- p . We may call this *epistemic modal fallibilism*:

EMF

S fallibly knows that p iff i) S knows that p , and ii) the epistemic chance of non- p is non-zero.⁹

According to most advocates of EMF (e.g., Dougherty & Rysiew, 2009, 2011; Fantl & McGrath, 2009a, 2009c) and proponents of PF (e.g., Reed, 2010, 2013), epistemic possibility or epistemic chance are equivalent to the notion epistemic probability.¹⁰ For example, according to Dougherty and Rysiew:

q is epistemically possible for S iff q has non-negligible probability on S 's total evidence (2009, p. 127).

Given such understandings, EMF can be taken as equivalent to PF. Here I assume this identification. I also endorse the equivalence between talks of epistemic modals and of epistemic probability. Hence, in the following, I use epistemic possibility/chance and epistemic probability interchangeably.

Proponents of EMF have come up with different characterisations of epistemic modals. It is reasonable to think that the suggested characterisations extend to the notion of epistemic probability given that the two notions are taken to be equivalent. According to the definition suggested by Dougherty and Rysiew (2009), q is epistemically possible for one iff not- q isn't entailed by S 's evidence (p. 127). Note that this account expresses the idea that knowledge is compatible with evidence that does not entail what is believed, which makes it equivalent to LF. But this seems to be a problematic way of construing epistemic possibility.

First, this view (as the other views formulated in section 7.1.1) implies the counterintuitive consequence that no necessarily false proposition

⁹ See e.g. Dougherty and Rysiew (2009, 2011), Fantl and McGrath (2009a, 2009c)

¹⁰ Proponents of the corresponding epistemic modal conception of infallibilism also equate epistemic possibility with epistemic probability, see e.g. Hawthorne (2004) and Dodd (2011).

can be epistemically possible, because the negation of that proposition is a necessary truth that is entailed by any other proposition. This also implies that one cannot fallibly know necessary truths, for the epistemic possibility of their negations would be zero. In order to avoid this problem, we should consider an account allowing that at least some necessarily false proposition is epistemically possible.

Second, Reed (2013) points out another problematic consequence with this characterisation of epistemic possibility in EMF: almost all claims that something is epistemically impossible will be false, given that very few negations of the believed propositions are excluded by one's evidence. Again, this consequence is highly implausible. It implies that when we claim that something cannot be true, most of times what we are saying is, strictly speaking, false. For example, I cannot truly say that it is epistemically impossible that in twenty seconds I will reach the opposite side of the universe.

A more plausible understanding of epistemic probability and epistemic modals should implement the following idea: the epistemic probability of a proposition reflects how good one's epistemic position is with respect to that proposition, regardless of whether the proposition is a necessary truth, a necessary falsity or neither of them. When the relevant type of evidence for a logical truth is inductive or testimonial, the epistemic probability of the relevant proposition could be less than one. Likewise, the epistemic probability of a necessary falsity could be non-zero.

Given these considerations, Fantl and McGrath (2009a, 2009c) provide a better account of epistemic chance. They suggest that epistemic chance governs one's rational betting behaviours: When it is epistemically necessary for one that p , it is rational to bet on p at any odds. But apparently, it is not rational to stake absurdly high amounts on many contingent propositions, such as the proposition that Plato taught Aristotle. On that basis, Fantl and McGrath argue that our epistemic chance of much of what we take ourselves to know is less than 1 (2009a, p. 59).¹¹ Their account can also accommodate cases of fallible knowledge about necessary truths and cases where necessary falsity is epistemically

¹¹ See Hawthorne (2012, section 8) for a further critical discussion of Fantl and McGrath's understanding of epistemic probability.

possible. Since they use rational betting behaviour to characterise epistemic probability, and betting behaviour is often taken to be a way of measuring rational credence, they commit to the view that rational credence should be equivalent to epistemic probability.

As argued in the previous chapter, my view on these matters is substantially different from that of Fantl and McGrath and many other epistemologists. In my view, far from being indicators of one's epistemic chance, betting behaviours could be associated with abnormal conditions in which rational credence does not equate to epistemic probability—especially when we consider extreme bets involving very high stakes (see sections 5.2, 5.3 and chapter 6). Moreover, it has been convincingly argued that betting dispositions are heavily affected by factors other than credence, such as risk aversion, motives besides money, and the format of the bet (e.g., Christensen, 2001; Earman, 1992; Weatherson, 1999). Others have argued that the very practice of gambling involves prudential and moral dimensions that inevitably misrepresent our actual degree of confidence (Salas, 2019). Eriksson and Rabinowicz (2013) convincingly argued that intuitive judgments about betting at most capture the degree of belief in a conditional that a proposition p would be the case if the agent were to bet on p , where the belief in this conditional itself is conditioned on certain stipulations (e.g. the opportunity to bet, how suspicious the bettor is, etc.) that inevitably alter the subject's epistemic state, so that an agent's choices in gambling cannot reflect her original credences. Some have also argued that betting scenarios where the costs of losing are extremely high almost inevitably affect one's evidence for a proposition, thereby modifying our initial epistemic position (and reasonable doxastic attitudes) toward a proposition (Dodd, 2017; Eriksson & Rabinowicz, 2013; Fassio, 2020, 2021; Hacking, 1965; Salas, 2019).

According to my account, epistemic probability is purely a factor of truth-conducive considerations. For a cognitively rational agent, the formation of credence is governed by mere truth-conducive considerations *only in normal circumstances*. Recall that by normal circumstances, I mean circumstances where no psychological factors could temporarily block access to one's possessed information or opinion and one's need-for-closure is neutral. In normal circumstances, a cognitively rational agent's credence accurately measures the strength of one's epistemic position

with respect to a proposition, being affected exclusively by truth-relevant considerations. In this framework, the epistemic probability of p for a subject is equivalent to the degree of her rational credence in p in normal circumstances. By contrast, contexts in which an agent is proposed a very high bet are precisely the type of abnormal circumstances in which one's rational credence should be affected by psychological factors and change accordingly, given bounded rationality standards (see sections 5.5 and 6.2).

In my view, it is difficult to find a method for measuring the rational credence of a subject in normal circumstances. Most measuring methods could involve abnormal circumstances that could modify rational confidence. However, we can assume that it is still possible to identify suitable measuring methods measuring degrees of credence in normal circumstances. For example, we can compare our degree of credence in different propositions and consider on which proposition we would be willing to place a low bet if we had the chance. If we apply such comparative methods, we can see that in normal circumstances our rational credence in most of the propositions that we know is less than one, and the epistemic probability for us in these propositions is also less than one, which is just what PF says. Thus my view is compatible with a version of PF.

The main problem with PF and EMF is that certain utterances of sentences that according to PF and EMF are literally true sound odd. This casts doubts on the truth of PF and EMF. Here is David Lewis:

If you are a contented fallibilist, I implore you to be honest, be naïve, hear it afresh. "He knows, yet he has not eliminated all possibilities of error." Even if you've numbed your ears, doesn't overt, explicit fallibilism still sound wrong? (Lewis, 1996, p. 550)

Likewise, David Hume objects to the use of probability in characterising knowledge:

But knowledge and probability are contrary and disagreeing natures, that they cannot well run insensibly into each other, and that because they will not divide, but must be either entirely present, or entirely absent. (*Treatise*, I. iv, 1, 1968, p. 181).

In defending PF and EMF, one must provide a non-semantic account for the infelicity of these sentences, so-called *concessive knowledge attributions*, whose truth is entailed by PF and EMF. Concessive knowledge attributions are sentences of the form “*S* knows that *p*, but it is possible that *q*” where *q* obviously entails not-*p*, and the ‘possible’ refers to epistemic possibility (Rysiew, 2001). For example, “I know it’s raining, but it’s possible it’s sunny.” Concessive knowledge attributions are typically infelicitous. The question is how to explain their infelicity. In section 7.3, I will discuss such accounts.

7.1.3 The fail-to-be-knowledge conception

Reed (2002, 2012, 2013) proposes a fail-to-be-knowledge conception of fallibilism:

FKF

S fallibly knows that *p* iff i) *S* knows that *p* on the basis of justification *j* and ii) *S*’s belief that *p* on the basis of *j* could have failed to be knowledge.

More specifically, according to Baron Reed, there are two ways that *S*’s belief that *p* could have failed to be knowledge despite being held with the same justification. First, *S*’s belief could have been false. This is the point highlighted by LF and MF. Second, *S*’s belief could have been accidentally true, as in Gettier cases.

FKF avoids the problem with necessary truths, for it allows one to have fallible knowledge about necessary truths. Although one’s belief in a necessary truth could not have been false, it could have been true just by accident. For example, one could acquire a belief about a mathematical truth through testimony from a reliable source about mathematics, but the belief may fail to be knowledge in a way familiar in Gettier cases. In other words, knowledge about necessary truths, like knowledge about contingent matters, can be Gettierised, and thus be fallible knowledge.

Reed takes FKF to be equivalent to PF. Nonetheless, we have reasons to think that FKF doesn’t commit to a probabilistic conception of justification or evidence. For example, according to the normic theory

of justification suggested by Smith (2010, 2016), justification requires one's belief in p being normically supported by one's evidence E , in the sense that the circumstance in which E is true and p is false requires more explanation than the circumstance in which E and p are both true. That means that if p turned out to be false despite E , some special explanation of this error would be needed, whereas no special explanation would be needed if E and p were both true. The degree of justification is measured in terms of the degree of normic support.¹² According to this view, mere probabilistic support cannot grant any normic support, and is thus irrelevant to epistemic justification.

7.1.4 A common challenge

The various forms of fallibilism that have been discussed so far are subject to one common challenge. It has been argued that evidence in good cases is utterly different from evidence in bad cases—where, recall, the terms ‘good case’ and ‘bad case’ are used in a technical sense: in good cases the subject knows; in bad cases the subject is misled, though possibly fully rational. For example, disjunctivists claim that veridical and non-veridical sensory experience does not share some highest common factor. Some argue that in the good cases, one knows that there is a table in front of her on the basis of factive evidence such as “I see that there is a table”, whereas in the bad case the evidence is “I seems to see that there is a table” (see e.g. McDowell, 1995; Pritchard, 2014).

In addition, some epistemologists, such as Williamson (2000), hold that evidence that grounds contingent perceptual knowledge, memory knowledge and testimonial knowledge entails the truth of the relevant proposition, whereas evidence in counterpart bad cases does not. In particular, according to Williamson, all and only knowledge is evidence. The equation of total knowledge and evidence implies that the epistemic

¹² According to Smith's normic theory of comparative justification, the strength of a normic support relation is determined by the normalcy gap between $E \wedge p$ and $E \wedge \neg p$. The larger the normalcy gap between $E \wedge p$ and $E \wedge \neg p$, the stronger the normic support relation between E and p . See Smith (2016, chapter 5).

probability for a known proposition is 1.¹³ For if I know that p , knowledge that p is part of my evidence. If this is the case, knowledge will be always infallible, contrary to logical, logical modal, probability and epistemic modal fallibilism. The same problem applies to the fail-to-be-knowledge fallibilism when this is conceived as equivalent to probability fallibilism.

In order to preserve their views, fallibilists have to deny disjunctivist assumptions (at least for what concerns perceptual knowledge) and the equivalence of evidence to knowledge. Unfortunately, providing a full argument against these assumptions would require a full book on its own. Here I will simply take for granted that these assumptions are false.¹⁴

7.1.5 The non-maximal conception

Lastly, we have the non-maximal conception of fallibilism. This fallibilism can be formulated as follows:

NMF

S fallibly knows that p iff i) S knows that p and ii) the strength of S 's epistemic position with respect to p is not maximal.

NMF is taken to be a weaker position than other forms of fallibilism discussed so far (Fantl & McGrath, 2009a, 2009c). First, it seems that NMF is entailed by LF, LMF, PF, EMF and FKF. If one has fallible knowledge in the logical or the probability sense, then one's evidence or justification cannot be perfect. It would be better if it were truth-entailing or having epistemic probability being 1, no matter how epistemic probability is understood. If one has fallible knowledge in the sense that the belief could have been accidentally true, then one's justification can be strengthened

¹³ For a proposition that is not known (e.g., a Gettiered belief), its epistemic probability is almost invariably less than one. I say 'almost invariably' because of cases in which one has probability one in a proposition in spite of not knowing that proposition. For instance, if someone has a ticket in a lottery with infinite tickets, her chance of losing will be zero, but she doesn't know that her ticket is a loser.

¹⁴ For objections to the thesis that all and only knowledge is evidence, see for example Arnold (2013), Comesaña and Kantin (2010), Fantl and McGrath (2009c), Joyce (2004), Rizzieri (2011).

to a degree that excludes the possibility of being Gettierised, although in practice it would be hard to achieve that.

Second, it seems that NMF entails none of LF, LMF, PF, EMF and FKF. We can envisage that one *infallibly* knows something in the logical sense or the probability sense but can still fallibly know the same thing in the non-maximal sense. For example, if one knows that Socrates' wife is a shrew on the basis of entailing evidence (e.g., that he remembers this), this is still compatible with his strength of epistemic position being imperfect. For there can always be some way to make one's epistemic position even stronger, such as gathering confirmation from others or auditory and visual information.

NMF is not prone to the typical problems affecting other stronger versions of fallibilism. NMF doesn't have problems with necessary truths, for it implies that one's knowledge of mathematical truths gained through testimony counts as fallible in the sense that one's epistemic position about that proposition can still be improved. NMF is also compatible with disjunctivism and Williamsonian epistemology, for even truth-entailing evidence doesn't imply that one's epistemic position cannot be stronger.

NMF seems to be such a weak position that all non-sceptics about knowledge can accept it. However, whether this position is so weak will also depend on how one conceives the notion of maximal justification. If by maximal justification, we mean justification supported by all the possible evidence the agent can actually gather given one's current circumstances, then it is not clear that NMF is weaker than other forms of fallibilism. For example, in many cases of knowledge we acquired in the past (e.g., I was standing on that precise spot ten years ago), the justification for the known proposition is maximal. This is because all possible evidence has already been gathered, no matter how weak that evidence is. According to NMF, such knowledge would count as infallible even if grounded on relatively weak evidence. By contrast, if by maximal justification we mean the justification that an ideal agent who has omniscient cognitive power such as God can achieve, then NMF may still be conceived as the weakest form of fallibilism.¹⁵ Following Fantl and McGrath (2009a, 2009c), I assume that NMF is the weakest form of fallibilism.

¹⁵ Thanks to Davide Fassio for discussion on this point.

7.1.6 Fallibilism and the knowledge norm of practical reasoning

So far, I have examined several versions of fallibilism, some of which are more plausible than others. In addition, I have tentatively endorsed a version of PF. The question now is how fallibilism fares with the knowledge norm of practical reasoning. There are reasons to think that fallibilism is not consistent with the knowledge norm, especially with the sufficiency version of this norm, SUFF. Counterexamples to SUFF reviewed in chapter 1 are advanced on the basis of an either explicit or implicit commitment to fallibilism. As Reed says:

If fallibilism is true and our knowledge is grounded in something less than certainty, it is never a given for us that p , even when we know it is true. So we can take ourselves to know that p while still recognising that there is a chance that it is false that p . When that chance carries with it very bad consequences if it were to become actual [...] it may be most rational *not* to act as if p . (Reed 2010, p. 229)

Likewise, the possibly weakest form of fallibilism NMF doesn't square with the knowledge norm either. A practical situation in which stakes are very high might call for a higher epistemic standard surpassing the epistemic standard that is sufficient for knowledge.

Fantl and McGrath (2009a, 2009b) show that PF or EMF is compatible with SUFF if we endorse a pragmatist account where the threshold on how probable p must be for you to know that p varies depending on the practical circumstances. In their view, "your probability for p is knowledge-level *iff* the probability that not- p doesn't stand in the way of p 's being put to work as a basis for belief and action." (2009a, p. 65, 2009b, p. 26). This allows them to reply to philosophers like Reed that the subject loses knowledge that p in cases in which one's epistemic position is not sufficient for acting on p , given that the relevant chance of error stands in the way of the proposition's being used as a basis for action. However, this view commits to a form of pragmatic encroachment, and thus is incompatible with moderate invariantism. Their view will be discussed in

more detail later (section 7.3.2).

Although some fallibilists have provided counterexamples to the knowledge norm of practical reasoning, they must recognise that the norm has a strong intuitive appeal.¹⁶ Therefore fallibilists who deny the norm need to explain its intuitive appeal in some other way. More specifically, they need to explain why ‘know’ and its cognates are the prominent terms used in epistemic assessment of action.

In sum, while fallibilism seems to be highly plausible, it is affected by at least two main problems. Fallibilism seems to conflict with two types of ordinary intuitions: i) intuitions about how we ordinarily assess our actions and reasoning based on knowledge; and ii) the infallibilist intuition put forward by Lewis, according to which knowledge that *p* is incompatible with the possibility of being wrong. In the next section, I shall introduce and criticise two well-known fallibilist accounts of the first type of intuition. My own view is that these two intuitions shouldn’t be considered separately. On the contrary, these two intuitions are related. Moreover, the explanation of one of them is related with the explanation of the other. In section 7.3, I will argue for such a joint explanation. More precisely, I shall argue that an explanation of the second type of intuitions also provides the basis for an explanation of the first type.

7.2 Threshold makers and communicative heuristic accounts

Recall that, as we explained in section 1.2, the fact that ‘knowledge’ and its cognates are the terms most frequently employed in epistemic assessments of rational action is often used to motivate the idea that knowledge is the epistemic norm for appropriate actions and deliberations. The intuitive data suggest the following thesis, formulated by Gerken (2015, 2017):

Prominence of ‘knowledge’

In normal cases of epistemic assessment of action, ordinary

¹⁶ Indeed many of them do recognize the appeal of the norm. See for example McGlynn (2014), Gerken (2015, 2017).

speakers frequently use the term ‘knowledge’ and its cognates.

This poses an explanatory challenge for opponents of the knowledge norm of practical reasoning. If knowledge does not govern rational action, why is it the prominent epistemic term in folk talks? Moreover, as we saw in the previous section, this fact also poses a challenge to the fallibilist. According to the latter, knowledge marks a quite high but fallible threshold of justificatory support, one that might not be sufficient to rationalize one’s actions and decisions when stakes are really high. Thus, according to this view, in our epistemic assessments of action and practical reasoning, we shouldn’t privilege knowledge over other epistemic conditions.

Two fallibilists who tried to deal with the challenge have been Reed (2013) and Gerken (2015, 2017). Both of them hold that the term ‘knowledge’ is frequently used to refer to a sort of threshold—it conveys that the subject has at least the degree of warrant sufficient to underwrite rational action. They both endorse certain kinds of conversational pragmatic accounts of knowledge ascriptions, though their accounts differ in certain details.⁷² The general idea is that a relevant knowledge ascription in epistemic assessments pragmatically (as opposed to constitutively) implies that the subject’s knowledge is well-grounded enough to make the action in question rational. Likewise, a relevant denial of knowledge in epistemic assessments pragmatically implies that the subject’s knowledge is not enough well-grounded to make the action in question rational. According to their views, the implicature holds in virtue of Grice’s Cooperative Principle (Grice, 1975), which says that one should make one’s contribution to conversation as relevant and useful as possible. When it matters to the conversation whether it is epistemically good enough to rely on a certain belief in practical reasoning, the Cooperative Principle requires someone who has a grasp of the subject’s epistemic position to provide informative judgment. The question is then why it has to be ‘knowledge’.

Gerken (2017, section 8.3.c) provides an account for why knowledge

⁷² Reed’s account is purely pragmatic, whereas Gerken’s account is a combination of psychological and pragmatic components, hence psycholinguistic.

rather than other epistemic terms such as justified or warranted belief functions as a *communicative heuristic* in folk epistemology. On Gerken's view, the concept of knowledge is deployed by default in cognitive heuristics in forming epistemic judgments. The word 'knowledge' then inherits this default status in communication given that communication itself is a cognitive task. In his view, knowledge ascriptions provide a good trade-off between accuracy and communicative effectiveness. 'Knowledge' appears to be the epistemic term that makes the communication effective and sufficiently informative. By contrast, using more accurate epistemic vocabulary instead of 'knowledge' risks being ineffective in communication, given that the more accurate terms could be out of the hearer's reach.

This assumption is reinforced by ontogenetic considerations: human beings acquire a basic competence of the word 'know' very early (see also Nagel (2013) and McGlynn (2016) for critical examinations of this claim). This, in turn, may be partly explained by the fact that a grasp of factive terms requires less cognitive competence than a grasp of non-factive terms such as 'belief' and 'justification'.

Cancellability is often taken to be one important mark, albeit a defeasible one, of pragmatic implicature (Blome-Tillman, 2008). And it seems that the cancellation can be made felicitously in abnormal cases of assessment. For example, in Brown's surgeon case (see section 1.6) the nurse might felicitously assert that "[the doctor] knows that it is the left kidney that is diseased. But I do not mean to say that she can just go ahead and operate before double-checking the records" (Gerken, 2015, p. 14).¹⁸

However, Reed and Gerken's accounts are not complete, for they cannot address all the relevant data concerning the use of knowledge ascriptions. In particular, their accounts face the challenge of explaining why knowledge ascriptions have more illocutionary force than ascriptions of justified or warranted belief. Imagine that in a normal case (low stakes, no urgency, etc.) in which one is in an epistemic position good enough

¹⁸ As a defender of the warrant account of practical reasoning, Gerken has an additional package of explanations based on the warrant account for the prominence of 'knowledge' and why the warrant account is explanatory better than the knowledge account. See Gerken (2015, 2017, chapter 6).

to recommend relying on p as a premise in practical reasoning, one says, “I am justified enough to believe that p ” rather than “I know that p ”. The hearer would be reasonably more hesitant to take for granted the recommendation of acting on p . The reason seems that justified belief doesn’t provide a guarantee of the truth of p , while one would definitely prefer to be assured of the truth of p for future actions and deliberations. In other words, assertions of “ S is justified enough to believe that p ” leaves open the possibility that p might be wrong, whereas the corresponding knowledge ascription seems to exclude that possibility. This is precisely the intuition behind Lewis’s quote in section 1.3; intuitively, it sounds wrong to say that one knows that p but p might not be the case. In this regard, the knowledge norm of practical reasoning is in a better position than its competing views to accommodate the fact that knowledge, rather than other doxastic attitudes falling short of knowledge, is taken to provide a *sufficient* condition for rational action in most ordinary conversational practices.

If ‘knowledge’ works as a threshold for the degree of warrant required for rational action, why can saying “ S is justified/warranted enough to believe that p ” not be as effective as saying “ S knows that p ” in recommending action?¹⁹ Reed and Gerken may resort to the factivity of knowledge to explain the extra illocutionary force of knowledge ascription, but that would be a further additional element in their explanation, different from their original accounts. This would undermine the uniformity of their accounts.²⁰ It’s preferable to have some account that can uniformly address the prominence of ‘knowledge’ and its cognates in epistemic assessments and the extra illocutionary force of knowledge ascriptions, if this account is available. In the next section, I will propose

¹⁹ Of course, on the assumption that the hearer has a good grasp of the terms ‘justified’ and ‘warranted’.

²⁰ The above problem doesn’t undermine the threshold maker and communicative heuristic account for the necessity claim of the knowledge norm. There are reasons to think that the data motivating the necessity claim are less robust. In cases of criticisms to one’s action, it is perfectly natural to substitute ‘know’ in complaints with other epistemic terms, such as ‘certain’, ‘have good reason’ or ‘justified believe’. Furthermore, even if the subject’s belief turns out to be false, if it is well justified there seems to be no good reason to criticise her action. See Gerken (2017, chapter 3) for more details.

such an account.

7.3 The infallibilist intuition and the knowledge norm of practical reasoning

In this section I will suggest an original account of i) the ordinary epistemic assessments commonly used to support the knowledge norm of practical reasoning and ii) the infallibilist intuition according to which knowledge that p is incompatible with error possibilities. I will also illustrate how this account can accommodate a number of other data and show a range of advantages with respect to alternative accounts.

I will start considering a standard infallibilist explanation of the sufficiency claim of the knowledge norm of practical reasoning SUFF. After examining some prominent fallibilist accounts of the infallibilist intuition and considering problems with these accounts, I will propose a mixed psychological-pragmatic account. The account relies on credal pragmatism—more in particular, on some constitutive dispositional properties of occurrent belief. In the end, I will show how the same account can explain the intuitiveness of the infallibilist intuition and the ordinary epistemic assessments of action and practical reasoning commonly used to support the knowledge norm of practical reasoning.

7.3.1 The infallibilist intuition and SUFF

As briefly mentioned in section 7.1.2, concessive knowledge attributions (CKAs henceforth) having the form “ S knows that p , but [however, yet, although, etc.] it’s possible [maybe, perhaps, there is a chance, etc.] that q (where q entails not- p)” often strike as absurd. Here are some examples:

- a) I know that Harry is a zebra, but it’s possible that Harry is a painted mule.
- b) John knows that Harry is a zebra, but it’s possible that Harry is a painted mule.
- c) John knows that Harry is a zebra, but it’s possible for John that Harry is a painted mule. (Stanley, 2005, p. 126)

The oddity of CKAs is recognized and broadly agreed upon among philosophers. For example, Lewis writes:

If you claim that *S* knows that *p*, and yet you grant that *S* cannot eliminate a certain possibility in which not-*p*, it certainly seems as if you have granted that *S* does not after all know that *p*. To speak of fallible knowledge, of knowledge despite uneliminated possibility of error, just sounds contradictory. (Lewis, 1996, p. 549)

The infelicity of CKAs reveals an *infallibilist intuition* about knowledge in our ordinary talks, according to which if one knows that *p*, then there is no possibility that not-*p*.

Recognising the existence of this infallibilist intuition helps to explain the intuitiveness of SUFF. The deduction from the infallibilist intuition to SUFF can be put as follows:

- I) If *S* knows that *p*, then there is no possibility for *S* that not-*p*;
- II) If there is no possibility for *S* that not-*p*, then it is appropriate for *S* to rely on *p* as a reason in her practical reasoning;
- C) Therefore, if *S* knows that *p*, it is appropriate for *S* to rely on *p* as a reason in her practical reasoning.

No matter whether the above reasoning is sound (in section 7.3.4 I will argue that it isn't), it looks like a very plausible argument whose premises are supported by ordinary epistemological intuitions. I've already discussed the intuitive appeal of I), i.e. the infallibilist intuition. II) is highly appealing as well. If there is no possibility for the subject that not-*p*, then it seems that nothing should hinder the subject from using *p* as a reason in practical reasoning. I will come back to this assumption in due time in section 7.3.4.

The infallibilist intuition also explains the stronger illocutionary force of knowledge ascriptions over ascriptions of other epistemic properties such as justified or warranted belief. This is because a parallel infallibilist

intuition does not hold for these epistemic properties. Intuitively, ascriptions of justified or warranted belief are compatible with error possibilities. Unlike CKAs, ordinary assertions equivalent to concessive justified belief attributions often sound felicitous. Consider the following examples:

- i) I have good reasons/evidence to believe that Harry is a zebra, but Harry might be a painted mule
- ii) Mary's belief that Amsterdam is the capital of the Netherlands is more than reasonable; still, she might be wrong.

Since the infallibilist intuition seems to obtain only for knowledge (or notions entailing knowledge), for other epistemic conditions, an equivalent of the argument I)–C) is not available.

While there is a path of reasoning leading from the infallibilist intuition and other intuitively plausible assumptions to SUFF, this intuition is arguably false. This intuition is inconsistent with the sort of fallibilism, i.e., epistemic modal fallibilism and probability fallibilism, which I endorsed and defended in section 7.1.2. As I argued there, the most plausible form of fallibilism entails that a subject's knowledge is compatible with there being an epistemic possibility or probability for that subject that not-*p*. That the intuition is false is also shown by contexts in which it sounds perfectly felicitous to assert that someone knows that *p*, but there is a chance for her that she is wrong. I considered some examples in section 7.2. For example, in Brown's surgeon case the nurse might felicitously assert, "She knows that it is the left kidney that is diseased. But I do not mean to say that she can just go ahead and operate before double-checking the records." (Gerken, 2015, p. 14; Worsnip, 2015)

I have shown how the infallibilist intuition, in combination with another seemingly plausible assumption II, can explain the intuitive appeal of SUFF and the stronger illocutionary force of knowledge ascriptions over ascriptions of other epistemic properties such as justified and warranted belief. However, I have also argued that the infallibilist intuition is false: knowledge is compatible with an epistemic possibility of error. An error theory of the intuition is needed—where, again, 'error theory' is here broadly conceived (as a theory that postulates false intuitions as opposed

to postulating specific performance error). Examples are pragmatic accounts (e.g., in terms of implicatures) and psychological accounts (e.g., in terms of biases). In the next section, I will examine some prominent pragmatic accounts of the infallibilist intuition and show that they are problematic. In section 7.3.3, I will then provide my own account.

7.3.2 Pragmatic accounts of the infallibilist intuition

The infallibilist intuition suggests an infallibilist folk epistemological principle. Let's call the corresponding version of infallibilism *epistemic modal infallibilism* (EMI). According to EMI, if *S* knows that *p*, the epistemic probability of *p* is 1. Proponents of EMI suggest that we should take as literally true the common-sense intuitions about the connection between epistemic possibility and knowledge as reflected in CKAs. CKAs are infelicitous simply due to their falsity (Hawthorne, 2012; Dodd, 2009, 2011; Stanley, 2005b). However, for fallibilists, CKAs rightly express the fallibilist idea that knowledge is compatible with the existence of error possibilities. According to these philosophers, CKAs are infelicitous for non-semantic reasons (Anderson, 2014; Dougherty & Rysiew, 2009; Fantl & McGrath, 2009a, 2009c; Reed, 2010, 2013; Rysiew, 2001).

Rysiew (2001) and Dougherty and Rysiew (2009)'s pragmatic account is a prominent non-semantic account of the oddity of CKAs. According to Dougherty and Rysiew, the oddity of CKAs is due to the incoherence that would arise in communication. In particular, CKAs violate general conversational rules, such as the Gricean's Co-operative Principle. According to their diagnosis:

[E]ither 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 [Co-operative Principle], 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. (Dougherty & Rysiew, 2009, pp. 128–129, *Italics added*)

According to Dougherty and Rysiew, by uttering a CKA, the speaker is representing herself as being ‘of two minds’ on the issue of whether p . If the possibility of not- p were not significant for the speaker, why would the speaker bother to mention it? By uttering that not- p is possible, according to Dougherty and Rysiew, one conveys that one has some real grounds for supposing not- p might be the case and/or that one isn’t confident that p . And from that, the hearer may infer that the speaker doesn’t take himself to know either that p or that not- p . This would clash with the first part of a CKA.

For Dougherty and Rysiew (ms), an error possibility can be significant and hence could be regarded as worth mentioning, due to different kinds of reasons—moral, practical, epistemic, aesthetic, and gustatory. The kind of relevance of the error possibility in conversation must be supplied by the context of utterance. However, they also acknowledge that without specification of the kind of relevance, the hearer is prone to understand the type of relevance as epistemic, given that the possibility of error is mentioned following upon a knowledge claim (ms, p. 5). Against pragmatic encroachment, Dougherty and Rysiew deny that error possibilities that become significant for merely practical reasons can threaten knowledge. Epistemic significance is the only sort of significance relevant for knowledge. In their view, an error possibility can be practically significant but might not be epistemically significant. In particular, they conceive epistemic significance to be connected with liability to epistemic defeat: the greater the liability to defeat, the greater the epistemic significance.²¹

Fantl and McGrath endorse a similar pragmatic account, but they disagree with Dougherty and Rysiew on what counts as an epistemically significant possibility of error. As defenders of pragmatic encroachment, Fantl and McGrath (2009a, 2009c) propose a pragmatic, interest-relative account of significance, whereby a chance of error “is significant just in

²¹ In their view, an epistemically significant error possibility ‘may well’ prevent one from knowing, even though it does not necessarily prevent it. By acknowledging the benign vagueness of ‘know(s)’ in natural language and the existence of borderline cases of knowledge, Dougherty and Rysiew (ms) deny that there is a precise cut-off point of how probable a proposition needs to be for one to be in a position to know. On their view, an epistemically significant error possibility flags a threat to knowledge, but it might be difficult to tell whether it truly defeats knowledge.

case it is high enough to make it improper to put p to work as a basis not only for belief, but...for action as well" (2009a, p. 65/2009c, p. 25). On their view, when an error possibility is practically significant, it is also epistemically significant and constitutes a knowledge-defeater.

Both the positions of Dougherty and Rysiew and Fantl and McGrath face certain problems. On Fantl and McGrath's view, a practically significant error possibility is also epistemically significant and knowledge-defeating. However, as argued by Dougherty and Rysiew, an error possibility can be practically significant without being epistemically significant.²² It is plausible for one to separate the epistemic significance and the practical significance in conversation. If the context specifies that an error possibility that q (where q entails not- p) is only significant in a practical sense, then one can admit the error possibility of q while also felicitously ascribe knowledge that p to someone. For example, in the surgeon case, the nurse might well ascribe knowledge to the surgeon as follows: "Of course, she knows which kidney it is. But imagine what it would be like if she removed the wrong kidney. She cannot afford the possibility that she might remove the wrong kidney." (Brown, 2008a; Gerken 2015, p. 15, 2017, chapter 6) Such cases of felicitous CKAs show that it is plausible to separate the practical significance and epistemic significance, and so it constitutes an objection to Fantl and McGrath's view.²³ Here I assume the validity and relevance of the intuitive judgment in Brown's surgeon case. I will come back to possible explanations of this case in section 7.3.4.

Against Dougherty and Rysiew's approach, Dodd (2011) has argued that their view cannot explain the following two facts:

²² It is worth acknowledging that a form of pragmatic encroachment weaker than that defended by Fantl and McGrath may only require that a practical factor *can* be epistemically significant.

²³ As I said in section 1.5, some philosophers challenge the intuitiveness of cases such as the surgeon case. For example, according to Neta (2009, pp. 697–698), it is mandatory for the surgeon to double-check the patient's records no matter her epistemic position, on pain of violating professional ethical requirements. So Neta would disagree on the fact that the surgeon should check because she cannot afford the possibility that she might remove the wrong kidney.

ENTAILMENT

When a speaker *S* thinks that a proposition *p* is epistemically possible for her, *S* will agree (if asked) that for all she knows, *p* is true—that *p* is consistent with her knowledge.

INFELICITY

It's infelicitous for a speaker to say "*p* might be true, but I'm not willing to say that for all I know, *p* is true".

Semantic explanations of ENTAILMENT and INFELICITY provide support to EMI. If the fact that *p* is epistemically possible for *S* implies that *p* is consistent with *S*'s knowledge, then what is epistemically possible for her is determined in large part by what is known. Concerning INFELICITY, in asserting "*p* might be true", *S* is saying that *p* is epistemically possible for her. If whether a proposition is epistemically possible for *S* is largely a function of what is known, then it would be odd for *S* being willing to affirm the epistemic possibility of *p*, without being willing to affirm that *p* is consistent with what she knows. Thus, EMI provides a natural explanation of ENTAILMENT and INFELICITY. By contrast, the fallibilist, and in particular the upholder of EMF, will find it difficult to provide a straightforward account of ENTAILMENT and INFELICITY, since they deny that being epistemically possible must be consistent with one's knowledge.

Recall that according to Dougherty and Rysiew's view of epistemic modals, *S*'s thought that *q* is epistemically possible entails that *q* is compatible with her evidence. But assuming that evidence is not simply the knowledge one has, the latter claim doesn't entail that *q* is compatible with all what *S* knows. So, in their view, in ENTAILMENT if *S* thinks that a proposition *q* is epistemically possible for her, it is not legitimate for her to agree (if asked) that for all she knows, *q* is true. Likewise, concerning the assertion mentioned in INFELICITY, Dougherty and Rysiew are committed to saying that it is legitimate, despite seemingly infelicitous, for *S* to say "*q* might be true, but I'm not willing to say that for all I know, *q* is true".

As pointed out by Dodd, Dougherty and Rysiew's pragmatic account for CKAs in terms of representing oneself as 'being of two minds' on

whether p does not help in explaining ENTAILMENT and INFELICITY. This is because ENTAILMENT and INFELICITY only include “ q might be true” or “For all I know, q ”, which is part of CKAs representing one’s doubt about the truth of p . It doesn’t include the other part, namely, the assertion that one knows that p , which represents one’s confident belief in p . It may be possible to give another separate pragmatic account for ENTAILMENT and INFELICITY, but then Dougherty and Rysiew won’t have a simple unified account of why these sentences are infelicitous. In the next subsection, I will propose a novel psychological account of the infallibilist intuition that avoids the problems that affect the accounts examined here.

7.3.3 A pragmatic-psychological account of CKAs and related data

A unified account that can explain ENTAILMENT, INFELICITY and the oddity of CKAs is called for. I will now suggest such an account. This account is partially psychological and partially pragmatic. Let me start with a sketchy analysis of the psychological status of someone who thinks that a proposition is epistemically possible. Cases in which one thinks a proposition is epistemically possible can be classified into two types. In type I) cases, one believes that p and realises the existence of the epistemic possibility that q (where q entails not- p) but doesn’t take it as knowledge undermining. Note that one does not need to be an epistemologist to do so. For example, a plumber can realise the existence of the epistemic possibility that all tubes are living beings but not take this possibility as undermining knowledge that they are not.²⁴

This type of case can be further divided into two subtypes. According

²⁴ Here I would like to flag that there is a substantive issue concerning what counts for a subject as ‘taking’ an error possibility as undermining knowledge. The notion of ‘taking’ should not be interpreted in a hyper-intellectualized way. Here taking an error possibility as knowledge undermining does not require possession of concepts such as those of ‘genuine alternative’ or ‘knowledge undermining’. This ‘taking’ is more like a feeling of incompatibility between a considered possibility and an attitude. I cannot provide a full account of this notion here. For the sake of exposition, I will stick to this terminology. For more on this issue, see Gerken (2011, 2013, chapter 2).

to one subtype I.a, one does not take the epistemic possibility that q as epistemically significant (i.e., one that is sufficient to undermine knowledge), but as one that is sufficiently practically relevant to be mentioned (an example is Brown's surgeon case). In the other kind of situation I.b, some error possibilities might be mentioned without infelicity for they cannot work as genuine knowledge defeaters because they are too far-fetched (for example, the brain-in-a-vat scenario for a lay person).²⁵ The central difference between the two types of case is that while in both types the subject can take there to be error possibilities which are not epistemically significant, in the latter case the error possibility is even not practically relevant, while in the former it is.

What ENTAILMENT and INFELICITY say does not seem to apply to those two kinds of situations (I.a and I.b). In those cases, one can felicitously admit the existence of error possibilities while deny that the error possibility is inconsistent with her knowledge. For example, in the surgeon case, the error possibility is taken to be insufficient to undermine knowledge but sufficiently practically relevant to be mentioned and considered in practical deliberation. The surgeon can well say, "The kidney that needs to be removed might be on the other side, that's why I have to double check, but I'm not willing to say that for all I know, it is true that the kidney might be on the other side."²⁶ Similarly, it may be felicitous to assert that I know that I am in Hangzhou even though there is a very tiny and abstract error possibility that I am a brain in a vat. Such error possibilities are so far-fetched that they don't deserve to be taken seriously to the extent of doubting any of our knowledge, even when mentioned (see also Blome-Tilmann (2009b, p. 247, 2014, p. 19), Williams (2001, p. 15)).

²⁵ The notion of far-fetchedness invoked is of course psychological. Following Gerken (2017, p. 145), whether a scenario counts as a far-fetched error possibility for a subject partly depends on whether the subject has previous exposure to it or whether it is novel or surprising. Individual and cultural variability in judgments of far-fetchedness should be expected. For more on the distinction between far-fetchedness and non-far-fetchedness, see Bach (2005), Dodd (2010), Dougherty and Rysiew (2009), Frances (2005), Levin (2008), Rysiew (2001) and Vogel (1999). In particular, on how lay people perceive sceptical scenarios, such as brain-in-a-vat or evil demon, see Pritchard (2001), MacFarlane (2005), Davis (2007, p. 436) and Adler (2012, p. 264).

²⁶ I am aware that it is very unlikely to have assertions like this in ordinary life.

CKAs stated in the above two types of situations don't strike me as so odd. Since ENTAILMENT, INFELICITY and oddity of CKAs do not hold in the types of context identified above, their universality is undermined.

Cases in which ENTAILMENT, INFELICITY and the oddity of CKAs apply are those in which the subject takes an error possibility to be knowledge undermining. Let us call these type II cases. When such an error possibility is mentioned in conversation, one conveys that the error possibility is epistemically significant, in the sense that it is psychologically salient and counts as a genuine alternative to the target proposition. The oddity of CKAs can then be explained in terms of the nature of occurrent belief. As argued in the previous chapter, having an occurrent belief that *p* constitutively involves settling one's mind about whether *p* (the question whether *p* is presently closed for the subject), and involves a series of dispositions such as being disposed to assert *p* and self-ascribe knowledge that *p*, and rely on *p* in one's practical reasoning. Now, in CKAs, when one self ascribes knowledge that *p*, one conveys that one occurrently believes that *p*. But by acknowledging the error possibility afterwards, one conveys that that error possibility is psychologically salient and counts as a genuine alternative to *p*. By doing this, she conveys that the question whether *p* is still open for her, and thus that she has not an occurrent belief that *p*.²⁷

A similar, though slightly more complex explanation can be given for the oddity of third-person CKAs. By ascribing to someone else knowledge that *p*, the speaker conveys that *p* is true and that is a settled question for her whether *p*. But by acknowledging an error possibility as epistemically significant for the person to whom the knowledge has just been ascribed ("she might be wrong"), it represents that error possibility as significant also for herself, conveying that the question whether *p* is open for her, and thus that she lacks occurrent belief.

In type II cases, ENTAILMENT and INFELICITY are true and are precisely what we should expect given the nature of occurrent belief.

²⁷ On this point, my view is similar to Dougherty and Rysiew's explanation of CKAs, appealing to the representation 'of being two minds'. A crucial difference is that my account is psychological rather than pragmatic. It appeals to the dispositional nature of occurrent belief, rather than to Gricean implicatures.

When the subject considers an error possibility epistemically significant, it is natural and epistemically rational to reopen the question and lose closed-mindedness about the relevant proposition, thereby losing occurrent belief in the target proposition. Hence, *in type II cases*, the subject won't take p to be part of her knowledge. Similarly, by asserting " p might be true, but I'm not willing to say that for all I know, p is true", one takes an error possibility as epistemically significant, conveying that she leaves open the question of whether p (i.e., to be open-minded about whether p), while at the same time denying willingness to claim to be open-minded about whether p . After all, by saying that "for all I know, p is true", one commits to the claim that p is consistent with what she knows, and so cannot exclude the possibility of p .

Recall that Dougherty and Rysiew's account does not have the resources to explain ENTAILMENT and INFELICITY. Their account only concerns what is manifested or implicated by positive knowledge ascriptions and the mention of p as an error possibility. By contrast, my account also explains what is going on in the psychology of the subject when she mentions p as an error possibility, and hence she takes that error possibility as epistemically significant. In my account, taking an error possibility as epistemically significant manifests a lack of cognitive closure and thus a lack of occurrent belief. This provides us with an explanation of ENTAILMENT and INFELICITY in addition to CKAs.

However, even though an epistemic significant error possibility can (and often does) defeat one's occurrent belief (together with the dispositions it involves), it might not defeat one's knowledge. In my view, whether the subject still knows the relevant proposition depends on whether she still has the relevant *dispositional belief*, which depends on the baseline level of NFC (or NTAC). As argued in chapter 6, it is plausible that, for a subject with high baseline NTAC, an entertained error possibility destroys both one's dispositional and occurrent belief.²⁸ But it is also plausible that, for a subject with high baseline NFC, the salience of error possibility doesn't

²⁸ This happens in cases in which the subject loses the dispositions to rely on the relevant proposition as a premise in reasoning also in normal circumstances. Whether this happens depends on how cautious the subject is and how easily it is for the subject to change her mind given the existence of counter-possibilities.

affect the dispositional belief at all. In such a case one can still have the corresponding knowledge, even without having an occurrent belief and the dispositions that go with it.

In short, my account of the infelicity of asserting the various claims considered above can be summarized as follows:

Account of the infelicity of CKAs and related data

A knowledge ascription that p conveys that one occurrently believes that p , where occurrent belief constitutively involves being closed-minded about whether p (the question whether p is presently closed for the subject). This rationally precludes taking the error possibility that not- p as epistemically significant in conversation—which would convey that the question of whether p is open for the subject.

‘Rationally’ here concerns consistency in what is asserted. One cannot coherently assert or convey in assertion that the question of whether p is both open and not open for her on pain of self-contradiction.

As it has been argued above, when a subject has occurrent belief that p , for her there are no epistemically significant error possibilities regarding p . So from the perspective of the subject who occurrently believes that p , there is either i) no error possibility incompatible with p or ii) there is an error possibility incompatible with p , but this error possibility is not significant—which means that this error possibility is insufficient to undermine any disposition constitutively involved in occurrent belief, such as dispositions to self-ascribe knowledge or reopen the question whether p . In the latter type of case ii), we can include cases in which one could recognize the existence of an error possibility but deny that it is inconsistent with her knowledge. This happens when, for example, the error possibility is taken to be too far-fetched, or the practical context requires an epistemic position stronger than knowledge (as in the Surgeon case). But in many other circumstances, recognizing the existence of an error possibility and at the same time ascribing knowledge will involve manifesting certain dispositions typical of someone who lacks occurrent belief and others typical of someone who possesses it, hence manifesting contradictory dispositions.

The presence of contradictory dispositions not only explains the oddity of CKAs, INFELICITY and ENTAILMENT, but also explains the absurdity of corresponding mental versions of CKAs, such as the absurdity of taking an error possibility as significant and at the same time taking oneself to know. This conjunction of attitudes is absurd because it involves manifesting the disposition to reopen the question of whether p (by taking an error possibility as epistemically significant) while at the same time manifesting a disposition to maintain the question closed (by taking oneself to know, and thus taking oneself to occurrently believe). Similar explanations are available for mental versions of INFELICITY and ENTAILMENT—circumstances in which the subject thinks about the relevant claims rather than asserting them.

7.3.4 An account of the infallibilist intuition and of ordinary epistemic assessments of action

The account considered in section 7.3.3 can also easily explain the two sets of intuitive judgments considered in previous sections, namely: i) the *infallibilist intuition* that knowledge that p is incompatible with error possibilities and ii) *ordinary epistemic assessments of action and practical reasoning* usually deployed to motivate NEC and SUFF.

If a subject occurrently believes that p , she will have several dispositions including a disposition not to reopen the question of whether p . Thereby, for the subject who occurrently believes that p , there will not be epistemically significant error possibilities, viz., possibilities which, amongst other things, could be able to reopen the question of whether p , and could undermine one's dispositions to rely on p in her reasoning. This intuition explains both why knowledge that p seems to be incompatible with error possibilities, and why in many cases we deem knowledge as a necessary and sufficient epistemic ground for rational action.

For what concerns i), thinking to relevant error possibilities would reopen the question of whether p , defeating the dispositions typical of occurrent belief, including self-ascribing knowledge and taking oneself to know. This explains the infallibilist intuition that knowledge is incompatible with error possibilities.

For what concerns ii), if someone has an occurrent belief, on the

one hand, she has the disposition to self-ascribe knowledge and to take herself to know. On the other hand, the question of whether p is (at least temporarily) closed for her, and she can make as if p were the case, reasoning and acting from it. Given these assumptions, it seems rational for one to act on p in practical reasoning when one takes oneself to know that p . Likewise, in normal circumstances, it seems unreasonable to take oneself to know that p (manifesting dispositions of one who has occurrent belief) while at the same time taking oneself as having insufficient grounds to rely on p in practical reasoning (manifesting the absence of dispositions typical of occurrent belief).²⁹ This explains the intuitiveness of SUFF. Similarly, for what concerns NEC, it seems unreasonable not to take oneself to know (manifesting the absence of dispositions typical of occurrent belief) while at the same time taking oneself as having rational grounds to fully rely on p in practical reasoning (manifesting dispositions of one who has occurrent belief).

Of course, all this is compatible with atypical cases in which one recognizes the existence of some abstract or far-fetched error possibility but denies its relevance to knowledge. In such cases, the error possibility doesn't stand as a reason capable of inhibiting one from acting on the relevant knowledge. In such cases the error possibility is insufficient to undermine the occurrent belief, and thus to destroy the disposition involved in occurrent belief to rely on the believed proposition. The explanation is also compatible with cases such as the surgeon case, in which one might agree that the error possibility is not epistemically significant, and thus it doesn't undermine the occurrent belief, but it is a possibility that must be taken into consideration in one's deliberation given abnormal practical circumstances, often due to specific prudential,

²⁹ Of course, in such cases the speaker should also have the disposition to ascribe justified or warranted belief. But recall that knowledge ascriptions have stronger illocutionary force than ascriptions of justified or warranted belief (see section 7.3.1). For that reason, knowledge ascription is more effective in communication given the relevant speaker's purposes. Thus, there is no surprise that folks often use knowledge ascriptions in assessing or directing their actions and decisions. And it is expectable that 'knowledge' and its cognates are the prominent epistemic terms we use in epistemic assessment of practical reasoning and action.

moral or social duties of the agent.³⁰

While the above account explains the intuitive appeal of taking knowledge as a necessary and sufficient condition for rational action, it also shows the limits of this idea. The connection that the account displays is not between knowledge and rational reliance on a proposition in practical reasoning, but between knowledge ascriptions that *p* and a rational disposition to rely on *p* in practical reasoning. This connection depends on the nature of occurrent belief, which involves both a disposition to (self-)ascribe knowledge and to rely on the believed proposition in practical reasoning. These two dispositions follow from the closed-mindedness constitutive of occurrent belief, which also explains the infallibilist intuition. However, since occurrent belief is neither necessary nor sufficient for knowledge (as I argued in chapter 6), it may well be that in some cases the subject knows but reasonably lacks a disposition to rely on the known proposition in practical reasoning, and vice versa. This happens, for example, in some high-stakes cases, in which the subject lacks occurrent belief, as well as all the dispositions occurrent belief constitutively involves (to ascribe knowledge, assert, rely on the known proposition in practical reasoning, not to reopen the question of whether *p*, etc.).

We are now in a position to re-examine the argument from the infallibilist intuition to SUFF discussed in section 7.3.1. The argument is unsound due to the falsity of premise I):

I) If *S* knows that *p*, then there is no possibility for *S* that not-*p*.

³⁰ Brown (2008a, 2008b), Reed (2010) and Gerken (2011, 2015, 2017) take such cases as showing that knowledge does not provide sufficient grounds to use the key proposition as a reason in one's practical reasoning. In my view, on the contrary, when the subject has an occurrent belief, for her the question of whether *p* is closed in favor of *p*. It seems thus reasonable for her to rely on *p* in her reasoning, even though the abnormal practical circumstances do not allow acting as if *p*. This is because even if the subject reasonably relies on the proposition as a *pro tanto* reason, this reason is defeated by other practical reasons (e.g., a professional duty to check if one's epistemic position is not extremely well grounded or doesn't include certain bits of evidence), despite the fact that for the subject the question whether *p* is closed.

Premise I), i.e. the infallibilist intuition, is false because knowledge that p is compatible with there being a possibility for S that not- p . The real (rational) incompatibility is only between taking oneself to know that p (manifesting occurrent belief that p) and leaving open the question of whether p (manifesting absence of occurrent belief that p)—where ‘rational’ here concerns once again consistency in what is thought or asserted.

Furthermore, the above account explains why, while I) is false, it appears to be true. The reason is that if S takes herself to know that p , then she is manifesting a typical disposition of occurrent belief. And occurrent belief also involves a disposition to keep closed the question whether p . Thus, it looks to S as if there is no possibility that not- p .

The account also explains the intuitive appeal of the second premise in the argument for SUFF considered in section 7.3.1:

II) If there is no possibility for S that not- p , then it is appropriate for S to rely on p as a reason in her practical reasoning.

II) seems intuitively true because, as just said, the disposition (constitutive of occurrent belief) to hold closed the question of whether p makes it look to S as if there is no possibility for S that not- p .³¹ Moreover, another disposition constitutive of occurrent belief in addition to closure is the disposition to rely on p as a premise in her practical reasoning. Thus, if the subject S who holds an occurrent belief that p takes there not to be a possibility that not- p , S will also consider it appropriate to rely on p as a reason in her practical reasoning (and inappropriate to do the opposite). This intuitive appeal of I) and II) explains why the argument from the infallibilist intuition to SUFF still seems to be sound and compelling.

These considerations provide us with an argument for the intuitive appeal of SUFF. The argument is as follows:

- I*) If S takes herself to know that p , then she doesn't take there to be any possibility that not- p ;
- II*) If S doesn't take there to be any possibility that not- p , then

³¹ I remain open here on whether II) is actually true or not.

it seems appropriate from her perspective to rely on p ;
 C*) Therefore, if S takes herself to know that p , then it seems
 appropriate from her perspective to rely on p .

While this can be seen as a further independent explanation of the intuitive appeal of SUFF in addition to the one considered at the beginning of this subsection, this argument relies on exactly the same assumptions about the dispositional nature of occurrent belief. As I see it, the argument constitutes another way of putting forward the same considerations about the nature of occurrent belief involved in other explanations considered in this section.

7.4 Concluding remarks

The fact that ‘knowledge’ and its cognates are used prominently in epistemic assessment has been used to motivate the knowledge norm of practical reasoning. Even if, as I have argued in chapter 3, the knowledge norm is false, the ordinary epistemic assessments used to motivate the norm still pose an explanatory challenge for fallibilist moderate invariantists. Based on the views and ideas defended in chapters 5 and 6, in this chapter I have developed a psychological account of the intuitive appeal of the knowledge norm of practical reasoning in terms of dispositions constitutive of occurrent belief. Moreover, I have argued that this account provides plausible explanations of a wide variety of data, including the infallibilist intuition, CKAs, ENTAILMENT and INFELICITY.

If we combine these results with those considered in the previous chapters—in particular, the ability of credal pragmatism to account for practical factor effects on knowledge ascriptions and on certain doxastic states—we can now appreciate how this view provides a broad and complete picture of the relations between doxastic attitudes, knowledge, knowledge ascriptions and practical matters. The resulting view combines epistemic modal fallibilism and moderate invariantism. At the same time, it is capable of explaining away the data apparently supporting epistemic modal infallibilism (as argued in this chapter) and the practical effects

on knowledge ascriptions (as argued in chapter 6). These various results provide further *prima facie* arguments for credal pragmatism. This view fares better than alternative moderate invariantist views, such as doxastic pragmatism and pragmatic accounts of knowledge ascriptions in terms of explanatory power and explanatory parsimony.

Concluding Remarks

In this book, I have argued against pragmatic encroachment on knowledge and defended a version of moderate invariantism. I first considered two types of arguments for pragmatic encroachment on knowledge. One appeals to practical factor effects on knowledge ascription; the other is supported by the knowledge norm of practical reasoning. I critically examined existing moderate invariantist responses to these arguments. These include pragmatic accounts (section 1.5) and doxastic accounts (chapter 4) of practical factor effects on knowledge ascription, objections to the knowledge norm of practical reasoning (chapter 1, section 6), and Williamson's moderate invariantist approach (chapter 2).

In other parts of the book, I developed new lines of responses to pragmatic encroachment on knowledge. In chapter 3, I argued that the knowledge norm of practical reasoning, as well as other epistemic norms of practical reasoning, face specific counterexamples. These counterexamples show that there is no general epistemic norm that can apply to every instance of practical reasoning, further undermining the knowledge norm and the arguments for pragmatic encroachment based on it.

In chapters 5–7, I defended a new type of doxastic pragmatism, credal pragmatism. The basic idea of this view is that our credences are systematically sensitive to practical factors. I also argued that such sensitivity manifests bounded rationality (chapter 5). In chapter 6, I considered a more elaborated version of the view which provides characterisations of two types of full belief: occurrent belief and dispositional belief. In this picture, occurrent belief is essential for providing premises in practical reasoning, as it involves (defeasible) dispositions to rely on p as a premise in reasoning and to assert that p in the actual circumstance. Occurrent belief depends on the actual degree of credence and is sensitive to practical factors' influences. By contrast, dispositional belief depends on credences in normal circumstances, where

non-truth-relevant factors do not interfere with the regulation of doxastic attitudes. This picture is compatible with moderate invariantism. This is because the doxastic attitude relevant for knowledge is dispositional belief, which is governed by purist rationality standards. Since dispositional belief tends to be insensitive to circumstances involving abnormal psychological and practical factors, knowledge inherits such stability and insensitivity to contingent practical factors.

Credal pragmatism has a remarkable explanatory power. It can explain away aspects of the relation between knowledge and practical matters, practical factor effects on knowledge ascription, and the intuitive normative role of knowledge in practical reasoning. These features can be explained by specific dispositional properties constitutive of occurrent belief. In particular, chapter 7 showed that such an account can explain not only the intuitive appeal of the knowledge norm of practical reasoning, but also the infallibilist intuition and related data. If the arguments in this book are correct, credal pragmatism provides a comprehensive picture of the relations between doxastic attitudes, knowledge, knowledge ascriptions and practical matters. Those who, like myself, find pragmatic encroachment on knowledge unappealing, should consider credal pragmatism as a serious alternative to other forms of moderate invariantism.

Bibliography

- Adler, J. E. (2012). Contextualism and fallibility: pragmatic encroachment, possibility, and strength of epistemic position. *Synthese*, 188(2), 242–272.
- Alston, W. (1988). An internalist externalism. *Synthese*, 74(3), 265–283.
- Alston, W. (1992). Infallibility. In J. Dancy and E. Sosa (Eds.), *A Companion to Epistemology* (p. 206). Oxford: Blackwell.
- Alston, W. (1996). Belief, acceptance, and religious faith. In D. Howard-Snyder and Jordan (Eds.), *Faith, Freedom, and Rationality* (pp. 3–27). London: Rowman & Littlefield.
- Alter, A. L., Oppenheimer, D. M., Epley, N., & Eyre, R. N. (2007). Overcoming intuition: metacognitive difficulty activates analytic reasoning. *Journal of experimental psychology: General*, 136(4), 569–576.
- Anderson, C. (2014). Fallibilism and the flexibility of epistemic modals. *Philosophical Studies*, 167(3), 597–606.
- Anderson, C. (2015). On the intimate relationship of knowledge and action. *Episteme*, 12(3), 343–353.
- Annis, D. (1977). Knowledge, belief and rationality. *The Journal of Philosophy*, 74(4), 217–225.
- Arnold, A. (2013). Some evidence is false. *Australasian Journal of Philosophy*, 91(1), 165–172.
- Audi, R. (1998). *Epistemology*. New York: Routledge.
- Audi, R. (2008). Belief, faith, and acceptance. *International Journal for Philosophy of Religion*, 63(1/3), 87–102.
- Ayer, A. J. (1956). *The Problem of Knowledge*. London: Penguin.
- Bach, K. (2001). Speaking loosely: Sentence nonliterality. *Midwest Studies in Philosophy*, 25(1), 249–263.
- Bach, K. (2005). The emperor's new 'knows'. In G. Preyer and G. Peter (Eds.), *Contextualism in philosophy: knowledge, meaning, and truth* (pp. 51–89). Oxford: 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 J. K. Campbell, M. O'Rourke, and H. S. Silverstein (Eds.), *Knowledge and Skepticism* (pp. 105–36). Cambridge, MA: MIT Press.
- Barnes, J. (1982). The beliefs of a pyrrhonist. *Proceedings of the Cambridge Philological Society*, n.28(1982), 1–29.
- Baron, J., and Hershey, J. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54(4), 569–579.
- Baumann, P. (2011). Wams: Why worry? *Philosophical Papers*, 40(2), 155–177.
- Baumann, P. (2012). Knowledge, practical reasoning and action. *Logos and Episteme*, 3(1), 7–26.
- Benton, M. A. (2014). Knowledge norms. *Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/kn-norms/>. Accessed 4 October 2015
- Berker, S. (2008). Luminosity regained. *Philosophers' Imprint*, 8(2), 1–22.
- Bird, A. (1998). Dispositions and antidotes. *Philosophical Quarterly*, 48(191), 227–234.
- Bird, A. (2007). Justified judging. *Philosophy and Phenomenological Research*, 74(1), 81–110.
- Black, C. (1971). Knowledge without belief. *Analysis*, 31(5), 152–158.
- Black, T. (2005). Classic invariantism, relevance and warranted assertability manoeuvres. *Philosophical Quarterly*, 55(219), 328–336.
- Blome-Tillmann, M. (2008). Conversational implicature and the cancellability test. *Analysis*, 68(2), 156–60.
- Blome-Tillmann, M. (2009a). Contextualism, subject-sensitive invariantism, and the interaction of 'knowledge'-ascriptions with modal and temporal operators. *Philosophy and Phenomenological Research*, 79(2), 315–331.
- Blome-Tillmann, M. (2009b). Knowledge and presuppositions. *Mind*, 118(470), 241–294.
- Blome-Tillmann, M. (2013). Knowledge and implicatures. *Synthese*, 190(18), 4293–4319.
- Blome-Tillmann, M. (2014). *Knowledge and Presupposition*. Oxford: Oxford University Press.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA:

- Harvard University Press.
- BonJour, L. (1998). *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- BonJour, L. (2010). The myth of knowledge. *Philosophical Perspectives*, 24(1), 57–83.
- Boult, C. (2017). Epistemic normativity and the justification-excuse distinction. *Synthese*, 194(10), 4065–4081.
- Braithwaite, R. B. (1932–1933). The nature of believing. *Proceedings of the Aristotelian Society*, 33, 129–46.
- Bratman, M. (1992). Practical reasoning and acceptance in a context. *Mind*, 101(401), 1–16.
- Bröder, A., & Newell, B. R. (2008). Challenging some common beliefs: Empirical work within the adaptive toolbox metaphor. *Judgment and Decision Making*, 3(3), 205–214.
- Broome, J. (2013). *Rationality Through Reasoning*. Oxford: Wiley-Blackwell.
- Brown, J. (2005). Adapt or die: The death of invariantism? *Philosophical Quarterly*, 55(219), 263–285.
- Brown, J. (2006). Contextualism and warranted assertibility manoeuvres. *Philosophical Studies*, 130(3), 407–435.
- Brown, J. (2008a). Knowledge and practical reason. *Philosophy Compass*, 3(6), 1135–1152.
- Brown, J. (2008b). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167–189.
- Brown, J. (2010). Fallibilism and the knowledge norm for assertion and practical reasoning. In J. Brown and H. Cappelen (Eds.), *Assertion* (pp. 153–74). Oxford: Oxford University Press.
- Brown, J. (2012a). Practical reasoning, decision theory and anti-intellectualism. *Episteme*, 9(1), 43–62.
- Brown, J. (2012b). Impurism, practical reasoning, and the threshold problem. *Noûs*, (47)1, 179–192.
- Brown, J. (2013). Experiential philosophy, contextualism and SSI. *Philosophy and Phenomenological Research*, 86(2), 233–61.
- Buchak, L. (2014). *Risk and Rationality*. Oxford: Oxford University Press.
- Buckareff, A. (2004). Acceptance and deciding to believe. *Journal of Philosophical Research*, 29, 173–190.

- Buckwalter, W. (2010). Knowledge isn't closed on Saturday: a study in ordinary language. *Review of Philosophy and Psychology*, 1(3), 395–406.
- Buckwalter, W. (2014). The mystery of stakes and error in ascriber intuitions. In J. Beebe (Ed.), *Advances in Experimental Epistemology* (pp. 145–174). London: Bloomsbury.
- Buckwalter, W., and Schaffer, J. (2015). Knowledge, stakes, and mistakes. *Noûs*, 49(2), 201–234.
- Buckwalter, W., and Turri, J. (2017). Descartes's schism, Locke's reunion: completing the pragmatic turn in epistemology. *American Philosophical Quarterly*, 54(1), 25–46.
- Burge, T. (2003). Perceptual Entitlement. *Philosophy and Phenomenological Research*, 67(3), 503–548.
- Burge, T. (2010). *Origins of Objectivity*. Oxford University Press.
- Burnyeat, M., and Frede, M. (Eds.). (1997). *The Original Sceptics*. Indianapolis and Cambridge, Mass.: Hackett.
- Campbell, C. A. (1967). Towards a definition of belief. *The Philosophical Quarterly*, 17(68), 204–220.
- Campbell-Moore, C., and Salow, B. (2020). Avoiding risk and avoiding evidence. *Australasian Journal of Philosophy*, 98(3):495–515.
- Chaiken, S., and Trope, Y. (1999). *Dual-Process Theories in Social Psychology*. New York: Guilford Press.
- Chisholm, R. (1989a). Probability in the theory of knowledge. In M. Clay and K. Lehrer (Eds.), *Knowledge and Skepticism* (pp. 119–130). Boulder, CO: Westview.
- Chisholm, R. (1989b). *Theory of Knowledge* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.
- Choi, S. (2008). Dispositional properties and counterfactual conditionals. *Mind*, 117(468), 795–841.
- Choi, S., and Fara, M. (2014). Dispositions. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014.). <http://plato.stanford.edu/archives/spr2014/entries/dispositions/>. Accessed 23 July 2015.
- Christensen, D. (2001). Preference-based arguments for probabilism. *Philosophy of Science*, 68(3), 356–376.
- Christensen, D. (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press.
- Clarke, R. (2013). Belief is credence one (in Context). *Philosopher's*

- Imprint*, 13 (11), 1–18.
- Cohen, J. (1989). Belief and acceptance. *Mind*, 98(391), 367–389.
- Cohen, J. (1992). *An essay on belief and acceptance*. New York: Clarendon Press.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2(Epistemology), 91–123.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13(s13), 57–89.
- Comesaña, J., and Kantin, H. (2010). Is evidence knowledge? *Philosophy and Phenomenological Research*, 80(2), 447–454.
- Conee, E., and Feldman, R. (2004). Making sense of skepticism. In E. Conee and R. Feldman (Eds.), *Evidentialism: Essays in Epistemology*, (pp. 277. nt). Oxford: Oxford University Press.
- Cresto, E. (2010). Belief and contextual acceptance. *Synthese*, 177(1), 41–66.
- Crisp, T. M. (2005). Hawthorne on knowledge and practical reasoning. *Analysis*, 65(286), 138–140.
- Daly, C. J. (2008). Fictionalism and the attitudes. *Philosophical Studies*, 139(3), 423–440.
- Dancy, J. (2004). *Ethics Without Principles*. New York: Oxford University Press.
- Davis, W. A. (2007). Knowledge claims and context: loose use. *Philosophical Studies*, 132(3), 395–438.
- de Finetti, B. (1990). *Theory of probability* (first edition 1974., Vol. I). New York: Wiley.
- de Sousa, R. B. (1971). How to give a piece of your mind: or, the logic of belief and assent. *Review of Metaphysics*, 25(1), 52–79.
- Dennett, D. (1978). How to change your mind. In *Brainstorms. Philosophical Essays on Mind and Psychology* (pp. 300–209). Montgomery, Vermont: Bradford Books.
- Derksen, A. A. (1978). The alleged lottery paradox resolved. *American Philosophical Quarterly*, 15(1), 67–74.
- DeRose, K. (1991). Epistemic possibilities. *The Philosophical Review*, 100(4), 581–605.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.

- DeRose, K. (1999). Contextualism: An explanation and defense. In J. Greco and E. Sosa (Eds.), *The Blackwell Guide to Epistemology* (pp. 187–205). Malden, MA: Blackwell Publishers.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Issues*, 111(2), 167–203.
- DeRose, K. (2009). *The Case for Contextualism*. New York: Oxford University Press.
- DeRose, K. (2011). Contextualism, contrastivism, and X-Phi surveys. *Philosophical Studies*, 156(1), 81–110.
- Dimmock, p. (2019). Knowledge, belief, and egocentric bias. *Synthese*, 196(8), 3409–3432.
- Dimmock, P., and Huvenes, T. T. (2014). Knowledge, conservatism, and pragmatics. *Synthese*, 191(14), 3239–3269.
- Dinges, A. (2016). Epistemic invariantism and contextualist intuitions. *Episteme*, 13(2), 219–232.
- Dinges, A. (2021). Knowledge and non-traditional factors. Prospects for doxastic accounts. *Synthese*, 198, 8267–8288.
- Dinges, A., & Zakkou, J. (2020). Much at stake in knowledge. *Mind and Language*, 36(5), 729–749.
- Dodd, D. (2010). Confusion about concessive knowledge attributions. *Synthese*, 172(3), 381–396.
- Dodd, D. (2011). Against fallibilism. *Australasian Journal of Philosophy*, 89(4), 665–685.
- Dodd, D. (2017). Belief and Certainty. *Synthese*, 194, 4597–4621.
- Dougherty, T., and Rysiew, p. (2009). Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research*, 78(1), 128–32.
- Dougherty, T., and Rysiew, p. (2011). Clarity about concessive knowledge attributions: reply to Dodd. *Synthese*, 181(3), 395–403.
- Dougherty, T., and Rysiew, p. (ms). *Pragmatics without pragmatism: reply to Fantl & McGrath*.
- Douven, I., and Schupbach, J. N. (2015). The role of explanatory considerations in updating. *Cognition*, 142(Supplement C), 299–311.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67(24), 1007–1023.
- Dub, R. (2015). Delusions, acceptances, and cognitive feelings. *Philosophy*

- and *Phenomenological Research*, 91(3), 723–759.
- Durant, J. (forthcoming). Knowledge-based decision theory and the new evil demon. In F. Dorsch and J. Durant (Eds.), *The New Evil Demon*. Oxford University Press.
- Earman, J. (1992). *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory*. MIT Press.
- Egan, A. (2008). Seeing and believing: perception, belief formation and the divided mind. *Philosophical Studies*, 140(1), 47–63.
- Eriksson, L., & Rabinowicz, W. (2013). The interference problem for the betting interpretation of degrees of belief. *Synthese*, 190(5), 809–830.
- Ekins, B. R., Rollins, D. E., Duffy, D. P., and Gregory, M. C. (1985). Standardized treatment of severe methanol poisoning with ethanol and hemodialysis. *Western Journal of Medicine*, 142(3), 337.
- Eklund, M. (2015). Fictionalism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2015). <http://plato.stanford.edu/archives/win2015/entries/fictionalism/>. Accessed 16 November 2015.
- Engel, P. (1999). Dispositional belief, assent, and acceptance. *Dialectica*, 53(3–4), 211–226.
- Engel, P. (2013). Doxastic correctness. *Aristotelian Society Supplementary Volume*, 87(1), 199–216.
- Eriksson, L., and Hájek, A. (2007). What are degrees of belief? *Studia Logica*, 86(2), 185–215.
- Evans, J. (2007). *Hypothetical Thinking: Dual Processes in Reasoning and Judgment*. New York: Psychology Press.
- Evans, J., and Frankish, K. (Eds.). (2009). *In Two Minds*. Oxford: Oxford University Press.
- Epley, N., & Gilovich, T. (2005). When effortful thinking influences judgmental anchoring: Differential effects of forewarning and incentives on self-generated and externally provided anchors. *Journal of Behavioral Decision Making*, 18(3), 199–212.
- Fantl, J., and McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111(1), 67–94.
- Fantl, J., and McGrath, M. (2007). On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research*, 75(3), 558–589.
- Fantl, J., and McGrath, M. (2009a). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

- Fantl, J., and McGrath, M. (2009b). Critical study of John Hawthorne's Knowledge and Lotteries and Jason Stanley's Knowledge and Practical Interests. *Noûs*, 43(1), 178–192.
- Fantl, J., and McGrath, M. (2009c). Advice for fallibilists: put knowledge to work. *Philosophical Studies*, 142(1), 55–66.
- Fantl, J., and McGrath, M. (2012). Arguing for shift epistemology. In J. Brown and M. Gerken (Eds.), *Knowledge Ascriptions* (pp. 55–74). Oxford: Oxford University Press.
- Fantl, J., & McGrath, M. (2017). Clarifying pragmatic encroachment: a reply to charity Anderson and John Hawthorne on knowledge, practical adequacy, and stakes. In T. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 6, pp. 258–266). Oxford: Oxford University Press.
- Fassio, D. (2015). The aim of belief. *Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/beli-aim/>. Accessed 2 July 2016.
- Fassio, D. (2020). Moderate skeptical invariantism. *Erkenntnis*, 85(4), 841–870.
- Fassio, D. (2021). A (partial) defense of moderate skeptical invariantism. In C. Kyriacou & K. Wallbridge (Eds.), *Skeptical Invariantism Reconsidered* (pp. 154–171). New York: Routledge.
- Fassio, D., and Gao, J. (2020). Belief, credence and statistical evidence. *Theoria* 86 (4):500–527.
- Fassio, D., and Gao, J. (2021). Do we really need a knowledge-based decision theory? *Synthese*, 199(3-4), 7031–7059.
- Feltz, A., and Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical Psychology*, 23(5), 683–706.
- Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Fogelin, R. (1994). *Pyrrhonian Reflections on Knowledge and Justification*. Oxford: Oxford University Press.
- Francis, K., Beaman, P., and Hansen, N.(2019). Stakes, scales, and skepticism. *Ergo*, 6(16), 427–487.
- Frances, B. (2005). When a skeptical hypothesis is live. *Noûs*, 39(4), 559–595.
- Frankish, K. (2004). *Mind and Supermind*. Cambridge: Cambridge University Press.

- Freund, T., Kruglanski, A. W., and Schpitzajzen. (1985). The freezing and unfreezing of impressional primacy: Effects of the need for structure and the fear of invalidity. *Personality and Social Psychology Bulletin*, 11, 479–487.
- Fumerton, R. (1995). *Metaepistemology and Scepticism*. Lanham, MD: Rowman and Littlefield.
- Fumerton, R. (2004). Epistemic probability. *Philosophical Issues*, 14(1), 149–164.
- Ganson, D. (2008). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139(3), 441–458.
- Ganson, D. (2019). Great expectations: belief and the case for pragmatic encroachment. In B. Kim & M. McGrath (Eds.), *Pragmatic Encroachment in Epistemology* (pp. 10–34). New York, Oxon: Routledge.
- Gao, J. (2015). Advances in experimental epistemology. *International Studies in the Philosophy of Science*, 29(1), 101–105.
- Gao, J. (2016). *Belief, Knowledge and Action*. The University of Edinburgh.
- Gao, J. (2017). Rational Action Without Knowledge (and Vice Versa). *Synthese*, 194(6), 1901–1917.
- Gao, J. (2019a). Against the Iterated Knowledge Account of High-Stakes Cases. *Episteme*, 16(1), 92–107.
- Gao, J. (2019b). Credal pragmatism. *Philosophical Studies*, 176(6), 1595–1617.
- Gao, J. (2021). Self-deception and pragmatic encroachment: A dilemma for epistemic rationality. *Ratio*, 34(1), 20–32.
- Gao, J. (2023). Should credence be sensitive to practical factors? A cost-benefit analysis. *Mind & Language*, 38(5), 1238–1257.
- Gao, J. (2024). Credal sensitivism: threshold vs. credence-one. *Inquiry: An Interdisciplinary Journal of Philosophy*, 67(1), 463–484.
- Gao, J. (forthcoming). Varieties of pragmatic encroachment. In K. Sylvan, M. Steup, E. Sosa & J. Dancy (Eds.), *The Blackwell Companion to Epistemology*, Third Edition: Wiley-Blackwell
- Gao, J., Gerken, M., & Ryan, S. (2017). Does contextualism hinge on a methodological dispute? In J. Ichikawa (Ed.), *The Routledge Handbook of Epistemic Contextualism* (pp. 81–93). New York: Routledge.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178(3), 529–547.
- Gerken, M. (2013a). Epistemic focal bias. *Australasian Journal of*

- Philosophy*, 91(1), 41–61.
- Gerken, M. (2013b). *Epistemic Reasoning and the Mental*. Hampshire, New York: Palgrave Macmillan.
- Gerken, M. (2015). The role of knowledge ascriptions in epistemic assessment. *European Journal of Philosophy*, 23(1), 141–161.
- Gerken, M. (2017). *On Folk Epistemology*. Oxford: Oxford University Press.
- Gerken, M. (2018). Against knowledge first epistemology. In A. Carter, E. Gordon, and B. Jarvis (Eds.), *Knowledge First, Approaches to Epistemology and Mind* (pp. 46–71). Oxford: Oxford University Press.
- Gibbons, J. (2010). Things that make things reasonable. *Philosophy and Phenomenological Research*, 81(2), 335–361.
- Gigerenzer, G. (2008). *Rationality for Mortals: How People Cope with Uncertainty*. Oxford: Oxford University Press.
- Gigerenzer, G., and Todd, p. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple Heuristics That Make Us Smart*. New York, Oxford: Oxford University Press.
- Gilbert, D. T. (1989). Thinking lightly about others: Automatic components of the social inference process. In J. S. Uleman and J. A. Bargh (Eds.), *Unintended Thought* (pp. 189–211). New York: Guildford Press.
- Goldman, A. (1979). What is justified belief? In G. S. Pappas (Ed.), *Justification and Knowledge* (pp. 1–25). Dordrecht: Reidel.
- Graham, p. (2012). Epistemic entitlement. *Noûs*, 46(3), 449–482.
- Greco, D. (2014). Could KK be OK? *Journal of Philosophy*, 111(4), 169–197.
- Greco, D. (2015). How I learned to stop worrying and love probability one. *Philosophical Perspectives*, 29(1), 179–201.
- Grice, H. P. (1975). Logic and conversation. In Donald Davidson and Gilbert Harman (Eds.), *The logic of Grammar* (pp. 64–75). Encion, CA.: Dickenson.
- Grimm, S. (2011). On intellectualism in epistemology. *Mind*, 120(479), 705–733.
- Grimm, S. (2015). Knowledge, practical interests, and rising tides. In J. Greco and D. Henderson (Eds.), *Epistemic Evaluation: Purposeful Epistemology* (pp. 117–137). Oxford: Oxford University Press.
- Hacking, I. (1965). *Logic of Statistical Inference*. Cambridge: Cambridge

- University Press.
- Hacking, I. (1967). Possibility. *Philosophical Review*, 76(2), 143–168.
- Hannon, M. (2015). Stabilizing Knowledge. *Pacific Philosophical Quarterly*, 96(1), 116–139.
- Harker, J. (1980). A note on believing that one knows and Lehrer's proof that knowledge entails belief. *Philosophical Studies*, 37(3), 321–324.
- Harman, G. (1973). *Thought* (3rd printing, 1977.). Princeton: Princeton University Press.
- Harris, A. J. L., Corner, A., & Hahn, U. (2009). Estimating the probability of negative events. *Cognition*, 110, 51–64.
- Haselton, M., Bryant, G., Wilke, A., Frederick, D., and Galperin, A. (2009). Adaptive rationality: An evolutionary perspective on cognitive bias. *Social Cognition*, 27(5), 733–763.
- Hauska, J. (2008). Dispositions and nomal conditions. *Philosophical Studies*, 139, 219–232.
- Hawthorne, J. (2000). Implicit belief and a priori knowledge. *The Southern Journal of Philosophy*, 38(S1), 191–210.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. (2012). Knowledge and epistemic necessity. *Philosophical Studies*, 158, 493–501.
- 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*, 128(3), 591–620.
- Heaton, A., and Kruglanski, A. (1991). Person perception by introverts and extraverts under time pressure: need for closure effects. *Personality and Social Psychology Bulletin*, 17, 161–65.
- Hertwig, R., & Pedersen, A. (2016). Find foundations for bounded and adaptive rationality. *Minds & Machines*, 26, 1–8.
- Hetherington, S. C. (1999). Knowing failably. *The Journal of Philosophy*, 96(11), 565–587.
- Hill, C. S., and Schechter, J. (2007). Hawthorne's lottery puzzle and the nature of belief. *Philosophical Issues*, 17(1), 102–122.
- Holton, R. (2008). Partial belief, partial intention. *Mind*, 117(465), 27–58.
- Hookway, C. (1990). *Scepticism*. London: Routledge.

- Horty, J. (2012). *Reasons as Defaults*. Oxford: Oxford University Press.
- Hyman, J. (1999). How knowledge works. *The Philosophical Quarterly*, 49(197), 433–451.
- Iacono, L. (2008). Warranted assertability maneuvers and the rules of assertion. *Pacific Philosophical Quarterly*, 82, 460–469.
- Ichikawa, J. (2012). Knowledge norms and acting well. *Thought*, 1(1), 49–55.
- Ichikawa, J., Jarvis, B., and Rubin, K. (2012). Pragmatic encroachment and belief-desire psychology. *Analytic Philosophy*, 53(4), 327–343.
- Jackson, E. (2020). The relationship between belief and credence. *Philosophy Compass*, 15(6), e12668.
- James, W. (1956). *The Will to Believe: And Other Essays in Popular Philosophy, and Human Immortality*. New York: Dover Books.
- Jamieson, D., and Zanna, M. p. (1989). Need for structure in attitude formation and expression. In A. R. Pratkanis, S. J. Breckler, and A. G. Greenwald (Eds.) *Attitude Structure and Function* (pp. 383–406). New Jersey: Lawrence Erlbaum Associates.
- Jeffrey, R. (1970). Dracula meets Wolfman: Acceptance vs. partial belief. In *Induction, Acceptance and Rational Belief* (pp. 157–185). Dordrecht: Springer Netherlands.
- Jeffrey, R. (1983). *The Logic of Decision* (2nd ed.). Chicago, IL: University of Chicago Press.
- Jeffrey, R. (2004). *Subjective Probability-The Real Thing*. Cambridge: Cambridge University Press.
- Jeshion, R. (2000). On the obvious. *Philosophy and Phenomenological Research*, 60(2), 333–355.
- Johnston, M. (1992). How to speak of the colors. *Philosophical Studies*, 68(3), 221–263.
- Joyce, J. (2004). Williamson on evidence and knowledge. *Philosophical Books*, 45(2), 296–305.
- Joyce, J. M. (2005). How probabilities reflect evidence. *Philosophical Perspectives*, 19, 153–178.
- Kassin, S. M., Rigby, S., & Castillo, S. R. (1991). The accuracy-confidence correlation in eyewitness testimony: Limits and extensions of the retrospective self-awareness effect. *Journal of Personality and Social Psychology*, 698–707.

- Kelley, H. (1987). Attribution in social interaction. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, and B. Weiner (Eds.), *Attribution: Perceiving the Causes of Behavior* (pp. 1–26). Morristown, NJ: General Learning Press.
- Kelly, C. M. and Lindsay, D. S. (1993). Remembering mistaken for knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language*, 32(1), 1–24.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: a critique. *Philosophy and Phenomenological Research*, 66(3), 612–640.
- Kelly, T. (2014). Evidence can be permissive. In M. Steup, J. Turri, and E. Sosa (Eds), *Contemporary Debates in Epistemology* (2nd edition, pp. 298–312). Malden, MA: Blackwell.
- Kelp, C., and Simion, M. (2017). Criticism and blame in action and assertion. *Journal of Philosophy*, 114 (2), 76–93.
- Keynes, J. M. (1921). *A Treatise on Probability*. New York: Macmillan.
- Kinderman, P., Dunbar, R., and Bentall, R. 1998. Theory-of-mind deficits and causal attributions. *British Journal of Psychology*, 89 (2), 191–204.
- Knobe, J., and Schaffer, J. (2012). Contrastive knowledge surveyed. *Noûs*, 46(4), 675–708.
- Koriat, A. Ma'ayan, H. and Nussinson, R. (2006). The intricate relationships between monitoring and control in metacognition: lessons for the cause-and-effect relation between subjective experience and behavior. *Journal of Experimental Psychology: General*, 135(1), 36–69.
- Kruglanski, A. W. (2004). *The Psychology of Closed Mindedness*. New York: Psychology Press.
- Kruglanski, A. W., and Freund, T. (1983). The freezing and un-freezing of lay-inferences: Effects on impression primacy, ethnic stereotyping and numerical anchoring. *Journal of Experimental Social Psychology*, 19, 448–468.
- Kruglanski, A. W., and Webster, D. M. (1991). Group members' reactions to opinion deviates and conformists at Varying degrees of Proximity to decision deadline and of environmental noise, 61(2), 212–225.
- Kruglanski, A. W., and Webster, D. M. (1996). Motivated closing of the mind: 'Seizing' and 'freezing.' *Psychological review*, 103(2), 263.

- Kruglanski, A. W., Webster, D. M., and Klem, A. (1993). Motivated resistance and openness to persuasion in the presence or absence of prior information. *Journal of Personality and Social Psychology*, 65(5), 861–876.
- Kyburg, H. (1971). Epistemological probability. *Synthese*, 23(2), 309–26.
- Lackey, J. (2010). Acting on knowledge. *Philosophical Perspectives*, 24(1), 361–382.
- Lee, M. D., and Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the ‘take the best’ and the ‘rational’ models. *Psychonomic Bulletin & Review*, 11(2), 343–352.
- Lerner, J. S., and Tetlock, p. E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, 125(2), 255–275.
- Lehrer, K. (1983). Belief, Acceptance, and Cognition. In H. Parret (Ed.), *On Believing*. De Gruyter.
- Lehrer, K. (2000). *Theory of Knowledge*. Boulder, CO: Westview Press.
- Leng, M. (2015). Fictionalism in the philosophy of mathematics. In *Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/mathfict/>. Accessed 16 November 2015.
- Levi, I. (1980). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. Cambridge, MA: MIT Press.
- Levin, J. (2008). Assertion, practical reason, and pragmatic theories of knowledge. *Philosophy and Phenomenological Research*, 76(2), 359–384.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Littlejohn, C. (2009). Must we act only on what we know? *Journal of Philosophy*, 106(8), 463–473.
- Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.
- Littlejohn, C. (2013). The russellian retreat. *Proceedings of the Aristotelian Society*, 113(3pt3), 292–320.
- Littlejohn, C. (forthcoming). A plea for epistemic excuses. In F. Dorsch and J. Durant (Eds.), *The New Evil Demon*. Oxford: Oxford University Press.
- Locke, D. (2015). Practical certainty. *Philosophy and Phenomenological Research*, 90(1), 72–95.
- Loewenstein, G., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as

- feeling. *Psychological Bulletin*, 127(2), 267–286.
- Lutz, M. (2014). The pragmatics of pragmatic encroachment. *Synthese*, 191(8), 1717–1740.
- MacFarlane, J. (2005). The assessment sensitivity of knowledge attribution. In T. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology* (pp. 197–234). Oxford: Oxford University Press.
- MacFarlane, J. (2011). Epistemic modals are assessment-sensitive. In A. Egan and B. Weatherson (Eds.), *Epistemic Modality* (pp. 144–178). Oxford: Oxford University Press.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Maher, P. (1990). Acceptance without belief. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1990, 381–392.
- Makinson, D. C. (1965). The paradox of the preface. *Analysis*, 25(6), 205–207.
- Marcus, R. (1990). Some revisionary proposals about belief and believing. *Philosophy and Phenomenological Research*, 50, 132–153.
- Margolis, J. (1973). *Knowledge and Existence*. New York: Oxford University Press.
- Markovits, J. (2011). Internal Reasons and the Motivating Intuition. In M. Brady (Ed.), *New Waves in Metaethics* (pp. 141–165). Palgrave Macmillan.
- Martin, C. B. (1994). Dispositions and conditionals. *The Philosophical Quarterly*, 44(174), 1–8.
- May, J., Sinnott-Armstrong, W., Hull, J. G., and Zimmerman, A. (2010). Practical Interests, relevant alternatives, and knowledge attributions: an empirical study. *Review of Philosophy and Psychology*, 1(2), 265–273.
- Mayseless, O., and Kruglanski, A. W. (1987). What makes you so sure? Effects of epistemic motivations on judgmental confidence. *Organizational Behavior and Human Decision Processes*, 39(2), 162–183.
- McAllister, D., Mitchell, T., and Beach, L. (1979). The contingency model for the selection of decision strategies: An empirical test of the effects of significance, accountability, and reversibility. *Organizational Behavior and Human Performance*, 24(2), 228–244.
- McGrath, Matthew. 2017. Pragmatic encroachment: its problem are yours. In C. McHugh, J. Way, and D. Whiting (Eds.) *Normativity:*

- Epistemic and Practical* (pp. 162–178). Oxford: Oxford University Press.
- McGlynn, A. (2013). Believing things unknown. *Nóus*, 47(2), 385–407.
- McGlynn, A. (2014). *Knowledge First?* Palgrave, Macmillian.
- McGlynn, A. (2017). Mindreading knowledge. In A. Carter, E. Gordon, and B. Jarvis (Eds.), *Knowledge First: Approaches in Epistemology and Mind* (pp. 72–94). Oxford: Oxford University Press.
- McHugh, C., and Whiting, D. (2014). The normativity of belief. *Analysis*, 74(4), 698–713.
- McKenna, R. (2011). Interests contextualism. *Philosophia*, 39(4), 741–750.
- McKenna, R. (2015). Assertion, complexity, and sincerity. *Australasian Journal of Philosophy*, 93(4), 1–17.
- Mellor, D. H. (2000). The semantics and ontology of dispositions. *Mind*, 109, 757–780.
- Mellor, D. H. (2005). *Probability: A Philosophical Introduction*. London: Routledge.
- Merricks, T. (1995). Warrant entails truth. *Philosophy and Phenomenological Research*, 55(4), 841–855.
- Moser, P. K. (1988). The foundations of epistemological Probability. *Erkenntnis*, 28(2), 231–251.
- Mosterín, J. (2002). Acceptance without belief. *Manuscrito*, 25(2), 313–35.
- Mueller, Andy (2021). The knowledge norm of apt practical reasoning. *Synthese*, 199 (1-2): 5395–5414.
- Mumford, S. (1998). *Dispositions*. Oxford: Oxford University Press.
- Mumford, S. (2001). Realism and the conditional analysis of dispositions: Reply to Malekorn. *The Philosophical Quarterly*, 51, 375–378.
- Myers-Schulz, B., and Schwitzgebel, E. (2013). Knowing that *p* without believing that *p*. *Nóus*, 47(2), 371–384.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2), 279–294.
- Nagel, J. (2010a). Epistemic anxiety and adaptive invariantism. *Philosophical Perspectives*, 24(1), 407–435.
- Nagel, J. (2010b). Knowledge ascriptions and the psychological consequences of thinking about error. *Philosophical Quarterly*, 60(239), 286–306.
- Nagel, J. (2011). The psychological basis of the Harman-Vogel Paradox.

- Philosophers' Imprint*, 11(5), 1–28.
- Nagel, J. (2013). Knowledge as a mental state. In *Oxford Studies in Epistemology* (pp. 275–310). Oxford: Oxford University Press.
- Nelkin, D. K. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review*, 109(3), 373–409.
- Neta, R. (2007a). Anti-intellectualism and the knowledge-action principle. *Philosophy and Phenomenological Research*, 75(1), 180–187.
- Neta, R. (2007b). Contextualism and a puzzle about seeing. *Philosophical Studies*, 134(1), 53–63.
- Neta, R. (2009). Treating something as a reason for action. *Noûs*, 43(4), 684–699.
- Neta, R. (2012). The case against purity. *Philosophy and Phenomenological Research*, 85(2), 456–464.
- Newell, B. R. (2005). Re-visions of rationality? *Trends in Cognitive Sciences*, 9(1), 11–15.
- Newell, B. R., and Lee, M. D. (2011). The right tool for the job? Comparing an evidence accumulation and a naive strategy selection model of decision making. *Journal of Behavioral Decision Making*, 24(5), 456–481.
- Nickerson, R. S. (1999). How we know—and sometimes misjudge—what others know: Imputing one's own knowledge to others. *Psychological bulletin*, 125(6), 737.
- Otte, R. (2006). Counterfactuals and epistemic probability. *Synthese*, 152(1), 81–93.
- Owens, D. (2003). Does belief have an aim? *Philosophical Studies*, 115(3), 283–305.
- Pargetter, R. (1984). Laws and modal realism. *Philosophical Studies*, 46(3), 335–47.
- Perry, J. (1980). Belief and Acceptance. *Midwest Studies In Philosophy*, 5(1), 533–542.
- Petersen, E. N. (2014). Denying knowledge. *Canadian Journal of Philosophy*, 44(1), 36–55.
- Phelan, M. (2013). Evidence that stakes don't matter for evidence. *Philosophical Psychology*, 27(4), 1–25.
- Phillips, L. D., and Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology*, 72(3),

- 346–354.
- Pinillos, Á. (2012). Knowledge, experiments, and practical interests. In J. Brown and M. Gerken (Eds.), *Knowledge Ascription* (pp. 193–220). Oxford: Oxford University Press.
- Pinillos, Á., and Simpson, S. (2014). Experimental evidence in support of anti-intellectualism about knowledge. In J. Beebe (Ed.), *Advances in Experimental Epistemology* (pp. 9–44). London: Bloomsbury.
- Plantinga, A. (1982). How to be an anti-realist. *Proceedings and Addresses of the American Philosophical Association*, 56(1), 47–70.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: Oxford University Press.
- Price, H. H. (1969). *Belief*. London: Allen & Unwin.
- Pritchard, D. (2000). Doubt undogmatized: Pyrrhonian scepticism, epistemological externalism, and the ‘metaepistemological’ challenge. *Principia—Revista Internacional de Epistemología*, 4, 187–214.
- Pritchard, D. (2001). Contextualism, scepticism, and the problem of epistemic descent. *Dialectica*, 55(4), 327–349.
- Pritchard, D. (2010). Contextualism, skepticism and warranted assertibility manoeuvres. In M. O’Rourke and H. Silverstein (Eds.), *Knowledge and Skepticism* (pp. 85–103). Cambridge, MA: MIT Press.
- Proust, J. (2012). The norms of acceptance. *Philosophical Issues*, 22(1), 316–333.
- Pryor, J. (2004). What’s wrong with Moore’s argument? *Philosophical Issues* 14 (1): 349.
- Pryor, J. (2005). There is immediate justification. In M. Steup and E. Sosa (Eds.), *Contemporary Debates in Epistemology*, (pp. 181. mp). Oxford: Blackwell.
- Radford, C. (1990). Belief, acceptance, and knowledge. *Mind*, 99(396), 609–617.
- Ramsey, F. (1926). *Truth and probability*. (Reprinted from *The Foundations of Mathematics and Other Logical Essays*, by R. B. Braithwaite, Ed., 1950, New York: The Humanities Press.)
- Reed, B. (2002). How to think about fallibilism. *Philosophical Studies*, 107(2), 143–157.
- Reed, B. (2010). A defense of stable invariantism. *Noûs*, 44(2), 224–244.
- Reed, B. (2013). Fallibilism, epistemic possibility, and epistemic agency. *Philosophical Issues*, 23(1), 40–69.

- Reichenbach, H. (1949). *The Theory of Probability*. Berkeley: University of California Press.
- Rey, G. (1988). Toward a computational account of akrasia and self-deception. In B. McLaughlin and O. Rorty (Eds.), *Perspectives on self-deception* (pp. 264–296). London: University of California Press.
- Rey, G. (2007). Meta-atheism: religious avowal as self-deception. In Antony (Ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life* (pp. 243–265). Oxford: Oxford University Press.
- Reynolds, S. (2013). Justification as the appearance of knowledge. *Philosophical Studies*, 163(2), 367–383.
- Rieskamp, J., and Otto, p. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135(2), 207–236.
- Ring, M. (1977). Knowledge: the cessation of belief. *American Philosophical Quarterly*, 14(1), 51–59.
- Rizzieri, A. (2011). Evidence does not equal knowledge. *Philosophical Studies*, 153(2), 235–242.
- Robinson, L. B., & Hastie, R. (1985). Revision of beliefs when a hypothesis is eliminated from consideration. *Journal of Experimental Psychology: Human Perception and Performance*, 11(4), 443–456.
- Roeber, B. (2013). *Believing, Knowing, Acting*. Rutgers, The State University of New Jersey, New Brunswick.
- Roeber, B. (2014). Minimalism and the limits of warranted assertability maneuvers. *Epistems*, 11(3), 245–260.
- Roeber, B. (2018). The pragmatic encroachment debate. *Noûs*, 52(1), 171–195.
- Roet, A., Kruglanski, A. W., Kossowska, M., Pierro, A., & Hong, Y. (2015). The motivated gatekeeper of our minds: New directions in need for closure theory and research. *Advances in Experimental Psychology*, 52, 221–283.
- Rose, D. Machery, E. Stich, S. et al. (2019). Nothing at stake in knowledge. *Noûs*, 53 (1), 224–247.
- Rose, D., and Schaffer, J. (2013). Knowledge entails dispositional belief. *Philosophical Studies*, 166(1), 19–50.
- Ross, J., and Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*, 88(2), 259–288.

- Rottenstreich, Y., & Hsee, C. K. (2001). Money, kisses, and electric shocks: On the affective psychology of risk. *Psychological Science*, 12(3), 185–190.
- Royzman, E. B., Cassidy, K. W., and Baron, J. (2003). 'I know, you know': Epistemic egocentrism in children and adults. *Review of General Psychology*, 7(1), 38–65.
- Rubin, K. (2015). Total pragmatic encroachment and epistemic permissiveness. *Pacific Philosophical Quarterly*, 96, 12–38.
- Russell, B. (1948). *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.
- Russell, G. & Doris, J. (2008). Knowledge by indifference. *Australasian Journal of Philosophy*, 86 (3):429–437.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35(4), 477–514.
- Rysiew, P. (2007). Speaking of knowing. *Noûs*, 41(4), 627–662.
- Rysiew, P. (2008). Rationality disputes-psychology and epistemology. *Philosophy Compass*, 3(6), 1–24.
- Salas, J. G. de p. (2019). Extreme Betting. *Ratio*, 32(1), 32–41.
- Samuels, R., Stich, S., and Faucher, L. (2004). Reason and rationality. In I. Niiniluoto, M. Sintonen, & J. Wolenski (Eds.), *Handbook of Epistemology* (pp. 131–179). Dordrecht: Kluwer Academic Publishers.
- Savage, L. (1954). *The Foundations of Statistics*. Hoboken, NJ: Wiley.
- Schroeder, M. (2012). Stakes, withholding, and pragmatic encroachment on knowledge. *Philosophical Studies*, 160(2), 265–285.
- Schulz, M. (2017). Decisions and higher-order knowledge. *Noûs*, 51(3), 463–483.
- Schwitzgebel, E. (2010). Acting contrary to our professed beliefs or the gulf between occurrent judgment and dispositional belief. *Pacific Philosophical Quarterly*, 91(4), 531–553.
- Schwitzgebel, E. (2021). Belief. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition). <http://plato.stanford.edu/archives/sum2015/entries/belief/>. Accessed 17 June 2023.
- Sellars, W. (1964). Induction as vindication. *Philosophy of Science*, 31(3), 197–231.
- Sextus Empiricus. (1994). *Outlines of Pyrrhonism*. (Annas and Barnes,

- Trans.).
- Shin, J. (2014). Time constraints and pragmatic encroachment on knowledge. *Episteme*, 11(2), 157–180.
- Shope, R. (2002). Conditions and analyses of knowing. In p. Moser (Ed.), *The Oxford Handbook of Epistemology* (pp. 25–70). Oxford: Oxford University Press.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63 (2), 129–139.
- Slooman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological bulletin*, 119(1), 3–22.
- Smith, M. (2016). *Between Probability and Certainty*. Oxford: Oxford University Press.
- Sosa, E. (2011). *Knowing Full Well*. Princeton: Princeton University Press.
- Sperber, D. (1996). *Explaining Culture: A Naturalistic Approach*. Oxford: Basil Blackwell.
- Sperber, D. (1997). Intuitive and reflective beliefs. *Mind and Language*, 12(1), 67–83.
- 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.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Stanley, J. (2005a). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stanley, J. (2005b). Fallibilism and concessive knowledge attributions. *Analysis*, (65), 126–31.
- Stanovich, K. E. (2011). *Rationality and the Reflective Mind*. Oxford: Oxford University Press.
- Stanovich, K. E., and West, R. F. (2000). Individual differences in reasoning: implications for the rationality debate? *The Behavioral and Brain Sciences*, 23(5), 645–665–726.
- Steglich-Petersen, A. (2006). No norm needed: on the aim of belief. *The Philosophical Quarterly*, 56(225), 499–516.
- Stein, E. (1996). *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*. New York: Oxford University Press.
- Steinberg, J. (2010). Dispositions and subjunctives. *Philosophical Studies*,

- 148, 323–341
- Stiller, J., and Dunbar, R. I. M. (2007). Perspective-taking and memory capacity predict social network size. *Social Networks*, 29(1), 93–104.
- Sunstein, C. R. (2002). Probability neglect: Emotions, worst cases, and law. *The Yale Law Journal*, 112, 61–107.
- Swinburne, R. (2001). *Epistemic Justification*. Oxford: Oxford University Press.
- Tang, W. H. (2015). Belief and cognitive limitations. *Philosophical Studies*, 172(1), 249–260.
- Tetlock, P. E., & Kim, J. I. (1987). Accountability and judgment processes in a personality prediction task. *Journal of Personality and Social Psychology*, 52(4), 700–709.
- Todd, P., and Brighton, H. (2016). Building the theory of ecological rationality. *Minds & Machines*, 26, 9–30.
- Tuomela, R. (2000). Belief versus acceptance. *Philosophical Explorations*, 3(2), 122–137.
- Turri, J. (2010). Epistemic invariantism and speech act contextualism. *Philosophical Review*, 119(1), 77–95.
- Turri, J. (2015). Skeptical appeal: the source-content bias. *Cognitive Science*, 39(2), 307–324.
- Turri, J., and Friedman, O. (2014). Winners and losers in the folk epistemology of lotteries. In J. Beebe (Ed.), *Advances in Experimental Epistemology* (pp. 45–69). London: Bloomsbury.
- Turri, J. (2017). Epistemic contextualism: An idle hypothesis. *Australasian Journal of Philosophy*, 95(1), 141–156.
- Turri, J., Buckwalter, W. & Rose, D. (2016). Actionability judgments cause knowledge judgments. *Thought: A Journal of Philosophy*, 5(3), 212–222.
- Vahid, H. (2006). Aiming at truth: doxastic vs. epistemic goals. *Philosophical Studies*, 131(2), 303–335.
- van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Oxford University Press.
- van Inwagen, P. (1996). Reflections on the chapters by Draper, Russell, and Gale. In D. Howard-Snyder (Ed.), *The Evidential Argument From Evil* (pp. 219–243). Bloomington, Indiana: Indiana University Press.
- van Leeuwen, N. (2014). Religious credence is not factual belief.

- Cognition*, 133(3), 698–715.
- Velleman, D. (2000). On the aim of belief. In *The Possibility of Practical Reason* (pp. 244–81). Oxford: Oxford University Press.
- Vogel, J. (1999). The new relevant alternatives theory. *Philosophical Perspectives*, 13(s13), 155–180.
- Vogt, K. (2010). Ancient skepticism. In (E. N. Zalta, Ed.) *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/win2011/entries/skepticism-ancient/>.
- von Mises, R. (1957). *Probability, Statistic and Truth* (revised English edition.). New York: Macmillan.
- Weatherson, B. (1999). Begging the question and Bayesianism. *Studies in History and Philosophy of Science*, 30, 687–697.
- Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19(1), 417–443.
- Weatherson, B. (2011). Defending interest-relative invariantism. *Logos & Epistemic*, 2(4), 591–609.
- Weatherson, B. (2012). Knowledge, bets and interests. In J. Brown and M. Gerken (Eds.), *Knowledge Ascriptions* (pp. 75–103). Oxford: Oxford University Press.
- Weatherson, B. (2016). Games, beliefs and credences. *Philosophy and Phenomenological Research*, 92(2), 209–236.
- Webster, D. M., Richter, L., and Kruglanski, A. W. (1996). On leaping to conclusions when feeling tired: Mental fatigue effects on impressional primacy. *Journal of Ethics and Social Philosophy*, 32(2), 181–195.
- Webster, Donna M. (1993). Motivated augmentation and reduction of the overattribution bias. *Journal of Personality and Social Psychology*, 65(2), 261–271.
- Webster, Donna M, and Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology*, 67(6), 1049–1062.
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, 16(s16), 267–297.
- Wedgwood, R. (2012). Outright Belief. *dialectica*, 66(3), 309–329.
- Wedgwood, R. (2013). Doxastic correctness. *Aristotelian Society Supplementary Volume*, 87(1), 217–234.
- Weisberg, J. (2020). Belief in psyontology. *Philosophers' Imprint*, 20(11),

- 1–27.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72(3), 655–684.
- Williams, B. (1973). Deciding to believe. In *Problems of the Self* (pp. 136–151). Cambridge: Cambridge University Press.
- Williams, M. J. (2000). Is contextualism statable? *Philosophical Issues*, 10(1), 80–85.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *Philosophical Quarterly*, 55(219), 213–235.
- Williamson, T. (2011). Knowledge first epistemology. In S. Bernecker and D. Pritchard (Eds.), *Routledge Companion to Epistemology* (pp. 208–218). New York: Routledge.
- Williamson, T. (2013). Knowledge first. In M. S. J. Turri (Ed.), *Contemporary Debates in Epistemology* (pp. 1–10). Oxford: Blackwell.
- Williamson, T. (forthcoming). Justifications, excuses, and sceptical scenarios. In F. Dorsch and J. Dutant (Eds.), *The New Evil Demon*. Oxford: Oxford University Press.
- Woozley, A. (1952). Knowing and not knowing. *Proceedings of the Aristotelian Society*, 53, 151–172.
- Worsnip, A. (2015). Possibly false knowledge. *Journal of Philosophy*, 112 (5), 225–246.
- Wright, C. (2004). Warrant for nothing (and foundations for free)? *Proceedings of the Aristotelian Society, Supplementary Volumes*, 78, 167–212.
- Yablo, S. (2006). Non-catastrophic presupposition failure. In J. J. Thomson and A. Byrne (Eds.), *Content and Modality: Themes From the Philosophy of Robert Stalnaker* (pp. 164–190). Oxford: Oxford University Press.
- Zhao, J., Crupi, V., Tentori, K., Fitelson, B., & Osherson, D. (2012). Updating: Learning versus supposing. *Cognition*, 124(3), 373–378.

Index

- acceptance 42, 44, 76–85, 87–89, 91, 93, 158–159, 161, 167
Alston, W 79, 184–185
Anderson, C. 1, 39, 202
Audi, R. 79, 184,
Bach, K. 31, 95, 107–109, 111, 115, 117, 128, 152, 207
BANK Case 10–11, 29, 33, 35–36, 45, 47, 63, 102, 107
Baumann, P. 34, 39
Bayesianism 61, 85–86, 123, 145–146, 152–153
belief
 dispositionalism 101, 159
 functionalism 95–96, 102
 see also dispositional belief, doxastic pragmatism, occurrent
 belief, outright belief, threshold view
Benton, M. 77
Bird, Alexandar 104, 107
Black, C. 115
Black, T. 32
Blome-Tillmann, M. 29–30, 32, 34, 36
BonJour, L. 184–185
bounded rationality 123, 144, 146–150, 153–154, 156, 171–173, 175, 177, 180,
 189, 217
Braithwaite, R. B. 158
Bratman, M. 79–81,
Brown, J. 13, 21, 25, 28, 32–33, 35–37, 39–40, 48, 59, 74–77, 197, 201, 204,
 207, 213
Buchak, L. 151
Buckwalter, W. 14–15, 45
Burge, T. 154
caution-oriented factors 124, 126, 133, 135, 138, 143, 166, 168
Clarke, R. 39, 153

- credence 2, 5, 31, 82, 96, 99–100, 110, 118, 121–122, 125–127, 130–131, 134–139, 141–146,
148–149, 152–157, 164, 168–175, 177, 180, 188–189, 217
see also partial belief, degree of belief
- Cohen, J. 79–80, 83
- Cohen, S. 13, 35, 47, 54, 184
- communicative heuristic 195, 197–198
- concessive knowledge attributions 156, 182, 190, 199–202, 204–206,
208–211, 215
- conversational implicature 32–38, 85, 196–197, 202, 208
- contextualism 23, 47–48, 170
- credal pragmatism 2–5, 121–122, 125–127, 130–140, 143–145, 148, 152–153,
155–157, 172–173, 175–176, 180–182, 199, 215–218
- degree of belief 101, 128, 188
- DeRose, K. 11, 13–14, 34–37, 45, 47, 129
- Dinges, A. 15, 45, 178
- Dimmock, P. 34–36, 138, 178
- dispositional belief 2, 5, 107, 115–117, 144, 149, 157–165, 168–180, 209–210,
217–218
- Dodd, D. 139, 186, 188, 202, 204–205, 207
- Dougherty, T. 186, 202–209
- doxastic pragmatism 31, 94–96, 99, 101, 105, 107, 114, 118, 121–127, 131, 133, 140, 142,
144, 216–217
see also credal pragmatism; threshold pragmatism
- Dub, R. 79–80
- Dutch Book 145
- Egan, A. 90
- Engel, P. 79–80
- epistemic rationality 144, 149–156, 175, 180
see also bounded rationality; ideal rationality; purist rationality
- excuse manoeuvre 39, 92–93
- fallibilism 181–183
logical conception of 183–185, 192–193
probability conception of 185–186, 189–190, 192–194
epistemic modal 186, 194, 205
fail-to-be-knowledge conception of 190, 193
non-maximal conception of 192–194

- Fantl, J. and McGrath, M. 1–3, 13–14, 16–19, 21–22, 25–27, 66, 71, 77, 100, 174, 176, 184, 186–188, 192–194, 202–204
- Fassio, D. 27, 80, 125, 188
- Frankish, K. 79, 158–161
- Ganson, D. 31, 95–97, 99–103, 105–106, 109, 124, 128
- Gao, J. 15, 21, 23, 27, 72, 93, 118, 139, 149, 156
- Gerken, M. 1, 3, 10, 13–14, 16, 28, 31–32, 39–42, 48, 74–77, 86, 92, 142, 154, 178, 182, 195–198, 201, 204–205, 207, 213
- Gettier 39, 74, 190, 192–193
- Gigerenzer, G. 63, 123, 137, 155, 167
- Goldman, A. 154, 185
- Greco, D. 65, 139
- Grice, H. P. 35, 196, 202
- Grimm, S. 22, 152
- Hannon, M. 22
- Harman, G. 184
- Hawthorne, J. 1, 16, 19, 21, 22, 38, 66, 78, 186–7, 202
- Hawthorne, J. and Stanley, J. 2, 16–18, 39, 54, 57, 66, 74, 78, 86, 92–3
- Hazlett, A. 32
- Hume, D. 158, 189
- Huvenes 34–36, 38
- high-stakes 10–16, 19, 22–24, 26–27, 29, 31–33, 35–37, 44–46, 49–51, 52, 62, 64–66, 69, 71, 95, 108, 110, 112–113, 126–127, 131, 133, 148, 169–170, 174–180, 188
- higher-order belief 66–71
- higher-order knowledge 44, 46, 49–53, 58–63, 65–69, 71
- higher-order reasoning 52, 58–59, 62–66
- Iacono, L. 37
- Ichikawa, J. 28, 41
- ideal rationality 123, 144, 146, 148, 150, 154, 156, 175
- infallibilist intuition 5, 182, 195, 199–202, 206, 211, 213–215, 218
- interest relative Invariantism 21
- invariantism, 3
- see subject sensitive Invariantism, moderate Invariantism, interest relative Invariantism
- justification 1, 11, 21, 25–26, 41–42, 44, 52, 94, 112, 127–128, 149, 153–154,

- 174, 184–185, 190–193, 197
- knowledge ascriptions 3–4, 9–11, 13–15, 22, 30, 33–37, 42, 48, 71, 94, 99, 112, 114, 121, 197, 200–201, 209, 212–213, 215–216, 218
- knowledge-first epistemology 16, 44
- knowledge norm of practical reasoning 3–4, 16–20, 25–26, 28, 38–39, 41–42, 46–49, 51–54, 56, 68, 71–72, 74–78, 82, 86–93, 156, 171, 180–182, 194–196, 198–201, 213–215, 217
- Kruglanski, A. 110, 129, 133–141, 147, 160, 165–169
- Lackey, J. 17, 39–41
- Lewis, D. 176, 185, 189, 195, 198, 200,
- Littlejohn, C. 38–39, 41–42, 74–75
- low-stakes 10–12, 14–15, 19–21, 23–24, 26–27, 31–33, 37–38, 42, 46–47, 50, 52, 57, 99, 102, 111, 113–114, 121, 124, 126–127, 129, 133, 140, 145, 154, 177, 179, 197
- Lutz, M. 32–33, 35–36
- MacFarlane, J. 3, 207
- May, J., Sinnott-Armstrong, W., Hull, J. G., and Zimmerman, A. 15, 45
- Mayseless, O. 129, 134–136, 138–141, 166–168
- McGlynn, A. 16, 40, 62, 195, 197
- McKenna, R. 55, 170
- moderate Invariantism 3–4, 9, 15, 23, 28, 31, 35, 37–38, 42–44, 45–46, 48–49, 53, 55, 71–72, 93, 107, 173, 181, 194, 215–218
- Mosterin, J. 79–81
- Mueller, A. 2, 20
- Myers-Schulz, B., and Schwitzgebel, E. 116, 161–162
- Nagel, J. 4, 13, 31, 45, 64, 95, 109–115, 117, 122, 128–130, 133, 141, 155, 160, 178–179, 197
- need-for-closure 111, 133–138, 140–142, 148–149, 151, 154–155, 163, 165–171, 175, 209
- need-to-avoid-closure 133–136, 138–141, 148–149, 151, 165–170, 174, 209
- normal circumstances 2, 28, 81, 100, 135, 149–152, 156, 163, 168, 171–172, 176, 188–189, 212, 217
- occurrent belief 2, 116–117, 157–166, 168, 170–180, 182, 199, 208–214, 217–218
- outright belief 31, 34, 94, 100–102, 107, 109, 112, 114–115, 121, 123, 125, 127–130, 144–145, 155, 157, 172, 179

- partial belief 92
- Pinillos, Á. 14–15
- Plantinga, A. 185
- practical factor effects 4–5, 9–10, 22–23, 31, 38, 42, 114, 117–118, 121, 157, 173, 215, 217–218
- practical factors, 1–3, 5, 10–11, 15, 20–22, 24–26, 31, 41, 94–95, 97, 118, 121–126, 130–131, 137, 141–145, 147–153, 155, 157, 164–177, 180, 217–218
see also caution-oriented factors; rush-oriented factors
- practical matters *see* practical factors, practical rationality
- practical rationality 1–2, 58, 71, 92
- pragmatic encroachment 21–23, 25, 27–31, 42, 44–45, 72, 93, 96, 128, 145, 149, 204, 217–218
- Pritchard, D. 32, 90, 170, 191, 207
- purism 1, 3, 21–22, 26, 94, 176,
- purist rationality 144, 149–150, 154, 156, 171, 173, 175, 178, 180, 218
- Radford, C. 115
- rational belief 40, 125
- rational credence 122, 125, 130, 144–145, 152, 188–189
- Reed, B. 11, 39, 48, 66, 107, 162, 174, 182–187, 190, 194, 196–198
- Rieskamp, J. and Otto, P. 63, 167
- Roeber, D. 34–36, 39, 48
- Rose, D. and Schaffer, J. 115–117, 158–162
- Ross, J., and Schroeder, M. 1–2, 14, 18, 21, 100
- rush-oriented factors 124, 126, 130, 135, 137–138, 143, 165–166, 168
- Rubin, K. 122, 145
- Russell, G. and Doris, J. 29
- Rysiew, P. 32–36, 186, 190, 202–209
- threshold pragmatism 5, 121–122, 125–128, 130–134, 137–139, 143, 155
- truth-relevant factors 1, 3, 21, 47, 146, 149, 155–156, 171
- Schroeder, M. 21
- Schulz, M. 17
- Schwitzgebel, E. 116, 158–159
- Shin, J. 15
- Smith, M. 191
- Sripada C.S. and Stanley, J. 14–15, 113, 178–179
- Stanley, J. 1–2, 11–14, 16, 18–19, 21–22, 25, 71, 111, 122, 184, 199, 202

subject sensitive Invariantism 21–22

SURGEON 39–40, 197, 201, 204, 207, 210, 212

threshold view 99, 100, 114, 125, 127, 129, 139, 164, 168, 170–172, 183

Trustful Airport 54, 56

Turri, J. 15, 31, 44

Unconfident Examinee 115–117, 161–162, 164, 169

Weatherson, B. 1, 2, 14, 17, 21–22, 24, 27–28, 30–31, 40, 71, 95–99, 101–103,
106, 109, 112, 127–128, 188

Webster, D. 133, 137–139, 141, 165–167, 169

Wedgwood, R. 80, 102–103, 124, 152

Williamson, T. 2, 4, 16, 19–20, 39, 43–64, 66–72, 90, 92, 117–118, 134, 191,
193, 217

Wright, C. 78–79

Ye, R. 27