


KNOWLEDGE, CHANCE, AND CONTRAST

Paul Dimmock

**A Thesis Submitted for the Degree of PhD
at the
University of St. Andrews**



2012

**Full metadata for this item is available in
Research@StAndrews:FullText
at:**

<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:

<http://hdl.handle.net/10023/3202>

This item is protected by original copyright

Knowledge, Chance, and Contrast

Paul Dimmock

Ph.D. Thesis
University of St Andrews

Submitted: 9th February 2012

1. Candidate's declarations:

I, Paul Dimmock, hereby certify that this thesis, which is approximately 78,000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September, 2007 and as a candidate for the degree of PhD in Philosophy in September, 2007; the higher study for which this is a record was carried out in the University of St Andrews between 2007 and 2012.

Date 06/02/2012 signature of candidate

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date 07/02/2012 signature of supervisor

3. Permission for electronic publication:

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the electronic publication of this thesis:

(i) Access to printed copy and electronic publication of thesis through the University of St Andrews.

Date 07/02/2012 signature of candidate

signature of supervisor

Acknowledgements

This thesis could not have been produced without the support of family, friends, and colleagues. I would like to especially thank Jessica Brown and Patrick Greenough for extensive advice, comments, and other guidance throughout my PhD studies. Thanks also to colleagues and friends at the University of St Andrews, and in the School of Philosophy at the Australian National University. Individual thanks to Ralf Bader, Björn Brodowski, Berit Brogaard, Yuri Cath, Stewart Cohen, Stuart Crutchfield, Derek Ball, Colin Caret, David Chalmers, Michael De, Dylan Dodd, Douglas Edwards, Ole Hjortland, Torfinn Huvenes, Jonathan Ichikawa, Dirk Kindermann, Julia Langkau, Federico Luzzi, Ezra Macdonald, Fiona Macpherson, Dilip Ninan, Jonathan Schaffer, Daniele Sgaravatti, Declan Smithies, Brian Weatherson, Timothy Williamson, and Crispin Wright.

I would also like to acknowledge the generous support of the Arts and Humanities Research Council (UK).

CONTENTS

<i>Introduction</i>	p. 1
1. Traditional Invariantism and Experimental Psychology	p. 9
2. Traditional Invariantism and Pragmatics	p. 58
3. Chance and Possibility	p. 92
4. Interest and Salience Relativity	p. 134
5. Scepticism and Contextualism	p. 170
<i>References</i>	p. 216

Introduction

The late 1980s and early 1990s saw the rise of *contextualist* theories of knowledge ascriptions (and denials). Contextualists about ‘knows’ maintain that utterances of the form ‘S knows p’ and ‘S doesn’t know p’ resemble utterances such as ‘Peter is here’ and ‘Peter is not here’, in the sense that their truth-conditions vary depending upon features of the context in which they are uttered. In recent years, contextualism about ‘knows’ has come under heavy attack. This has been associated with a proliferation of defences of so-called *invariantist* accounts of knowledge ascriptions, which stand united in their rejection of contextualism.¹

The central goal of the present work is two-fold. In the first instance, it is to bring out the serious pitfalls in many of those recent defences of invariantism. In the second instance, it is to establish that the most plausible form of invariantism is one that is *sceptical* in character. Of course, the prevailing preference in epistemology is for *non-sceptical* accounts. The central conclusions of the thesis might therefore be taken to show that – despite recent attacks on its plausibility – some form of contextualism about ‘knows’ must be correct. However, this project is not undertaken without at least the suspicion that embracing (a particular form of) sceptical invariantism is to be preferred to embracing contextualism. In the course of the discussion, I therefore not only attempt to rebut some standard objections to sceptical invariantism, but also to reveal – in at least a preliminary way – how the sceptical invariantist might best argue for the superiority of her account to that of the contextualist.

There are a number of judgments that have proven central in the debates between contextualists and invariantists. One contrast in our judgments concerning ‘knows’ – henceforth the ‘HIGH-LOW’ contrast – that has gained much attention is manifested in examples like Keith DeRose’s Bank Case.

¹ For prominent defences of contextualism, see Cohen (1988, 1999), DeRose (1992, 1995, 2009), and Lewis (1996). For criticism of contextualism, see *inter alia* Feldman (2001), Stanley (2004, 2005a), Hawthorne (2004a), MacFarlane (2005, *f/c*), and Williamson (2005a, 2005b). For recent defences of invariantism, see *inter alia* Hawthorne (2004a), Stanley (2005a), Fantl & McGrath (2007), Brown (2005, 2006), Rysiew (2001, 2007), Hazlett (2009), Schaffer (2004), Davis (2007), Nagel (2008, 2010a, 2010b), and Williamson (2005a, 2005b).

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

Bank Case B. My wife and I drive past the bank on a Friday afternoon, as in case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, ‘Banks do change their hours. Do you know that the bank will be open tomorrow?’ Remaining as confident as I was before that the bank will be open then, still, I reply, ‘Well, no, I don’t know. I’d better go in and make sure.’ (DeRose 2009, p. 1-2; see also DeRose 1992, p. 913)

It is commonly reported that at least when the cases are considered in isolation from one another, DeRose’s knowledge *ascription* seems true and appropriate in Bank Case A – the ‘LOW’ context – but that his knowledge *denial* seems true and appropriate in Bank Case B – the ‘HIGH’ context.

DeRose appears to occupy the same epistemic position in HIGH that he does in LOW. Indeed, it can seem that the only relevant differences between the HIGH context and the LOW context are that the conversational participants in HIGH have a lot at stake, and are considering the possibility that the bank has recently changed its hours. This fact has been exploited to suggest that our judgments concerning HIGH and LOW represent a strong argument in favour of contextualism about ‘knows’, and against competing invariantist accounts.

The contextualist assimilates the HIGH-LOW contrast to the contrast found in examples like the following, featuring ‘here’. Suppose that John is speaking on the telephone with a friend. John is having a party at his house, and his friend asks him ‘Is Peter there?’. John replies ‘Peter is here’. Suppose that Mary is also speaking on the telephone with a friend. Mary is also having a party at her house, and her friend asks her, ‘Is Peter there?’. She replies ‘Peter is not here’. Suppose that Peter is at John’s house, but not Mary’s. In that case – and just as in the Bank Case – Mary has uttered a sentence that is the negative form of the sentence uttered by John, but both utterances seem true and appropriate.

The natural treatment of our judgments in the ‘here’ case is a contextualist one. The contextualist about ‘here’ maintains that utterances of ‘Peter is here’ and ‘Peter is not here’ are context sensitive in such a way that John’s utterance expresses something like *Peter is at John’s house* whilst Mary’s utterance expresses something like *Peter is not at Mary’s house*. Both these propositions are true, and this – it is proposed – accounts for our judgments to the effect that the relevant utterances are true and appropriate. The contextualist about ‘knows’ takes the HIGH-LOW contrast to support the claim that something similar is true in the case of knowledge ascriptions (and denials). She suggests that utterances of the form ‘S knows p’ and ‘S doesn’t know p’ are context sensitive in such a way that DeRose’s knowledge ascription in LOW expresses a proposition like *DeRose knows that the bank will be open by low standards*, whilst his knowledge denial in HIGH expresses a proposition like *DeRose does not know that the bank will be open by high standards*. Both these propositions are – it is alleged – true, and this accounts for our judgments to the effect that the relevant utterances are true and appropriate.

The contextualist appears able to neatly explain the HIGH-LOW contrast, and to do so in a way that avoids the attribution of error to speakers. But the invariantist might not appear to fare so well. The invariantist maintains that ‘knows’ is not context sensitive, and therefore claims that DeRose’s knowledge ascription in LOW and his knowledge denial in HIGH express (respectively) the propositions that *DeRose knows that the bank will be open* and *DeRose does not know that the bank will be open*. The invariantist has also traditionally clung to the *intellectualist* thesis that the only difference between true belief and knowledge lies in such ‘truth-relevant’ factors as

the reliability of the subject's belief forming mechanisms, and the character of the evidence she possesses. However, as noted above, the relevant differences between DeRose's true belief in the HIGH context and his true belief in the LOW context appear to concern only *non*-'truth-relevant' factors – viz. how much is at stake, and the kind of error possibilities being considered. It follows that if the proposition expressed by DeRose's knowledge ascription in LOW is true, then the proposition expressed by his knowledge denial in HIGH is false, and *vice versa*. In consequence, it can seem that the invariantist (or at least, the invariantist who accepts intellectualism) must rule that *either* the judgment that DeRose's knowledge ascription in LOW is true and appropriate *or* the judgment that DeRose's knowledge denial in HIGH is true and appropriate is mistaken.

The HIGH-LOW contrast therefore represents a *prima facie* argument in favour of contextualism and against invariantism (or at least, against those forms of invariantism that include a commitment to intellectualism). The extent to which the HIGH-LOW contrast really does represent a strong argument in favour of contextualism and against invariantism presumably depends on the precise extent and nature of the error (if any) that the invariantist has to attribute to speakers in order to account for our judgments.

Outline of the Thesis

In the first three chapters of the thesis, I argue against *traditional invariantism*. In addition to intellectualism, the traditional invariantist maintains that our ordinary knowledge ascriptions – such as DeRose's knowledge ascription in LOW – are true. In regard to the HIGH-LOW contrast, the task for the traditional invariantist is therefore to explain the apparent truth and appropriateness of DeRose's (false) knowledge denial in HIGH, and to do so in a manner that does not involve the attribution of problematic error to speakers.

In Chapter 1, I explore the suggestion that such an explanatory undertaking is best carried out by appeal to various general biases in our *psychological processing*. In Chapter 2, I explore the prominent alternative to such a psychological account – viz. an account that locates the apparent truth and appropriateness of DeRose's (false)

knowledge denial in HIGH in some *pragmatic* feature(s) of language. Neither species of explanation is found to be plausible. It is difficult to see how *else* the traditional invariantist can explain our judgment to the effect that DeRose's knowledge denial in HIGH seems true and appropriate, other than to (rather lamely) claim that speakers exhibit a faulty grasp of the meaning of 'knows'. It is presumably much more problematic for a theory to have to attribute such an error to speakers, so let's describe an attribution of that kind of error as an attribution of *significant error*. Taken in conjunction, the arguments of Chapters 1 and 2 suggest that the traditional invariantist is unable to account for the HIGH-LOW contrast without attributing significant error to speakers.

In the course of the first two chapters, I shall also recount a familiar objection to the effect that traditional invariantism encounters problems with our judgments relating knowledge and further inquiry, such as the judgment that utterances of the form 'I know that p, but I must inquire further' seem infelicitous. (For example, in the HIGH context of the Bank Case, it seems infelicitous for DeRose to utter 'I know that the bank will be open, but we need to inquire further'.) It is argued that various recent attempts to account for the infelicitousness of such utterances – without attributing significant error – are unsuccessful.

In Chapter 3, I leave the HIGH-LOW contrast aside, and focus on some judgments that have gained rather less attention in the literature: those concerning the relationship between knowledge and epistemic chance and possibility. It has been remarked that utterances of the following form (typically) seem infelicitous: 'I know that p, but it might be that not-p' and 'I know that p, but there's a chance that not-p'. For example, it seems infelicitous for DeRose (in either LOW or HIGH) to utter 'I know that the bank will be open, but it might not be' or 'I know that the bank will be open, but there's a chance it won't be'. Our judgments concerning the (in)felicitousness of such utterances suggest that principles along the following lines are true:

(CK) If S knows that p, then there's no epistemic chance, for S, that not-p.

(PK) If S knows that p, then it's epistemically impossible, for S, that not-p.

I present a number of arguments that suggest that the traditional invariantist can accept these principles only at significant cost – a cost, once again, in the form of an attribution of significant speaker error. I also consider and reject recent proposals – including those that involve appeal to pragmatic features of language, and those that locate a kind of context sensitivity in our talk of epistemic chance and possibility – for explaining the infelicity of utterances such as ‘I know that p, but it might be that not-p’ *without* (at least straightforwardly) conceding the truth of principles like (CK) or (PK). The upshot is that the traditional invariantist is left with a dilemma – acceptance *or* rejection of (CK) and (PK) leads to the attribution of (significant) speaker error.

In the final sections of Chapter 3, I consider the principle of Multi-Premise Closure (MPC).

(MPC) If a subject S knows propositions p_1 to p_n and competently deduces q from p_1 to p_n – retaining throughout her knowledge of p_1 to p_n – then S knows q.

It is argued that considerations pertaining to (MPC) bring out a further dilemma for the traditional invariantist – both acceptance and rejection of (MPC) bring about additional problems for the traditional invariantist, problems that once again seem to lead the traditional invariantist into positing significant speaker error. (I also explore the relationships between traditional invariantist acceptance of (CK), (PK), and (MPC) and argue that endorsement of (CK) and (PK) entails endorsement of (MPC), and *vice versa*.)

The first three chapters thus constitute the backbone of the case against traditional invariantism, culminating in the conclusion that in respect to (i) the HIGH-LOW contrast, (ii) our talk concerning the relationship between knowledge and further inquiry, (iii) our talk concerning the relationship between knowledge and epistemic chance and possibility, and (iv) our judgments pertaining to the closure of knowledge under (multi-premise) deductive inference, the traditional invariantist has to posit significant speaker error.

In Chapter 4, I consider the possibility that the invariantist can assuage the difficulties encountered by the traditional invariantist – and retain a commitment to the truth of

such ordinary knowledge ascriptions as DeRose's utterance of 'I know that the bank will be open' in LOW – by rejecting the thesis of *intellectualism*. The resulting position is often termed *subject sensitive invariantism (SSI)*. This position appears to improve upon that of the traditional invariantist in re our judgments concerning the relationship between knowledge and further inquiry. However – as is familiar from the literature – there appear to be manifestations of the HIGH-LOW contrast that the subject sensitive invariantist is unable to account for (at least, not without the attribution of significant speaker error). The rejection of intellectualism also appears to entrain problematic predictions of its own. I explore attempts to account for the problematic manifestations the HIGH-LOW contrast – by appeal to both pragmatic and psychological mechanisms – and to alleviate the problematic consequences associated with the rejection of intellectualism. These attempts are found to be unsuccessful. Furthermore, it is shown in the final sections of Chapter 4 that SSI is beset by parallel dilemmas to those that beset traditional invariantist with respect to our judgments concerning epistemic chance and possibility (and principles (CK) and (PK)), and with respect to principle (MPC).

The upshot of Chapter 4 is therefore that the rejection of intellectualism does little to assist the invariantist in evading the attribution of significant error to speakers. In Chapter 5, I argue that a form of *sceptical invariantism* does appear able to avoid the attribution of such error, and that, in consequence, sceptical invariantism represents 'The Best Invariantism'. The form of sceptical invariantism that I explore maintains that in order to know some proposition *p*, a subject must occupy a *very* strong epistemic position with respect to *p*. It is shown that this position is afforded significant support by (a) our judgments pertaining to knowledge and the appropriateness of further inquiry, (b) our judgments pertaining to the relationship between knowledge and epistemic chance and possibility (and principles (CK) and (PK)), and (c) considerations pertaining to (MPC). In the course of that discussion, it becomes apparent that the sceptical invariantist faces no dilemmas with respect to the endorsement of principles like (CK), (PK), and (MPC). For example, the sceptical invariantist is able to account for the infelicitousness of utterances of the form 'I know that *p*, but it might be that not-*p*' and 'I know that *p*, but there's a chance that not-*p*' by invoking principles (PK) and (CK), and without having to attribute

significant error to speakers. This stands in stark contrast to the situation faced by traditional and subject sensitive invariantists.

Of course, such advantages would mean little were the sceptical invariantist unable to make sense of our everyday knowledge ascribing practices. Given that the sceptical invariantist maintains that we need to occupy a *very* strong epistemic position with respect to some proposition in order to know it, her view appears to have the upshot that we know very little. It is therefore incumbent upon the sceptical invariantist to explain such judgments as the judgment that DeRose's (false) knowledge ascription in LOW seems true and appropriate. It is suggested that this might be best achieved by appeal to some *pragmatic* mechanism. In outline, the pragmatic account has it that our everyday knowledge ascriptions – although false – convey such truths as that we have very good evidence for the proposition at issue, or that we are close enough to (strictly speaking) knowing the proposition at issue for current purposes. It is proposed that in virtue of conveying such truths, ordinary knowledge ascriptions, like DeRose's knowledge ascription in LOW, might seem true and appropriate. I defend this account against various possible objections, and also address some more general problems for sceptical invariantism. It is suggested that the sceptical invariantist can respond to those objections, and that – in consequence – she can avoid the attribution of significant error to speakers.

The sceptical invariantist's ability to make sense of our ordinary knowledge ascribing practices, to address potential objections, and to avoid the problems that beset traditional and subjective sensitive invariantism – all whilst avoiding the attribution of significant speaker error –, strongly suggests that sceptical invariantism represents 'The Best Invariantism'. This is the central contention of the thesis. However, such a conclusion clearly raises questions about the relationship between sceptical invariantism and contextualism. In the final section of Chapter 5, I therefore undertake to reveal – in at least a preliminary way – the kind of arguments and judgments that the sceptical invariantist might sensibly appeal to in attempting to show the superiority of her view to that of the contextualist.

Chapter 1: Traditional Invariantism and Experimental Psychology.

As noted in the Introduction, one of the most important pieces of evidence cited in favour of contextualism about ‘knows’, and against traditional forms of invariantism, concerns a certain contrast – ‘the HIGH-LOW contrast’ – in our judgments about cases. As with any appeal to judgments about cases, in evaluating this appeal we should be sensitive to subtle pragmatic effects and to quirks in our psychological processes of judgment. This chapter concerns itself with the proposal that the traditional invariantist is able to account for the HIGH-LOW contrast by appeal to certain experimentally supported features of our psychological processing. If the proposal is correct, it seems that the traditional invariantist is able to explain the HIGH-LOW contrast in such a way as to avoid the attribution of significant error to speakers. (That is, in such a way as to avoid claiming that speakers exhibit a faulty grasp of the meaning of ‘knows’, or some other, related expression.) The central goal of this chapter is to review the relevant psychological proposals, and argue on a variety of grounds that they fail.

The chapter is divided into five main sections. In §1, I set out the explanatory task faced by the traditional invariantist in regard to the HIGH-LOW contrast. In §2, I consider explanatory proposals centred around an appeal to the availability heuristic. In §3, I consider a related proposal, based on the claim that we tend to weight salient considerations more heavily. In §4, I introduce and discuss some further *prima facie* problematic judgments for the traditional invariantist, concerning the relationships between knowledge and inquiry, and knowledge and chance. In §5, I consider some more developed psychological proposals, found in the work of Jennifer Nagel (2008, 2010a, 2010b, 2011), which emphasise the role that outright belief, alongside our tendency to project pertinent features of our psychological situations onto others, might play in accounting for the judgments discussed.

§1. The Issues.

Let us first consider the salient features of the HIGH-LOW contrast. In both the HIGH and LOW cases, an attributor is considering whether a subject, who occupies at least a

moderate epistemic position with respect to some proposition p , knows that p .¹ In the typical HIGH case, the attributor has a lot at stake in whether or not p is true and is entertaining various relatively distant or unlikely error possibilities. In the typical LOW case, the attributor does not have a lot at stake and is not entertaining distant or unlikely error possibilities. It is widely alleged that it seems true and appropriate for the attributor in LOW to *attribute* knowledge that p to the subject, but true and appropriate for the attributor in HIGH to *deny* knowledge that p to the subject. Similar judgments are alleged to arise if the only variation across the HIGH and LOW cases is a difference in the salience of relatively distant or unlikely error possibilities to the attributor or if the only variation is in how much is a stake for the attributor.

There are a host of examples that fit this particular HIGH-LOW pattern – DeRose’s Bank Case, Cohen’s Airport Case, and variations on Dretske’s Zebra Case are amongst the most widely discussed. In what follows, I shall focus – for the most part, though not exclusively – upon DeRose’s Bank Case, presented in the Introduction. Recall that in that example, DeRose’s knowledge ascription seems true and appropriate in Bank Case A – the ‘LOW’ context – but his knowledge denial seems true and appropriate in Bank Case B – the ‘HIGH’ context. How does this kind of HIGH-LOW pattern fit with traditional invariantism?

The traditional invariantist denies that ‘knows’ is context sensitive and that whether a subject has knowledge is constitutively tied to factors like the costs of her being in error or the error possibilities salient to her. It is also a feature of her view that – alongside (ungettiered) true belief – a moderate epistemic position, like the epistemic position occupied by DeRose in the Bank Case, is typically sufficient for knowledge. For the traditional invariantist then, an attribution of knowledge in both the LOW case and the HIGH case will be true, and the corresponding denials will be false. The traditional invariantist position is therefore compromised if she cannot (at least) explain away the appearance of truth and appropriateness – henceforth, ‘correctness’ – in a denial of knowledge in HIGH. It is to our psychological processes of judgment that some philosophers have turned in their attempts to carry out this task.

¹ Note that the subject and the attributor can be one and the same individual – as in the Bank Case, given above.

A final note before we move on. The traditional invariantist typically assumes a (single-premise) closure principle for knowledge. For present purposes, we can characterise Single-Premise Closure (SPC) as the principle that if S knows p, and p entails q, then S knows q.² A central reason to accept such a principle is to fend off concerns that certain ‘abominable conjunctions’ will come out true – for example, ‘I know that the bank will be open, but I don’t know that it’s hours haven’t changed (in such a way that it’s closed)’.³

The acceptance of (SPC) places additional explanatory burdens upon the traditional invariantist. Given (SPC), if a subject knows some proposition p, she also knows that those error possibilities in which not-p holds do not obtain. For example, if DeRose knows that the bank will be open in HIGH, then he knows that the bank has not changed its hours. The problem is that we judge that the attributor in the HIGH context is also correct to deny that the subject knows that the kind of error possibilities typically salient in the HIGH context do not obtain. For example, it seems not only that DeRose is correct to deny, in the HIGH context, that he knows the bank will be open on Saturday, but also that he is correct to deny that he knows that the bank has not changed its hours. (In fact, the latter kind of judgment often seems more compelling than the former.) It therefore falls to the traditional invariantist to also explain our judgments to the effect that it seems correct, in the HIGH context, to deny that the subject knows that the error possibilities typically salient in that context do not obtain.

§2. The Availability Heuristic.

Initial psychological proposals focused heavily on the salience of error possibilities. One prominent suggestion was that we overestimate the likelihood of various error possibilities obtaining as a result of our reliance on the availability heuristic

² It is notoriously difficult to provide counter-example free formulations of epistemic closure principles. Such difficulties – e.g. concerning a lack of knowledge of the relevant entailment – are not central to the present discussion. See Hawthorne (2004a, p. 31-36) for some relevant discussion, and for a ‘competent deduction’ closure principle that improves on the one given in the main text.

³ See DeRose (1995, p. 28). For further defence of closure, see esp. Hawthorne (2004a, ch. 1; and 2005c).

(Hawthorne 2004a, p. 163-164).⁴ In judging how likely an event is, we often rely on how easy it is to imagine and recall events of that kind (on how ‘available’ they are). In general, this strategy works quite well. The kinds of events that occur more often are those we are more exposed to, and this exposure generally makes those kinds of events easier to imagine and recall. But reliance on availability can also lead to distortions in our judgments. In particular, it can lead to distortions in our judgments when how easy it is to imagine or recall an event is not an accurate reflection of how likely it is. For example, excessive media exposure to terrorist incidents can make such incidents very easy to imagine and recall. Consequently, we tend to overestimate the likelihood that they will take place.

There are a number of possible sources of excessive availability. Disproportionate exposure to particular kinds of event or reports of them is one. Thinking carefully about the reasons why a particular event might come about (and not thinking carefully about why certain alternatives to that event might come about) is another. The basic ‘availability defence’ of traditional invariantism is founded in the idea that, in the HIGH case, misleading availability leads us to overestimate the likelihood of some salient error possibility obtaining. (In contrast, we don’t even consider the likelihood of that possibility obtaining in the LOW case.)

As noted above, there are (at least) two *prima facie* problematic judgments concerning the HIGH-LOW cases that the traditional invariantist has to account for. In the Bank Case, the traditional invariantist has to explain our judgment that DeRose seems correct to claim that he doesn’t know the bank has changed its hours in HIGH, and our judgment that he seems correct to claim that he doesn’t know that the bank will be open tomorrow in HIGH.⁵ Presumably, the source of our judgment to the effect

⁴ Hawthorne (2004a, p. 162-165) was trying to defend a form of subject-sensitive invariantism, but his suggestions can clearly be adapted in defence of traditional invariantism. Insofar as one tries to defend subject-sensitive invariantism by appeal to the availability heuristic, one will encounter many of the same problems discussed here (see Chapter 4, §4 for relevant discussion). (In fairness, it should be noted that Hawthorne explicitly states (2004a, p. 164, fn. 9) that he does not intend an appeal to the availability heuristic to account for all the relevant phenomena.)

⁵ Notice that even those traditional invariantists who deny (SPC) will (presumably) still have to account for our judgment to the effect that it seems correct for DeRose to deny himself knowledge that the bank will be open in HIGH. It will become apparent in the following discussion that doing so by appeal to the availability heuristic, or more generally, by appeal

that DeRose seems correct, in HIGH, to deny himself knowledge that the bank has not changed its hours is straightforwardly our overestimation of the likelihood that it has. The source of our judgment to the effect that he seems correct to deny himself knowledge that the bank will be open in HIGH is then presumably *either* some (possibly implicit) closure-based inference *or* an *underestimation* of the likelihood that the bank will be open tomorrow, brought about by our overestimation of the likelihood of various alternatives. This general species of explanation may then be extended to HIGH cases in which only the stakes are raised (and no error possibilities are explicitly made salient) on the assumption that an increase in stakes naturally leads to additional error possibilities being entertained.

There are a range of difficulties with this proposal. I first consider some problems localised to appeals to the availability heuristic. I then go on to consider some more general problems for any explanation based around the idea that we overestimate the likelihood of pertinent error possibilities obtaining.

Research shows that if we are told about the effects of the availability heuristic, or we are given some reason to think that availability might have a misleading source, the distortions in our judgments tend to disappear (Nagel 2010a, p. 298). However, our knowledge judgments seem largely unaffected when we are told about the potential effects of the availability heuristic. Similarly, when we are given a reason to think that the relevant availability might be misleading – for example, if we are warned that the events featured in error possibilities might be very easy to imagine because we have been over-exposed to them in the media – our judgments do not seem significantly affected. This suggests that our putative reliance on the availability heuristic does not lie at the heart of those judgments.

There is also the difficulty of accounting for the *source* of the excess availability in question. It may well be that we are overexposed to reports of dramatic incidents like terrorist attacks and record lottery wins. But what about banks changing their hours? Or being confronted with outdated airline schedules? Or run-of-the-mill misprints in the newspaper? It doesn't seem so plausible that we are misleadingly over-exposed to

to the idea that we overestimate the likelihood of various error possibilities obtaining, is very problematic.

such events (or to reports of them). Nevertheless, the traditional invariantist still needs to explain why we judge that it would seem correct to deny knowledge that more run-of-the-mill error possibilities don't obtain. Similarly, she also needs to explain why, in HIGH contexts in which only more run-of-the-mill alternatives to *p* are raised, it typically seems correct to deny that the subject knows that *p*. For example, in the HIGH context of the Bank Case, the traditional invariantist needs to explain why it seems correct for DeRose to deny himself knowledge that the bank will be open, even though only the more run-of-the-mill possibility that the bank has changed its hours has been raised. If the traditional invariantist cannot make the case that there is some *other* source of excess availability in play in such cases, besides past over-exposure to the relevant kind of event (or reports of it), the availability defence might seem to be in trouble.

Hawthorne (2004b, p. 521-522) suggests that *discussing* possibilities may also increase our assessment of their likelihood, because discussing possibilities makes them easier to imagine.⁶ It might therefore be suggested that the reason we overestimate the likelihood of even run-of-the-mill error possibilities obtaining in the HIGH case is that those error possibilities are being discussed.

This suggestion faces some pressing challenges. First, as Nagel (2010a, p. 291-293) points out, even engaging in sustained discussion of a possibility can sometimes lead us to judge that the likelihood it obtains is *lower* than we would otherwise have judged it to be.⁷ The effects of sustained discussion do not always appear to be positive – they do not always lead to higher estimates of likelihood.

In any case, in typical HIGH contexts, it seems that sustained discussion of the error possibilities raised in those contexts is not even necessary to bring about our negative knowledge judgments. In at least many cases, it seems that the relevant error possibilities need only be *entertained* in order to have that negative effect. However, research does not support the conclusion that we uniformly overestimate the

⁶ Hawthorne is drawing on some remarks in Slovic *et al.* (1982, p. 465): one 'particularly important implication of the availability heuristic is that discussion of a low-probability hazard may increase its memorability and imaginability and hence its perceived riskiness'.

⁷ See Sherman *et al.* (1985) for relevant empirical work.

likelihood of possibilities we merely entertain. As remarked earlier, absent such factors as past media overexposure, our reliance on the availability heuristic tends to give rise to comparatively accurate judgments that a given event will occur. But in order to judge the likelihood that some event will occur, we presumably have to entertain the possibility that it does. This suggests that our reliance on the availability heuristic tends to give rise to comparatively accurate judgments that a given event will occur *even when we have entertained the possibility that it will*. Furthermore, empirical research indicates that in some cases we are actually prone to *underestimate* the likelihood that possibilities that we have entertained will obtain – for example, in those cases in which we have been misleadingly underexposed to the relevant kind of event.⁸

Such considerations cast significant doubt on the adequacy of Hawthorne's suggestion. They also cast significant doubt on a natural alternative proposal – viz. that the source of excess availability with respect to more run-of-the-mill error possibilities lies simply in the fact that, in the relevant HIGH contexts, we are entertaining those possibilities. Failure to find a plausible source for the misleading availability at issue represents another potential problem for the availability-based account.

There are also some more general problems for any explanatory strategy based around the idea that we overestimate the probability of various error possibilities obtaining (whether from misleading availability or some other source).⁹

First, in many of the relevant HIGH cases, it seems correct to deny that the subject knows even *very unlikely* error possibilities not to obtain. In Dretske's Zebra Case, for example, the likelihood that the animal in the pen is a cleverly disguised mule is presumably extremely small. Even if we overestimate the likelihood that such

⁸ Cf. Cohen (2004a, p. 489-490). For relevant psychological work, see e.g. Slovic *et al.* (1982).

⁹ Williamson (2005a) appears to exhibit some sympathies with the idea that our judgments lie in an overestimation of the likelihood of error possibilities obtaining:

One effect of fictional violence on television is to make viewers overestimate their chances of being the victims of violent crime: they suffer an illusion of danger. Might not an illusion of epistemic danger result from exposure to lurid stories about brains in vats, evil demons, painted mules, or gamblers who bet the farm? (2005a, p. 226)

possibilities obtain, our estimate is presumably still going to be very small. It's far from clear that this is sufficient to explain our judgments to the effect that it oftentimes seems correct to deny knowledge that such possibilities do not obtain. This problem is made particularly stark when we reflect that our estimate of the likelihood that, for example, the animal in the pen is not a cleverly disguised mule will almost certainly be higher than our estimate of the likelihood that the bank will be open tomorrow in the LOW context of the Bank Case. But then why does it seem correct to deny that the subject knows that the animal in the pen is not a cleverly disguised mule, whilst it nevertheless seems correct for DeRose to claim (in LOW) to know that the bank will be open?

Second, even when the relevant chances are made explicit in the HIGH case, our judgments about knowledge seem to remain much the same.¹⁰ For example, even if it is stipulated that the chance that the bank has changed its hours is 1 in 1000, it nevertheless seems correct for DeRose to deny that he knows the hours haven't changed. It's rather less clear how it could be overestimation of a relevant likelihood that is at the root of that judgment.

Taken together these concerns put significant pressure on the availability defence of invariantism.¹¹ The last two concerns also pose a significant challenge to any account that seeks to locate the source of our judgments in an overestimation of the likelihood of some error possibility obtaining. In §4, I shall consider some further *prima facie* problematic judgments for the traditional invariantist, and argue that such accounts are also unable to assist in alleviating the threat posed by those judgments. For the present, however, let's consider another proposal.

§3. The Saliency Account.

Williamson (2005a, 2005b) suggests a rather different account that is *prima facie* not subject to the same kinds of objections. Williamson observes that we are prone to weight salient considerations more heavily. He therefore proposes that when various

¹⁰ See Stanley (2005a, p. 101), Nagel (2010a, p. 299-300).

¹¹ Nagel (2010a) offers a wealth of further objections to the availability strategy. (However, I should remark that I do not find many of those objections especially persuasive.)

error possibilities are made salient to us, the existence of those possibilities weighs more heavily on our judgment:

When ... little is at stake and the sceptic is not bothering us, few possibilities of error are psychologically salient: no wonder we do not doubt that cases which are or seem similar to socially accepted paradigms of knowledge (perception, memory, testimony, ...) are indeed cases of knowledge (or whatever privileged epistemic status is at issue). When ... much is at stake or the sceptic is bothering us, more possibilities of error are psychologically salient: no wonder we give the newly salient cases more weight than before in assessing claims to knowledge (or whatever privileged status is at issue). (2005b, p. 112)

On the face of it, this story need not include any claim to the effect that we misjudge some relevant likelihood. The idea is simply that the fact that there *are* error possibilities – however likely or unlikely we may judge them to be – of various sorts is something to which we give weight in our judgments about HIGH, but not LOW. This explains the contrast in our judgments.

There are a number of difficulties with the view.

First – as came to the fore in the previous discussion – in some cases it is clear that the likelihood that some error possibility obtains is *very* small. Consider, for example, the possibility that Barack Obama is one of the unlucky few individuals who, despite appearing in good health, suffered a fatal heart-attack in the last few minutes.¹² Such a possibility is very unlikely indeed. Nevertheless, unless one is stood next to Obama or watching a live presidential broadcast (or similar), it can seem correct to deny that one knows that Obama did not suffer such an attack. In the same vein, once such a possibility has been raised, it can seem correct to deny that one knows that Obama is still president. (It may nevertheless seem correct to claim that one knows that it is exceedingly likely that Obama is still president.)

However, the traditional invariantist must allow that even *significantly more likely* error possibilities than the possibility that Obama has suffered such a fatal heart attack represent no threat to our knowledge. For example, in the Bank Case, the traditional invariantist allows that DeRose knows that the bank has not changed its hours.

¹² The example is adapted from one in Hawthorne (2004a, p. 3).

Nevertheless, it is significantly more likely that the bank has changed its hours than that Obama has suffered an unexpected and fatal heart attack in the last few minutes. Such low likelihood possibilities just don't seem to come close to threatening our knowledge on the traditional invariantist account. In a similar vein, notice that such possibilities arise even with respect to the kinds of 'paradigm cases' of knowledge that Williamson alludes to. For example, in paradigm cases of testimony, there is a small chance that the relevant individual (however trustworthy they may be in general) is lying, and in paradigm cases of perceptual knowledge, there is a small chance that the relevant faculties are delivering misleading appearances. It is therefore far less clear why taking into account very low-likelihood error possibilities should lead us to judge *either* that it seems correct for the attributor to deny that the subject knows that those low-likelihood alternatives themselves do not obtain (e.g. that Obama has not suffered an unexpected and fatal heart attack in the last few minutes), *or* that it seems correct for the attributor to deny that the subject knows the proposition at issue (e.g. that Obama is still president).

Second, one might expect that consideration of various ways in which the subject might be led into error would yield a *better* understanding of her epistemic predicament – in other words, to more accurate judgments about knowledge.¹³ Indeed, empirical research supports the conclusion that considering alternative possibilities leads to more accurate judgments concerning epistemic properties closely related to knowledge.

A wide variety of studies have shown that we are prone to allow aspects of our own epistemic position to unduly influence our assessments of the epistemic positions of others.¹⁴ In one particular manifestation of this, we become 'infected' by information about outcomes. If we are assured that the outcome of some process is *o*, we are liable to overestimate the extent to which the evidence available to the subject (which does

¹³ Cf. Nagel (2010a, p. 300-301). Nagel notes that it is not clear why taking error possibilities explicitly into account should lead us to less, rather than more, accurate judgments in regard to knowledge, but does not offer further argument to support the claim that it should lead to greater accuracy.

¹⁴ For reviews of relevant empirical studies, see e.g. Royzman *et al.* (2003), Birch and Bloom (2004), and Nickerson (1999). See Nagel (2008, 2010a, 2010b) for philosophical application of such work.

not include the assurance we have) supports o.¹⁵ This suggests that when provided, for example, with the outcome in the Bank Case – that the bank is open on Saturday – those who read the case description will be liable to overestimate how probable DeRose’s evidence makes that outcome.

Why is this relevant? It’s relevant because it turns out that thinking about alternatives is a good way to attenuate the effects of this bias.¹⁶ In terms of the effects on our judgments, this suggests that readers of the HIGH-LOW cases will be liable to overestimate how likely the subject’s evidence makes the proposition at issue when presented with just the LOW case, but liable to make more accurate judgments of that likelihood when presented with the corresponding HIGH case, which encourages thinking about alternative possibilities.¹⁷ For example, in the Bank Case, it suggests that readers are liable to make a more accurate judgment concerning the likelihood, on DeRose’s evidence, that the bank will be open, when confronted with the HIGH case, because they are encouraged to take into account the possibility that the bank has changed its hours.

The important point here is simply that we have reason to expect that thinking about alternatives – the very thing that the salience account alleges leads us to make mistakes – will lead to *more* accurate judgments with respect to epistemic properties in the neighbourhood of knowledge. In particular, we have reason to think that it will lead to more accurate judgments in regard to the likelihood of the proposition central to the discussion (e.g. that the bank will be open on Saturday) on the subject’s

¹⁵ For further details, see the discussion of epistemic egocentrism in §5.2.

¹⁶ See e.g. Koriat *et al.* (1980), Slovic & Fischhoff (1977). For more recent discussion, see e.g. Sanna & Schwarz (2003).

¹⁷ In the relevant empirical research (see previous fn.), ‘thinking about possibilities’ typically requires more than simply entertaining them – it involves such practices as thinking about (or writing down) which aspects of evidence might support the possibility in question, or generating the alternative possibilities oneself. It would seem to follow that if error possibilities are simply entertained in the HIGH case, we may not witness a very significant difference between our likelihood judgments when presented with just the LOW case, and our likelihood judgments when presented with the HIGH case. However, even if that’s the case, it’s not much comfort for Williamson. Firstly, it doesn’t mean our judgments aren’t more accurate (if only slightly). Secondly, it just shows that *despite* considering alternatives we may well remain *overoptimistic* with respect to DeRose’s possession of epistemic properties in the neighbourhood of knowledge. This is still in at least *prima facie* tension with saying that, in the same circumstances, we exhibit *excessive pessimism* with respect to DeRose’s possession of knowledge.

evidence. That we are told that thinking about alternatives has precisely the opposite effect on our judgments concerning the subject's *knowledge* of that proposition is therefore troubling.

At the very least, these two concerns put pressure on those sympathetic to a proposal centred around the claim that we weight salient considerations more heavily to expand on the details of the account. In the next section, I leave the HIGH-LOW contrast aside, and discuss some further *prima facie* problematic judgments concerning knowledge and inquiry, and knowledge and chance. It is shown that neither the proposal considered in this section, nor the proposal considered in the previous section, can assist in alleviating the threat to traditional invariantism posed by those judgments.

§4. Knowledge and Inquiry.

Utterances of the form 'I know that p, but I need to inquire further' and 'I know that p, but we need to check into that further' typically seem infelicitous. For instance, an utterance by DeRose in HIGH of 'I know that the bank will be open, but we need to inquire further' or 'I know the bank hasn't changed its hours, but we need to check into that further' seems straightforwardly infelicitous. (Indeed, in regard to a similar case, Cohen (1999, p. 60) describes such utterances as apparent contradictions.) The natural explanation of this infelicity is that utterances of the form 'I know that p, but I need to inquire further' express contradictions, and – in particular – that knowledge *closes inquiry*, in the sense that if a subject knows some proposition p, there's no need for her to inquire further into whether or not p is true.

Such a connection between knowledge and inquiry actually appears to fall out of recent proposals linking knowledge and appropriate action/practical reasoning. Of particular interest here is the suggestion that a principle in the vicinity of the following is true:¹⁸

¹⁸ See, *inter alia*, Hawthorne (2004a, p. 30), Hawthorne & Stanley (2008), Fantl & McGrath (2002, 2007, 2009), and Williamson (2005a, p. 227).

- (1) a. If S knows that p, then S can appropriately rely on p in her practical reasoning.
- b. If S knows that p, then S can appropriately treat p as a reason for acting.
- c. If S knows that p, then S can appropriately act as if p.

If knowing a proposition guarantees that one can rely on that proposition in one's practical reasoning (or if (1b) or (1c) holds), it's not clear on what grounds one might allege that, although a subject knows a proposition, further inquiry into its truth is still necessary. Thus, (1a) – and, for similar reasons, (1b) and (1c) – appear to entail that knowledge closes inquiry. Such principles therefore gain some support from our judgments concerning the infelicity of utterances of the form 'I know that p, but we need to inquire further'. These principles are also lent more direct support by the observation that utterances such as those of the form 'I know that p, but I cannot rely on p in my practical reasoning' also seem infelicitous. For example, it seems infelicitous for DeRose to utter 'I know that the bank will be open, but I cannot rely on that in my practical reasoning'. The natural explanation for this judgment is that some principle like (1a-c) is true.

Williamson (2005a) suggests that the traditional invariantist might be able to endorse (1a). If that's right, it seems that the traditional invariantist is able to maintain that knowledge closes inquiry, and is free to suggest that the reason utterances of the form 'I know that p, but I need to inquire further' and 'I know that p, but I cannot rely on p in my practical reasoning' typically seem infelicitous is that they express contradictions.

Perhaps the most pressing problems with traditional invariantist endorsement of (1a), (or the claim that knowledge closes inquiry), concern subjects who have a great deal at stake. Recall that the traditional invariantist denies that knowledge varies with how much is at stake, and maintains that a modest epistemic position (alongside (ungettiered) true belief) is all that is required for knowledge. The problem is that whilst in some circumstances a modest epistemic position can indeed seem sufficient

for a subject to close inquiry (or rely on some proposition in her practical deliberations), in other circumstances – those in which a great deal is at stake – it seems that a subject needs to occupy a very considerable epistemic position in order to appropriately close inquiry (or rely on that proposition in her practical deliberations). For example, the traditional invariantist maintains that in the HIGH context of the Bank Case, DeRose knows that the bank will be open. Nevertheless, it seems, in the HIGH context, both that DeRose needs to inquire further, and that he cannot rely on the proposition that the bank will be open in his practical deliberations.

A further problem arises in regard to certain counterfactual judgments. Notice that utterances of the form ‘If there were more at stake, it would not be appropriate to rely on that proposition in your deliberations’ and ‘If there were more at stake, you would need to check further’ typically seem true and appropriate, and that they do so even in circumstances in which the subject already meets the (modest) conditions on knowledge proposed by the traditional invariantist. For example, it seems true and appropriate for DeRose’s wife, in LOW, to utter ‘If you had more at stake, you would need to check further’. However, on the assumption that knowledge closes inquiry (and on the assumption that some principle like (1a) is true), such counterfactual claims should presumably seem false.

In response to such concerns, Williamson (2005a, p. 229-231) suggests that whilst knowledge is indeed sufficient to assure the rationality of relying on some proposition in one’s practical reasoning, there can also be certain *derivative* norms in play. In essence, his proposal is that in high stakes situations it may be necessary, in some derivative sense, to make sure that one’s practical deliberations are rationally respectable. How sure? That depends on the stakes – very high stakes might demand considerable assurance that it is rationally acceptable to rely on the relevant proposition(s) in one’s practical reasoning. Adopting Williamson’s suggestion, it might then be proposed that the reason why it seems that higher stakes subjects need to engage in inquiry *beyond* that required for knowledge (on the traditional invariantist account), is that our ordinary uses of expressions such as ‘appropriate’ or ‘need to’ express some derivative, stakes-sensitive, sense of propriety. This could also explain why counterfactuals such as ‘If there were more at stake, it would be

appropriate to check further' can seem true, even supposing that the subject meets the traditional invariantist conditions on knowledge.

This response is problematic. A central concern is that the response threatens to undermine the explanation for why utterances such as those of the form 'I know that p, but I need to inquire further' seem infelicitous. If our ordinary uses of expressions such as 'appropriate' and 'need to' are picking out some derivative, stakes-sensitive sense of propriety, then presumably they would do so in utterances such as 'I know that the bank will be open, but we need to inquire further'. But if that's the case, they ought to seem felicitous.

The obvious remedy is to suggest is that in the case of utterances such as 'I know that the bank will be open, but we need to inquire further', the reading on which the derivative sense of propriety is picked out is not a natural one. But why not? And why don't we find ready readings of counterfactuals such as 'If there were more at stake, we would need to check further' on which they pick up on the *non*-derivative sense of propriety (a reading on which they would oftentimes be false)?

Furthermore, it seems that the advocate of this remedy will be forced to impute various problematic equivocation errors to speakers. The reason is that it does not seem, for example, as though 'need to inquire further' has a different meaning when paired in a sentence with 'know' than when it features in counterfactuals such as 'If the stakes were higher, I would need to inquire further'. Consider an example. Suppose that, in the LOW context of the Bank Case, DeRose utters both 'If we stood to lose the house, we would need to check further' and 'If we stood to lose the house, then I would still know that the bank would be open, so we wouldn't need to check further'. On the proposed remedy, an utterance of the former counterfactual should pick out the derivative sense of propriety in its consequent, whilst an utterance of the latter counterfactual should pick out the non-derivative sense of propriety in its consequent. If that's so, they can be consistently endorsed. But it does not seem as though those counterfactuals can be consistently endorsed.¹⁹

¹⁹ It might be contended that this is not really a consequence of the natural remedy. In particular, it might be suggested that because these counterfactuals are themselves being considered together, their consequents either both concern the derivative sense of inquiry or

The appeal to derivative notions of propriety looks unlikely to relieve the difficulties associated with a traditional invariantist endorsement of (1a). Is there an alternative? It might be proposed that the kind of psychological explanations we have been discussing in regard to the HIGH-LOW contrast can assist the traditional invariantist in alleviating the problems apparently associated with endorsing (1a). Recall that it was suggested earlier that rising stakes leads to the consideration of additional error possibilities, and that – for various possible reasons – this leads us to mistakenly judge that people lack knowledge. It might therefore be proposed that (1a) is true, but that the reason that we (mistakenly) judge that certain high stakes subjects are unable to rely on the relevant proposition(s) in their practical reasoning, or to close inquiry with respect to those proposition(s), is tied to the fact that we (mistakenly) judge that those subjects don't know those proposition(s).

This strategy might most naturally be cashed out by assuming that knowledge that p is both sufficient *and necessary* for a subject to rely on p in her practical reasoning – that is, that a principle like (1a)* is also true.²⁰

- (1)* a. If S can appropriately rely on p in her practical reasoning, then S knows p

Similarly, the traditional invariantist could propose that knowledge is both sufficient *and necessary* to close inquiry. This would then enable the traditional invariantist to explain the apparent inappropriateness of relying on some proposition in practical reasoning, or the apparent inappropriateness of closing inquiry with respect to some proposition, in terms of an apparent lack of knowledge of that proposition. On the assumption that certain psychological mechanisms tend to make us judge that high

both concern the non-derivative sense. This might explain why we do not seem able to consistently endorse them both. But it still leads to a dilemma. If they both concern the derivative sense of propriety, then the inference in the consequent of the second counterfactual should seem invalid, which it doesn't. But if they both concern the non-derivative sense of propriety, the first counterfactual should seem false, which it doesn't.

²⁰ Williamson (2005a, p. 227) seems to endorse both (1a) and (1a)*. (See also Hawthorne (2004a, p. 30) and Hawthorne and Stanley (2008).) However, the criticisms to follow would also cause trouble for attempts to appeal to the psychological mechanisms considered thus far that do *not* involve appeal to (1a)*.

stakes subjects lack knowledge that they in fact have, the traditional invariantist can then readily explain our judgments to the effect that those high stakes subjects cannot close inquiry. For example, the traditional invariantist can propose that when confronted with the HIGH context of the Bank Case, psychological bias leads us to (mistakenly) judge that DeRose does not know the bank will be open, and therefore (mistakenly) judge that he needs to engage in further inquiry. In contrast, this is not a mistake we make in the LOW context, in which we judge that DeRose knows that the bank will be open, and therefore – given (1a) – judge that he does not need to engage in further inquiry.

One problem with this proposal is that it seems not only that our judgments about the appropriateness of closing inquiry are sensitive to the stakes. We also judge that whether or not it is appropriate to close inquiry is something that is itself determined by the stakes. In particular, we judge that high stakes subjects require significantly more evidence to close inquiry than do low stakes subjects. However, on the view under consideration further inquiry is not something that varies with the stakes. Furthermore, we judge that counterfactual utterances such as ‘If there more at stake, it would be necessary to inquire further’ are true and appropriate. But on the view under consideration, such claims are oftentimes false.

It might be replied that the traditional invariantist can also appeal to the psychological explanations to explain these judgments. The natural suggestion would be that we recognise that our judgments about appropriate inquiry exhibit sensitivity to how much is at stake (but presumably don’t also recognise that this represents an error on our part). This could account both for our judgments to the effect that appropriate inquiry is determined by the stakes, and for our judgments to the effect that counterfactuals like ‘If there were more at stake, it would be necessary to inquire further’ seem true.

This reply looks problematic. One central problem is that if it works in the case of our judgments concerning appropriate closure of inquiry (and appropriate reliance on a proposition in practical reasoning), it seems that it should *also* work in the case of our judgments about knowledge. After all, those judgments vary in just the same way that our judgments concerning inquiry do. (Indeed, on the proposal under consideration,

our judgments about inquiry are straightforwardly grounded in our judgments about knowledge.) It therefore seems that we should also expect counterfactuals like ‘If there were more at stake, you would not know’ to seem true and appropriate. However, such counterfactuals *do not* strike us as true and appropriate – they strike us as false. Similarly, it seems that we should also expect it to seem as though knowledge were determined (at least partially) by the stakes. But it does not.²¹

It therefore seems that *neither* appeal to derivative senses of propriety, *nor* appeal to the psychological proposals considered thus far, is able to assist the traditional invariantist in responding to the problems encountered with granting that knowledge closes inquiry, or endorsing a principle like (1a). This raises another question: Is it possible for the traditional invariantist to explain the infelicitousness of utterances like ‘I know that the bank will be open, but we must inquire further’, and ‘I know that the bank will be open, but I can’t rely on that in my practical reasoning’ *other than* by appeal to the claim that knowledge closes inquiry, or appeal to (1a)?

By rejecting (1a), the traditional invariantist is able to embrace the idea that the appropriateness of relying on some proposition in practical reasoning, and the appropriateness of closing inquiry, are genuinely stakes sensitive notions. She can therefore explain our judgments to the effect that they are. She is also able to claim that the reason that it seems inappropriate for high stakes subjects, like DeRose in HIGH, to close inquiry, or to rely on pertinent propositions in their practical reasoning, is because it really is inappropriate for them to do so. However, it seems that she must then accept that utterances of the form ‘I know that p, but we must inquire further’ are oftentimes true. For example, it seems that she must accept that an utterance of ‘I know that the bank has not changed its hours, but we need to inquire into that further’ by DeRose in HIGH is true. In the remainder of this section, I consider whether the psychological proposals considered thus far can assist the traditional invariantist in explaining why – although true – such utterances seem infelicitous.

²¹ The reader may notice that the view suffers from parallels of many of the problems that anti-intellectualist accounts suffer from (such problems are discussed in Chapter 4). The contrast is that whilst in the case of anti-intellectualism the problems are grounded in the fact that knowledge is alleged to be stakes sensitive whilst many of our judgments suggest it’s not, in the case of ‘stakes-insensitive appropriateness’ it’s rather that appropriate inquiry is alleged *not* to be stakes sensitive, whilst many of our judgments suggest that it *is*.

Consider first those psychological accounts centred around the availability heuristic (§2, above). The central claim of such accounts was that when confronted with the HIGH case, we overestimate the likelihood that the bank has changed its hours. Thus, the reason that we judge that DeRose does not know that the bank has changed its hours in the HIGH case, is because we wind up *underestimating* the extent to which his evidence supports the proposition that the bank has *not* changed its hours. The guiding idea is therefore that whilst we do not make a mistake about how likely the proposition that the bank has not changed its hours *needs to be* in order for DeRose to know that proposition, we do make a mistake about how likely it *is* that the proposition is true. This is how the proposal enables the traditional invariantist to avoid the attribution of significant error to speakers – that is, how it enables her to claim that speakers do not exhibit a faulty grasp of the meaning of ‘knows’.

However, just as it is a feature of the account that we don’t make a mistake about how likely the proposition that the bank has not changed its hours needs to be in order for DeRose to know that proposition, it seems to be no feature or consequence of the account that we should expect speakers to make a mistake about how likely the proposition that the bank has not changed its hours needs to be in order for DeRose to rationally close inquiry in regard to it. The account simply provides us with some reason to think we will make a mistake about how likely it *is* that the bank has not changed its hours. There is therefore no reason to expect that we will make a mistake about the fact that the epistemic position required for DeRose to close inquiry in HIGH is *better* than that required for knowledge (on the traditional invariantist account). It follows that there should be no apparent contradiction in DeRose’s utterance, in HIGH, of ‘I know that the bank has not changed its hours, but we need to inquire into that further’. The account centred around appeal to the availability heuristic therefore offers little assistance in explaining the infelicitousness of that utterance. (For similar reasons, any account centred around the claim that we overestimate the likelihood that pertinent error possibilities obtain would also seem to offer little assistance to the traditional invariantist.)

Let’s now turn to the account – the Saliency Account – centred around the claim that we tend to weight salient considerations more heavily (§3, above). Interestingly, it

seems that those sympathetic to the Saliency Account can provide an explanation for at least the infelicitousness of an utterance, by DeRose in HIGH, of ‘I know that the bank is open, but we need to check further’. She need only embrace two assumptions:

(a) The salience of the possibility that the bank has changed its hours does *not* lead us to make a mistake about the strength of DeRose’s epistemic position (characterised in terms other than knowledge), nor any mistake about the epistemic position (characterised in terms other than knowledge) required for DeRose to close inquiry. For instance, the salience of the possibility that the bank has changed its hours does not lead us to make a mistake about the likelihood on DeRose’s evidence that the bank will be open; nor does it lead us to make a mistake about how much evidence DeRose must gather in order to rationally close inquiry.

(b) The extent to which DeRose needs to improve his epistemic position in order to correct for the distorting influence, on our judgments about knowledge, associated with taking the possibility that the bank has changed its hours into account *is greater than or equal to* the extent to which DeRose needs to improve his epistemic position in order to rationally close inquiry. For instance, the amount of evidence DeRose needs to gather, in order for us to judge that it seems true for him to utter ‘I know that the bank will be open’ in HIGH, is greater than or equal to the amount of evidence DeRose needs to gather in order to be able to rationally close inquiry.

Given both of those assumptions, the prediction would be that DeRose’s utterance, in HIGH, of ‘I know that the bank will be open, but we need to inquire further’ should indeed seem contradictory.

There are, however, good reasons to be sceptical about both assumptions. First, it is not clear why the improvement in DeRose’s epistemic position required to correct for the distorting influence of considering the possibility that the bank has changed its hours should be at least as great as the improvement required in order for DeRose to rationally close inquiry. Notice that given his very high stakes, the latter improvement will have to be quite considerable – why expect the former improvement to be equally large?

Second, as came out in the discussion of §3, it seems uncomfortable to claim that considering alternatives has an excessively negative effect on our judgments concerning DeRose's knowledge, but not on our judgments concerning his possession of other, related, epistemic properties. It therefore seems troubling to maintain that considering the possibility that the bank has changed its hours should have a markedly different effect on our judgments about the appropriacy of DeRose closing inquiry than it does on our judgments about knowledge.²²

But such an assumption is essential. For suppose instead that considering the possibility that the bank has changed its hours has the same effect on our judgments concerning the appropriateness of further inquiry as it does on our judgments concerning knowledge. That is, suppose instead that the improvement in DeRose's epistemic position required in order to counterbalance the distorting influence of considering the possibility that the bank has changed its hours *on our assessment of whether DeRose knows that the bank will be open* is just identical to the improvement in his epistemic position required in order to counterbalance the effect of considering that possibility *on our assessment of whether it is appropriate for him to close inquiry into the matter of the bank's being open*. Just as in regard to the account centred around the availability heuristic, the Salience Account does not seem to supply us with any reason to expect that speakers will make a mistake about the comparative claim that the epistemic position DeRose requires for knowledge is weaker than the epistemic position required for him to rationally close inquiry. It therefore seems that, given our alternative assumption, there should be a way to improve DeRose's epistemic position in the case description (give him more evidence) such that it seems to us both that DeRose knows that the bank will be open, and that he needs to engage in further inquiry. In consequence, it seems that there should be no appearance of contradiction in an utterance, by DeRose in HIGH, of 'I know that the bank will be open, but we need to inquire further'.

A quick summary. Utterances of the form 'I know that p, but we need to inquire further', and 'I know that p, but I can't rely on p in my practical reasoning' typically

²² Note that in his original sketch of the Salience Account (quoted earlier) Williamson explicitly applies the account to 'other privileged epistemic' statuses – being rationally able to close inquiry is plausibly one such status.

seem infelicitous. This suggests that a principle like (1a) is true. However, the traditional invariantist faces some obvious problems in attempting to endorse that principle. It was shown that neither appeal to derivative senses of propriety, nor appeal to the kind of psychological stories considered in §2 and §3, enables the traditional invariantist to alleviate those problems. In order to avoid those problems, the traditional invariantist might therefore seek to reject (1a). However, this leaves her in need of an explanation for why utterances such as ‘I know that the bank will be open, but we need to inquire further’ seem infelicitous, despite being true. It was shown that the psychological mechanisms invoked to account for the HIGH-LOW contrast in §2 and §3 are unable to assist the traditional invariantist in meeting that explanatory task. The concern therefore arises that whether she accepts *or* rejects (1a), the traditional invariantist is forced to postulate significant speaker error. (In this case, the concern is that the traditional invariantist has to claim that we exhibit a faulty grasp of the meaning of ‘knows’, or of such expressions as ‘need to inquire further’, or both.)

§4.1 Knowledge and Chance.

Conjunctions linking knowledge and inquiry are not the only sources of conjunction-related discomfort for the traditional invariantist. A further class of problematic conjunctions – one that we shall look at more extensively in Chapter 3 – concerns the relationship between knowledge and chance. Notice that it typically seems infelicitous to utter conjunctions of the form ‘I know that p, but there’s a chance that not-p’. For example, it seems infelicitous for DeRose to utter ‘I know that the bank will be open, but there’s a chance it won’t be’, or ‘I know the bank’s hours haven’t changed, but there’s a chance that they have’ in either HIGH or LOW.

Such uses of ‘chance’ are most naturally read as being *epistemic* in character – e.g. as expressing a chance associated with the subject’s evidence. The natural explanation for the infelicitousness of such conjunctions therefore invokes a principle in the vicinity of (CK).²³

²³ Cf. Williamson (2000, ch. 9) and Hawthorne (2004a, p. 24-28).

(CK) If S knows that p, then there's no epistemic chance, for S, that not-p.

(CK) furnishes us with a straightforward explanation of the oddity of the conjunctions linking chance and knowledge – they express contradictions. However, the costs associated with traditional invariantist endorsement of such a principle are very high.

First, it seems the traditional invariantist must deny the compelling claim that, for nearly all of what we take ourselves to know, there's a chance, for us, that we are wrong. In a similar vein, it seems that – given the character of his evidence – there is a chance, for DeRose (in LOW and in HIGH), that the bank won't be open, and that the bank has changed its hours. But the traditional invariantist claims that DeRose knows (in LOW and in HIGH) that the bank will be open, and that the bank has not changed its hours. Thus, given (CK), it follows that there is *no* chance, for DeRose (in LOW or in HIGH), that the bank won't be open, or that the bank has changed its hours.

Further problems with a traditional invariantist endorsement of (CK) arise when we disjoin various propositions. Take a fairly large selection of propositions that a typical subject, Peter, takes himself to know about the future – that he will be at his important board meeting next week, that he will own a car next month, that he will not be leaving London before Tuesday, and so on. Let those propositions be p_1 to p_n . Suppose that Peter occupies at least a modest epistemic position with respect to each proposition p_i – for example, Peter has made very careful plans to be at his important board meeting, Peter has no intention to sell his car, and so on – and that, on the traditional invariantist account, he knows each p_i . Given (SPC), Peter knows that the negation of each p_i is false. Given (CK), the traditional invariantist must therefore admit that, for each proposition p_i , there is no chance, for Peter, that $\sim p_i$ is true.

Now consider the disjunction $\sim p_1 \vee \sim p_2 \dots \vee \sim p_n$: Peter will not be at his important board meeting next week, or Peter will not own a car next month, or Peter will not be leaving London before Tuesday, or The traditional invariantist is faced with a choice. She must either (i) admit that there is no chance, for Peter, that the disjunction is true, or (ii) try and argue that there is a chance, for Peter, that the disjunction is true. Neither option seems palatable. On the one hand, the traditional invariantist has to deny the compelling judgment that there is at least a chance, for Peter, that the

disjunction is true. Indeed, it seems that there is a fairly *high* chance, for Peter, that the disjunction is true.²⁴ On the other hand, she must admit, of each disjunct, that there's no chance, for Peter, that it's true, but concede that there is nevertheless a chance, for Peter, that the disjunction is true. However, speakers seem to involve themselves in contradiction if they claim, of each of a series of propositions, 'There's no chance that's true', but then claim, with respect to their disjunction, 'There's a chance *that's* true'.²⁵

The endorsement of (CK) seems to entrain significant problems for the traditional invariantist.²⁶ However, if she does *not* endorse (CK), the concern arises that she lacks an explanation for the infelicity of utterances of the form 'I know that p, but there's a chance that not-p'. Once again, the traditional invariantist faces a dilemma: acceptance *or* rejection of (CK) leads to trouble.

Can the psychological proposals considered thus far assist the traditional invariantist? It is far from clear how to appeal to the overestimation account to assist the traditional invariantist here – either in attenuating the problematic consequences of endorsing (CK), or in explaining our judgments concerning the felicity of utterances of the form 'I know p, but there's a chance that not-p' without appeal to that principle.

However, it may seem that the Salience Account can provide some assistance to the traditional invariantist – at least, on the assumption that (CK) is false, and that knowledge is therefore compatible with a chance of error. On that assumption, the traditional invariantist is able to avoid the problems highlighted above with endorsing (CK). For instance, she can straightforwardly respect our judgment to the effect that there is a chance, for DeRose (in HIGH and in LOW) that the bank will be open, by

²⁴ Note that if the propositions concern the future, there may well be quite a high (or even a very high, if the disjunction is sufficiently large) objective chance that the disjunction is true. This alone may seem incompatible with claiming that there is a zero (epistemic) chance, for Peter, that it's true. See Hawthorne (2004a, 2005a) for relevant discussion.

²⁵ Of course, it might seem true to claim that there is a chance, for Peter, that the disjunction is true, but no chance, for Peter, that the disjuncts are true, if Peter doesn't realise that the disjuncts entail the disjunction. But this is irrelevant to the present argument. Grant that there is no chance, for Peter, that the disjunction does not follow from the disjuncts, and it still seems that there is at least a chance, for Peter, that the disjunction is true. See Chapter 3, §2 for relevant discussion.

²⁶ For further problems associated with a traditional invariantist endorsement of (CK), see Chapter 3, §1-3 and §5.

claiming that there is such a chance. However, the traditional invariantist then needs to explain away our judgments to the effect that utterances like ‘I know that the bank will be open, but there’s a chance that it won’t be’ by DeRose in LOW or HIGH seem infelicitous.

Notice that in claiming that there is a chance, for himself, that the bank won’t be open, DeRose plausibly makes salient that there is a possibility in which he has (or at least appears to have) just the same evidence, but in which the bank will not be open. Perhaps we give undue weight, in judging whether or not DeRose knows that the bank will be open, to the fact that there is such a possibility. It could then be proposed that the reason an utterance of ‘I know that the bank will be open, but there’s a chance that it won’t be’ seems infelicitous is grounded in the fact that asserting the second conjunct leads us to judge that the first conjunct is false.

One central problem with this proposal is already familiar. Recall that the traditional invariantist proposes that only a modest epistemic position, like that occupied by DeRose in the Bank Case, is required for knowledge. The traditional invariantist therefore allows that knowledge is compatible with *many* error possibilities in which DeRose has (or appears to have) just the same evidence, but in which the relevant proposition is false, and also with a *comparatively sizeable* chance that one of those possibilities is realised. But given that, it’s not clear why making salient merely that there is *a* chance that such a possibility will be realised would make much difference to our judgments about knowledge. That there is such a chance doesn’t come close to threatening our possession of knowledge. In a related vein, it’s a feature of even ‘paradigm cases’ of knowledge that there exist possibilities in which we have (or at least appear to have) just the same evidence, but in which the relevant proposition is false. It’s therefore not clear that the account can explain why utterances of the form ‘I know that p, but there’s a chance that not-p’ seem infelicitous, even though they are oftentimes true.

It therefore seems that whether the traditional invariantist accepts *or* rejects principle (CK), she faces a significant explanatory challenge. Moreover, it is far from clear how to make effective use of the psychological mechanisms considered thus far in meeting that challenge. The concern therefore arises that not only in regard to utterances such

as those of the form ‘I know that p, but I need to inquire further’, but also in regard to utterances such as those of the form ‘I know that p, but there’s a chance that not-p’, the traditional invariantist is forced to postulate significant speaker error.²⁷

§5. The Sophisticated Belief Account.

In a recent series of papers, Jennifer Nagel (2008, 2010a, 2010b, 2011) has proposed some more sophisticated alternatives to the proposals introduced in §2 and §3. These alternatives seem to offer promise in alleviating the problems for traditional invariantism encountered in regard to the HIGH-LOW contrast, and also in regard to our judgments concerning the relationships between knowledge and inquiry, and knowledge and chance.

Nagel suggests that individuals feel varying levels of *epistemic anxiety*.²⁸ The level of epistemic anxiety a subject feels determines (at least partially) how much inquiry she will engage in before taking a question to be settled. Nagel suggests that it is a feature of human cognition that in high stakes scenarios we feel more epistemic anxiety than in low stakes scenarios. This fits neatly with the observation that we will typically seek more evidence before taking a question to be settled in high stakes scenarios than we do in low stakes scenarios. Nagel proposes that we can explain the HIGH-LOW contrast, and other *prima facie* problematic judgments besides, by appeal to (i) the idea that taking a question to be settled – or ‘outright belief’ – is a necessary condition on knowledge and (ii) some features of the relationship between epistemic anxiety and outright belief.

If no stipulation regarding outright belief is made in the case descriptions, Nagel suggests that the explanation for the HIGH-LOW contrast may lie in the fact that, in the HIGH case, the subject appears to lack outright belief. For example, given the high stakes, those confronted with the HIGH context of the Bank Case will assume that DeRose feels a high level of epistemic anxiety. It might therefore be proposed that

²⁷ I provide a more thorough treatment of the relationships between knowledge and epistemic chance, and the standing of traditional invariantism with respect to utterances such as those of the form ‘I know that p, but there’s a chance that not-p’, in Chapter 3.

²⁸ See esp. Nagel (2010b).

they will assume that DeRose does not take the question of whether the bank will be open on Saturday to be settled on the fairly modest amount of evidence available to him. On the other hand, those confronted with the LOW context will assume that DeRose feels a more typical (lower) level of epistemic anxiety. They will therefore assume that DeRose *does* take the question to be settled on the modest amount of evidence available to him. This contrast could then account for our contrasting judgments about knowledge. Furthermore, Nagel suggests that even in those HIGH contexts in which the subject is simply entertaining possibilities of error (but the stakes remain low), the salience of error possibilities – bringing about as it does greater reflection on the possible pitfalls of the available evidence – may plausibly be assumed to raise the subject’s level of epistemic anxiety.²⁹

An immediate challenge to the account is what happens if we simply stipulate that the subject described in the HIGH case takes the relevant question to be settled on the moderate amount of evidence she possesses. It seems that our judgment about the cases is just the same. For example, in the HIGH context of the Bank Case, it seems that DeRose would truthfully deny himself knowledge. But isn’t the traditional invariantist forced to admit that the subject knows in such a case? After all, the subject outright believes the (true) proposition on the basis of a moderate amount of evidence.

In response to this kind of case, Nagel first suggests that those confronted with the HIGH case will assume the subject is experiencing the high level of epistemic anxiety appropriate to her stakes. She then proposes that given the subject forms an outright belief *despite* feeling such anxiety, those judging the case will naturally assume that there must be some other factor *overshadowing* the subject’s high epistemic anxiety, such as a time pressure, or distracting influence. Nagel proposes that such factors compromise accuracy. In fact, she suggests, the apparent presence of such factors could make the subject in the HIGH scenario appear *less* accurate than the subject in the LOW scenario. This could then account for why it seems correct, for example, for DeRose to deny that he knows the bank will be open in the HIGH context, even when it is stipulated that he forms an outright belief. (And given DeRose appears less accurate

²⁹ See Nagel (2010a; 2010b, p. 430, n. 13)

in the HIGH scenario than he does in the LOW scenario, this does not undermine the traditional invariantist account of our judgments concerning LOW.)³⁰

An initial concern. It is natural to wonder whether the subject would really appear less accurate in the HIGH case (in which it is also stipulated that the subject forms an outright belief) than in the LOW case. After all, it is stipulated in both the LOW and HIGH contexts that the central piece of evidence possessed by the subject is the same. In the Bank Case, it is that the bank was open two weeks ago on a Saturday. But given this stipulation, is the HIGH subject really going to appear less accurate than the LOW subject?

In response to such concerns, Nagel suggests that it may be in regard to the *processing* of the evidence that the HIGH subject seems less accurate than the LOW subject. Her suggestion is that since the HIGH subject will be assumed to be under something like a time pressure, her processing of the evidence will appear to be more susceptible to error.

However, even granting that the HIGH subject would appear to us to be under, for example, a significant time pressure, it's not clear that her processing would appear less accurate than that of the LOW subject. Note that the conditions we are dealing with are, in the HIGH case, time pressure and high epistemic anxiety, and in the LOW case, no time pressure and low epistemic anxiety. It may be reasonable to expect a subject's processing to appear less accurate in conditions in which epistemic anxiety is low and the subject is under time pressure than in conditions in which epistemic anxiety is low and there is no such pressure. But it's not clear that we should expect processing to appear less accurate in conditions in which there is a significant time pressure but there is also a factor motivating the subject to be more careful (the high stakes and consequent high epistemic anxiety), than in conditions in which anxiety is low and there is no time pressure. Indeed, for the typical subject, their processing may be just as careless (or just as accurate) in the two scenarios. This may well represent a cause for concern with Nagel's account. But for the present, let's just set it aside.

³⁰ See Nagel (2010a; and esp. 2010b, p. 419)

The account does have a number of positive features not shared by the accounts previously considered. First, it seems to offer some promise in accounting for the infelicitousness of utterances such as ‘I know that the bank will be open, but I need inquire further’ by DeRose in HIGH. Presumably, a normal subject will only take a question to be settled once her epistemic position is such that she no longer needs to engage in further inquiry. In addition, if a subject *does* take a question to be settled whilst still needing to engage in further inquiry, that might suggest that there is some accuracy-compromising factor, such as a time pressure, overshadowing her epistemic anxiety. It can therefore be proposed that either through the apparent lack of outright belief, or through apparent inaccuracy, a subject who needs to engage in further inquiry will not appear to possess knowledge. This could account for the infelicitousness of utterances such as those of the form ‘I know that p, but I need to inquire further’.

This account may also provide an explanation for the infelicity of utterances of the form ‘I know that p, but there’s a chance that not-p’. Nagel suggests that, typically, when it is salient to a subject that there is a chance, for her, that not-p, her processing will be such that she does not take the question of whether p to be settled.³¹ The subject will at most form an outright belief concerning the chance that p is true. For example, if it is salient to DeRose that there is a chance, for him, that the bank won’t be open, then – at least typically – he will not form an outright belief that it will be. In addition, if the subject *does* form an outright belief in the proposition at issue, that suggests she is influenced by some accuracy-compromising factor, like a time pressure. It can therefore be proposed that either through apparent lack of outright belief, or apparent inaccuracy, if it is salient to the subject that there is a chance that not-p, she will appear not to know that p. This could ground our judgment to the effect that utterances of the form ‘I know that p, but there’s a chance that not-p’ seem infelicitous.³²

Nagel’s view therefore promises to explain – without the attribution of significant speaker error – the HIGH-LOW contrast, and our judgments concerning the

³¹ See Nagel (2011).

³² I return to discussion of the relationship between knowledge and chance on Nagel’s account in §5.4.

relationships between knowledge and chance, and knowledge and inquiry. Unfortunately for the traditional invariantist, the view suffers from some serious difficulties. In the next section, I present a number of them. In the following sections, I consider some possible replies.

§5.1 Objections to the Sophisticated Belief Account.

The most telling problems for the Sophisticated Belief Account involve third person cases, cases featuring subjects with atypical psychologies, and cases featuring various counterfactual judgments. Let's take the latter first.

- (2) He doesn't know the bank will be open. But if the stakes were lower, he would know.

Suppose that DeRose is in the situation described in the HIGH context of the Bank Case, but let it be stipulated that DeRose outright believes that the bank will be open, on the basis of having been there two weeks ago on a Saturday. Suppose that another subject, Jeff, is also in a situation like that of DeRose in the HIGH. Jeff has a lot at stake, and he is also considering the possibility that the bank has changed its hours. Jeff was also at the bank two weeks ago on a Saturday. However, unlike DeRose, Jeff has also just gone into the bank and confirmed the hours.

In such a case, it seems true and appropriate for Jeff to deny that DeRose knows that the bank will be open, but true and appropriate for him to claim that he (Jeff) knows that the bank will be open. Now suppose that Jeff is aware of DeRose's predicament in HIGH, and is also aware that if the stakes were lower for DeRose, then DeRose would have no more evidence than he (DeRose) currently does in regard to the proposition that the bank will be open, and that he (DeRose) would still take the question of whether the bank will be open on Saturday to be settled on the basis of that evidence. Suppose Jeff utters (2).

Such an utterance seems false. However, if facts about outright belief drive our negative judgments about knowledge, it seems that we should judge that utterance to be true. Given his current stakes, we will assume that DeRose has high epistemic

anxiety. In consequence, the fact that he has formed an outright belief, he will appear to be under the influence of some accuracy compromising factor, like a time pressure. But if the stakes were lower, we would presumably take it that DeRose's epistemic anxiety would be lower. It should therefore seem to us that he is able to take the question of whether the bank will be open to be settled on the basis of having been there two weeks ago on a Saturday without appearing to be under the influence of some accuracy compromising factor, like a time pressure. Jeff's utterance of (2) should therefore seem true.

A second problem. Suppose that DeRose and his wife are in the situation described in HIGH, but are also aware of some third party, Jim. Jim has just the same evidence as DeRose to suggest that the bank will be open on Saturday, but Jim has little at stake in the matter. It is typically reported that it seems correct for DeRose, in HIGH, to claim that Jim does not know the bank will be open.

However, Jim does not share DeRose's high epistemic anxiety, so Jim will presumably take the question of the bank's being open tomorrow to be settled on the moderate evidence available. Further, Jim need not appear to be subject to time pressure, distraction, or such like. In short, he's just like DeRose in LOW. So why do we judge DeRose's knowledge denial to be correct?

A third problem. Imagine an individual, John, with a slightly quirky psychology. John is like a typical subject in every respect, except in the respect that he is so constituted that he doesn't feel more epistemic anxiety when the stakes are high than when they are low. He's just not that epistemically anxious. John therefore readily forms outright beliefs on moderate evidence, even when the stakes are high.

Now suppose John is in a situation similar to that of DeRose in HIGH. John takes the question of whether the bank will be open on Saturday to be settled on the basis of the fact that he was at the bank two weeks ago on a Saturday, but John has a lot at stake in the question of whether the bank will be open, and he is also considering the possibility that the bank has changed its hours. Suppose John utters 'I don't know that the bank will be open tomorrow'. Such an utterance seems true. However, on Nagel's account, John's utterance should presumably seem false. The reason is that (i) John

has an outright belief on evidence that, by traditional invariantist lights, is sufficient for knowledge, and (ii) John does not feel heightened epistemic anxiety in high stakes scenarios – so we need not suppose that there are any accuracy compromising factors, like a time pressure, acting to overshadow high anxiety. It should therefore seem to us that John knows that the bank will be open.

These three problems strike at the core of Nagel's proposal. Notice also that the last two problems suggest that Nagel's account cannot capture simple variations on the HIGH-LOW contrast. Recall the third person case, described above. In that case, it seemed correct for DeRose, who has a lot at stake, and is considering the possibility that the bank has changed its hours, to claim that Jim does not know that the bank will be open. This just represents a typical HIGH case. In the corresponding LOW case, the situation would be identical except in the respect that DeRose would not have a lot at stake and would not be considering such distant possibilities of error. In such a situation, it seems correct for DeRose to claim that Jim knows that the bank will be open. Our argument therefore suggests that Nagel's account cannot capture this variation on the HIGH-LOW contrast, since her account seems unable to capture our judgment regarding the HIGH case.

The case described featuring the individual with the atypical psychology, John, also represents a HIGH case. It seems true for John, who has a lot at stake in whether or not the bank will be open, and is considering the possibility that the bank has changed its hours, to claim that he does not know that the bank will be open. In the corresponding LOW case, the situation would be identical except in the respect that John would have little at stake, and would not be considering the possibility that the bank has changed its hours. It seems that, in such a case, it would be true and appropriate for John to claim that he knows that the bank will be open. (Remember: his psychological quirk is just that in high stakes circumstances, he doesn't feel heightened epistemic anxiety. In low stakes (as well as high stakes) circumstances, he feels a fairly low level of epistemic anxiety, and therefore engages in epistemic behaviour typical to low stakes subjects.) Once again, it seems that Nagel's account is unable to capture our judgment concerning the relevant HIGH case.

The concern thus arises that Nagel's account is unable to capture simple variations on the HIGH-LOW contrast, and that it suffers from further problems besides – those featuring certain counterfactual judgments. However, Nagel has a response. Her response is to suggest that these *prima facie* problematic judgments arise because high stakes subjects are prone to project onto low stakes subjects the (high) levels of epistemic anxiety that they currently feel. This appeal to so-called 'epistemic egocentrism' forms a central component of her psychological defence of invariantism.

§5.2 Responding to Objections I: Epistemic Egocentrism.

Research has shown that we tend to treat others as though they shared features of our mental life – our values, knowledge, and so on.³³ This is something we do when we lack specific information about the other individuals whose mental lives we are concerned with. There are some arguments to suggest that, in such circumstances, this is not an irrational thing to do.³⁴ However, there is also evidence to suggest that even when we know some other individual does not share some feature of our mental life, we continue to treat her as if she did. Or at least that even when we know some other individual not to share some feature of our mental life, we allow our own possession of that feature to colour our evaluation of her. One of the most prominent examples of this phenomenon is the so-called 'hindsight bias'.³⁵

The hindsight bias provides an example of how possession of certain outcome information colours our assessment of those that we know to lack that information. A typical experimental set-up is as follows. Participants are provided with the information originally available to those observing a historical conflict between two tribes, A and B. Participants are then asked 'How likely would you have judged it to be that tribe A would triumph in that conflict?'. Prior to asking the question, one group of participants is also told that tribe A in fact triumphed, whilst the remaining participants are not given any outcome information. In more general terms: participants are asked how likely they would have judged a certain outcome, *o*, to be

³³ See e.g. Birch and Bloom (2004), Nickerson (1999), and Royzman *et al.* (2003).

³⁴ See Dawes (1989).

³⁵ See Fischhoff (1975). There is now an enormous literature on the hindsight bias. See e.g. Blank, Musch, and Pohl's (2007) introduction, and the subsequent papers, in the 2007 special issue of *Social Cognition* on the hindsight bias.

given some supplied evidence. One group of participants is also told that *o* was the actual outcome, whilst another group is not given any outcome information.

It has repeatedly been found that members of the group provided with the information that the outcome was *o* give significantly higher estimates of how likely they would have judged *o* to be (on the original evidence available) than do members of the group not provided with such information. In our example set-up, this suggests that those told that tribe A triumphed would provide an answer to the question ‘How likely would you have judged it to be that tribe A triumphed in that conflict?’ that is higher than those not provided with any outcome information. The upshot of such hindsight studies seems to be that being acquainted with outcome information colours our evaluation of the epistemic states of our (in this case imagined) past selves. In some sense, we project our knowledge of the outcome onto our imagined past selves, despite knowing that our imagined past selves do not share that knowledge.

Nagel suggests that in examples like those given in the previous section, it is plausible that the attributor is projecting some feature of his/her cognitive situation onto the subject (even though he/she knows that the subject doesn’t share that cognitive situation). One natural suggestion is that the attributor is projecting his/her heightened epistemic anxiety onto the subject. I will assume this in the discussion to follow, but the relevant criticisms could readily be adapted if it is alleged that what is projected is instead something like heightened stakes or increased awareness of error possibilities.

Why might attributors project their levels of epistemic anxiety onto others? A possible explanation is that some process of ‘anchor and adjustment’ is at work.³⁶ The proposal would be that attributors treat their current level of epistemic anxiety as an anchor (or starting point), and make adjustments from that anchor based on the known differences between the subject and themselves. It is a common feature of such anchor and adjustment that we under-adjust. This suggests that attributors in HIGH contexts will treat subjects in LOW contexts as though they have a higher level of epistemic anxiety than they really do. This ‘anchor and adjustment’ proposal also has

³⁶ For a review of several potential ‘cognitive process model’ explanations of the hindsight bias, some of which involve something akin to the anchoring idea, see Blank & Nestler (2007).

application in the hindsight case. The proposal would be that study participants use their current epistemic situation (which includes the information that the outcome was *o*) as an anchor, and then adjust by trying to remove the effects of the outcome information on their current epistemic situation. Once again, they under adjust, and so end up providing a significantly higher estimate of how likely they would have judged the outcome *o* to be.³⁷

Let's see how the appeal to epistemic egocentrism is supposed to assist with the problematic third-person case introduced in the previous section. In that example, DeRose is in a situation resembling HIGH, and is aware of Jim and Jim's situation, which resembles that of DeRose in LOW. It is typically reported that it seems correct for DeRose to deny that Jim knows the bank will be open.

The egocentric explanation for why *DeRose* might take Jim to lack knowledge is that, due to the fact that DeRose exhibits high levels of epistemic anxiety, he will treat Jim as having a higher level of epistemic anxiety than he (Jim) really does. In consequence, Jim will either appear not to have an outright belief, or, if he does so appear, DeRose will naturally take it that Jim only has such a belief because he's under the influence of some accuracy compromising factor, like a time pressure. Either way, it will seem to DeRose that Jim does not know.

But what about *us*? DeRose's utterance seems correct not merely from the perspective of DeRose, but also from the perspective of those people (like you and me) assessing it. It does not matter to the readers of the case whether the bank will be open or not. In consequence, it might seem implausible to suggest that those readers would feel a high level of epistemic anxiety with respect to the proposition that the bank will be open. However, Nagel suggests that, in fact, considering high stakes subjects, and reflecting on error possibilities, may lead those reading the cases to feel a heightened level epistemic anxiety with respect to the proposition that the bank will be open.³⁸ As a result of the egocentric bias, it can therefore be proposed that we also end up

³⁷ There are *many* proposals for why we see effects such as the hindsight bias. (See, e.g., the papers in the 2007 special issue of *Social Cognition* on the hindsight bias.) Not all of these proposals are so readily extendable to the kind of phenomenon Nagel is positing.

³⁸ See e.g. Nagel (2010b, p. 413)

treating Jim as though he had a higher level of epistemic anxiety than he really does. This could then explain our (mistaken) judgment to the effect that DeRose is correct to deny the Jim knows that the bank will be open.

There are a variety of problems with this appeal to epistemic egocentrism. One rather serious problem stems from the fact that it is not particularly plausible to claim that we treat others as sharing the *same* level of epistemic anxiety that we currently exhibit, even when we are aware that their stakes are different. There are a number of sources of evidence for this claim.

First, just note that counterfactuals such as ‘If there were less at stake, he would not engage in as much inquiry before he took the question to be settled’ seem obviously true. However, if we treated other subjects as though they shared the same levels of epistemic anxiety as us, even in situations in which we knew their stakes to be different, it should seem to us that – no matter what their stakes – others would need to check just as much as we do in order to take a question to be settled. Such counterfactuals claims should therefore seem false to us.

Notice that this observation regarding such counterfactual claims seems, by itself, to undermine the egocentric proposal as a solution to the counterfactual-based objection of the previous section. Recall that Jeff’s utterance, in reference to DeRose in HIGH, of ‘He doesn’t know. But if the stakes were lower, he would know’ seems false. The egocentric explanation in this case will rest on the claims that (i) Jeff (and those reading the case) exhibit high epistemic anxiety and (ii) Jeff (and those reading the case) treat the counterfactually envisaged low-stakes-DeRose as though he also possessed their high epistemic anxiety. However, if the correct explanation did rest on those claims, then presumably an utterance by Jeff, in reference to DeRose in HIGH, of ‘He doesn’t take the question to be settled. But if the stakes were lower, he would’ would *also* seem false. But it seems true.

Perhaps more importantly, the claim that we treat others as though they exhibit the same level of epistemic anxiety that we do is just not supported by the relevant empirical research. That research suggests that whilst we may treat another subject as just like us *when we have no information about whether they are like us or not*, we do

not treat other subjects as just like us *when we have information that suggests they are different*.³⁹ If we have such information, we may be liable to treat them as *more* like us than they really are, but not as *exactly* like us. For example, in the hindsight case, those study participants given the outcome information don't treat their (imagined) former selves as though they had just the same outcome information that they do. If they did, they'd presumably claim that their imagined former selves would give *very* high estimates of the likelihood of the relevant outcome – ones that simply assumed their former selves had the outcome information. Similarly, in the cases at hand, it may be plausible that high stakes subjects treat those they are aware to be in lower stakes scenarios as exhibiting *higher* levels of epistemic anxiety than they really do, but not that they treat those in low stakes scenarios as exhibiting *just the same* level of epistemic anxiety as they (the high stakes individuals) do.

The problem is that the latter claim is just what is required to account for all our judgments. Take the third-person case introduced in §5.1 once more. Recall that the egocentric explanation of why (high stakes) DeRose would judge that (low stakes) Jim lacks knowledge is based around the idea that DeRose takes Jim to exhibit a higher level of epistemic anxiety than he really does. However, DeRose is aware that the stakes are low for Jim. It is therefore plausible that, in this case, the relevant egocentric effect will only be *partial*, and DeRose will not take Jim to exhibit the same level of epistemic anxiety that he (DeRose) does, but rather a lower level of epistemic anxiety.⁴⁰ DeRose will therefore take Jim to require less evidence than he (DeRose) does in order to form an outright belief (without appearing under the influence of accuracy compromisers, like time pressure). In consequence, DeRose should presumably take Jim to require less evidence than he does in order to know that the bank will be open. Similarly, even supposing that *we* exhibit the same level of epistemic anxiety as DeRose does with respect to the proposition that the bank will be open, we will not, on the assumption that the egocentric effect is only partial, treat

³⁹ For relevant papers, see e.g. Nickerson (1999), and much of the literature on the hindsight bias and other egocentric phenomena (such as the false consensus effect).

⁴⁰ Notice that even supposing DeRose lacked specific information about Jim's stakes, we have reason to suspect that he would *still* not treat Jim as exhibiting just the same level of epistemic anxiety that he does. The reason is that DeRose is plausibly in an *unusual* stakes situation. Typical people are not in such a situation. Thus, DeRose automatically (as it were) has some evidence that Jim is not in such a situation, but rather has lower stakes than him. See Nickerson (1999) for relevant discussion.

Jim as though he shared that level of epistemic anxiety. Thus, it seems that we should also take Jim to require less evidence than DeRose requires in order to know that the bank will be open.

The upshot is that there should be cases in which Jim and DeRose occupy identical epistemic positions (and both have an outright, true belief that the bank will be open), but in which it seems true for DeRose to utter ‘Jim knows that the bank will be open, but I don’t’. However, such conjunctive claims seem inappropriate and false. In order to avoid predicting that such conjunctive claims should seem true, the view requires the implausible claim that attributors (and readers of the case) project onto others exactly the level of epistemic anxiety they feel.⁴¹

There are further problems with the egocentric proposal, even setting aside worries that the effect is only partial. First, there is evidence that egocentric effects such as the hindsight bias vary with intelligence. Those who are more intelligent exhibit less bias.⁴² Applied to the case at hand, this suggests that those who are more intelligent should find the utterances in the cases presented in the previous section less problematic – they should be better at avoiding egocentrism with respect to their levels of epistemic anxiety. Professional academics (hopefully) qualify as highly intelligent individuals. But it is they who have found utterances like ‘He doesn’t know. But if less were at stake, he would know’ so objectionable. Note that this research also re-inforces the idea that the relevant effects are only partial – the

⁴¹ It might be replied that it is possible to explain at least the appearance of *inappropriateness* in an utterance by DeRose of ‘Jim knows that the bank will be open, but I don’t’ by appeal to the idea that knowledge is required for appropriate assertion. On the assumption that DeRose doesn’t know that the bank will be open, it follows that he doesn’t know that Jim knows that the bank will be open (assuming knowledge is factive). On the assumption that knowledge is required for appropriate assertion, it therefore follows that DeRose’s utterance of the second conjunct of ‘Jim knows that that bank will be open, but I don’t’ entails that his utterance of the first conjunct is inappropriate. This could account for the inappropriateness of uttering such a conjunction.

However, given the egocentric effect is only partial, the account also predicts (for similar reasons) that there should be cases in which an utterance by some third party of ‘Jim knows that the bank will be open, but DeRose does not’ seems true and appropriate, despite the fact that Jim and DeRose occupy equivalent epistemic positions (and both truly believe that the bank will be open). If the third party is suitably situated (see e.g. DeRose, 2009, p. 196-197 for relevant details), the assertion based explanation cannot be extended to account for our judgment that such an utterance seems, in fact, inappropriate (and false).

⁴² See Stanovich and West (1998). See also Musch and Wagner (2007).

findings were that the hindsight effect is partial, and less pronounced in individuals of high intelligence.

Second, there are reasons to expect that there might be another egocentric effect working in an opposing direction. As we saw in the case of the hindsight bias, those who possess certain privileged pieces of information are liable to allow that information to skew their assessment of the epistemic positions of those who lack it. But readers of the cases are given the outcome information that the bank is open on Saturday. In consequence, it seems that readers of the cases might take the subject's evidence in the Bank Case (both in LOW and in HIGH) to support the proposition that the bank will be open to a greater extent than it really does. It follows that even supposing that we – in certain circumstances – take individuals with low stakes to exhibit higher levels of epistemic anxiety than they really do, we may *also* take them to have better evidence than they really do. The concern is that one effect may just cancel the other out.

Nagel's response to this concern is to argue that it's just not plausible that we would project our 'outcome information' in regard to the bank's being open when evaluating whether DeRose is correct to claim, in LOW, that he knows that the bank will be open, or when evaluating whether DeRose is correct to deny, in HIGH, that he has such knowledge.⁴³ She proposes that we assess whether a knowledge ascription (or denial) is correct by representing the transition from the subject's evidence to their state of belief, and then judging whether that transition from evidence to belief produces knowledge. Nagel claims that whatever egocentric effects might be in play, it simply isn't plausible to suggest that, in the course of such a process, we would take DeRose's evidence to include the outcome information itself – viz. the proposition that the bank will be open.

This response is problematic on at least two counts. First, it also seems implausible to suggest that, when representing the cognition of subjects we know to be in low stakes situations, we treat those subjects as exhibiting exactly the same levels of epistemic anxiety as high stakes subjects. But – as we have seen – that's exactly the claim that

⁴³ Nagel (2010b, p. 431, n. 18) credits the objection to John MacFarlane.

Nagel requires. Second, research on, for example, the hindsight bias shows that the effect of outcome information is rather more insidious than Nagel's reply might suggest. In the case of the hindsight bias, it's not that we directly take our imagined former self to have just the same outcome information that we do, but rather that our possession of the outcome information makes the original evidence seem as though it provides better support for the relevant outcome than it really does. In the Bank Case then, we have reason to suspect that outcome information will make DeRose's evidence appear better than it really is. However, it is not implausible to suggest that our outcome information regarding the bank's being open could influence our assessment of whether the transition from evidence to belief produces knowledge by making the evidence appear to support the proposition that the bank will be open to a greater extent than it really does. The basic concern thus stays in force. One egocentric effect, which leads to overamplification of the quality of the subject's evidence, may counterbalance the other proposed egocentric effect, which leads to overestimation of the subject's level of epistemic anxiety.⁴⁴

A final concern. It has been suggested that our failure to correct for certain egocentric effects, even when provided with information about the mental state of the individual at issue, may have (at least in part) a self-presentational basis. For example, it has been suggested that it is (at least in part) because we want to look like better predictors that we exhibit hindsight bias.⁴⁵ This has a flip-side: it suggests that in cases where exhibiting hindsight bias makes us look worse, we will be less inclined to exhibit it.

The problem for the egocentric proposal is that there are certain knowledge judgments – so-called 'retraction judgments' – which seem to involve negative evaluation of our

⁴⁴ As noted earlier – in the discussion of the Salience Account – one way to reduce the biasing effect of outcome information is by thinking carefully about alternative outcomes. This suggests that the egocentric effect that leads to overamplification of the evidence may be less pronounced in our judgments about the HIGH cases, in which error possibilities are raised. This, in turn, suggests that the effect associated with outcome information may not cancel out the effect associated with epistemic anxiety after all. However, recall that research suggests that subjects need to think quite carefully about alternatives, or to generate them themselves, in order to attenuate the biasing effects of outcome information – it's not enough to simply entertain alternatives. This suggests that the biasing effect of outcome information is probably not going to be significantly reduced for those reading the HIGH case.

⁴⁵ See, e.g., Campbell and Tesser (1983).

past selves. In high stakes situations, or situations in which more distant or unlikely error possibilities have been raised, we are liable to claim that our earlier knowledge attributions were false. However, given that our past selves will often have been in low stakes situations (in which error possibilities are not being considered), their levels of epistemic anxiety will have been correspondingly low. The appeal to epistemic egocentrism will therefore be required in order to account for our claiming that our earlier knowledge ascriptions were in error. But note that by projecting high epistemic anxiety onto our former selves, we end up painting our former selves in a negative light: we end up claiming that the knowledge attributions we made were in error (when they weren't). If self-presentational effects play an important role in egocentric biasing, that does not sit comfortably with appeal to such biasing in response to retraction phenomena. But it is not clear what alternative explanation of that phenomena might be available.

For a number of reasons then – perhaps chief among them the concern that the relevant egocentric effect is only partial – Nagel's appeal to epistemic egocentrism does not look like an adequate reply to the problem posed by the cases presented in §5.1.

§5.3 Responding to Objections II: Alternatives to Egocentrism.

Nagel does suggest an alternative for responding to our original objections (§5.1) that does not involve an appeal to egocentric bias. Recall the third person case. In that example, DeRose is in a situation resembling HIGH, and is aware of Jim and Jim's situation, which resembles that of DeRose in LOW. It seems correct for DeRose to deny that Jim knows the bank will be open.

Nagel suggests that in such a case, DeRose is concerned with the possibility of using Jim as an informant, and – as a result – is really considering the question of whether Jim would know that the bank will be open, were he to go and ask him.⁴⁶ However, if DeRose were to explain his predicament to Jim, and ask him whether he knows that

⁴⁶ For further details, see Nagel (2010b, p. 425-6). For a similar proposal (but put to work in defence of subject sensitive invariantism), see Stanley (2005a, p. 101-104). I discuss Stanley's proposal in Chapter 4.

the bank will be open, Jim's epistemic anxiety would be raised. This might then form the basis of an explanation for why it seems correct for DeRose to deny that Jim knows the bank will be open.

The proposal is limited at best. First, it seems that DeRose asking Jim whether the bank will be open need not have much effect on Jim's epistemic anxiety. It is entirely possible to treat people as informants without making them aware of how much is at stake for us, and thus without making it plausible that they will experience significant increases in epistemic anxiety. (This point is obscured somewhat by Nagel choosing a case involving the police – the police asking one a question may instantly raise one's epistemic anxiety.) It could therefore just be stipulated in the case description that were DeRose to approach Jim, he would not explain his situation to Jim. Such a stipulation does not seem to influence the relevant judgments – it still seems true and appropriate for DeRose to deny that Jim knows that the bank will be open.

Second, it should presumably be possible for DeRose to *clarify* his knowledge claim in such a way as to avoid potential confusions associated with his interest in the question of whether Jim would know that the bank will be open, were he to treat him as an informant. For example, it should presumably seem true and appropriate for DeRose to utter, 'Jim might know *right now* that the bank will be open tomorrow, but were I to go and ask him, he would not know that'.⁴⁷ However, such claims do not seem true and appropriate.

Finally, this account does not look very promising when it comes to accounting for cases featuring individuals with atypical psychologies. Recall John, who does not feel heightened epistemic anxiety when the stakes are high. Suppose that John is in a situation like that of DeRose in LOW. John was at the bank two weeks ago on a

⁴⁷ The 'might' in the first conjunct is inserted to fend off concerns (see fn. 41) tied to the idea that knowledge is required for appropriate assertion. Those concerns can also be avoided by supposing that DeRose (but not Jim) has collected additional evidence, such that, despite DeRose's high stakes, it seems true and appropriate for him to self-ascribe knowledge that the bank will be open. In that case, it seems that DeRose should be able to clarify Jim's epistemic situation as: 'Jim knows right now that the bank will be open, but if I were to go and ask him, he wouldn't know'. However, such an utterance seems clearly false. Given DeRose's improved epistemic position, no appeal to the claim that knowledge is required for appropriate assertion looks likely to assist in explaining that judgment. See fn. 41 for relevant discussion.

Saturday, but he has little at stake in the matter of whether the bank will be open, and is not considering any distant or unlikely error possibilities. Suppose also that DeRose is in a situation like that described in HIGH – he has a lot at stake and is considering the possibility that the bank has changed its hours – and that he is aware of John and John’s situation.

In that case, it seems true and appropriate for DeRose to claim that John does not know that the bank will be open. However, even supposing that DeRose explained his high stakes situation to John, John is not going to feel a significant increase in his epistemic anxiety. In consequence, the claim that he *does* feel such anxiety cannot be invoked to explain our judgment about this case.

This alternative to an appeal to epistemic egocentrism, in responding to the objections of §5.1, also looks unsuccessful. I conclude that The Sophisticated Belief Account is therefore unable to explain pertinent manifestations of the HIGH-LOW contrast – those featuring third-person cases, and those featuring subjects with atypical psychologies – and seems to entrain problems of its own – those concerning utterances such as ‘DeRose does not know that the bank will be open. But if he had less at stake, he would know’.

§5.4 Inquiry, Chance, and Outright Belief.

Recall that Nagel’s account promised not just to explain the HIGH-LOW contrast, but also to account for our judgments concerning the relationships between knowledge and inquiry, and knowledge and chance. In this section, I show that simple extensions of the cases presented in §5.1 reveal that Nagel’s account is also unable to capture our judgments concerning those relationships.

Recall the proposed explanation for our judgment to the effect that utterances of the form ‘I know that p, but I need to inquire further’ seem infelicitous. The proposal centred on the idea that, if a subject needs to inquire further in regard to some proposition p, her levels of epistemic anxiety will be such that, typically, she will not take the question of whether p to be settled. In addition, if she *does* take that question to be settled, that suggests that she is under the influence of some accuracy

compromising factor, like a time pressure. It follows that either through apparent lack of outright belief, or through apparent inaccuracy, a subject that has to inquire further in regard to p will appear not to know p . This could account for infelicity of an utterance of the form ‘I know p , but I need to inquire further’.

Unfortunately, this proposal also encounters problems in regard to individuals with atypical psychologies. Suppose John is exactly like DeRose in HIGH except in the respect that John feels atypical levels of epistemic anxiety, such that he takes questions to be settled on moderate levels of evidence, even when the stakes are high. Suppose that in regard to his epistemic position, the fact that he is considering the possibility that the bank has changed its hours, and that fact that the stakes are high for him, John is just like DeRose in HIGH. The problem for the account just sketched is that just as it seems infelicitous for DeRose in HIGH to utter ‘I know that the bank will be open, but I need to inquire further’, it *also* seems infelicitous for John, in his context, to utter ‘I know that the bank will be open, but I need to inquire further’. However, given John is stipulated not to feel high levels of epistemic anxiety when the stakes are high, it seems plausible both that he *does* take the matter of the bank’s being open to be settled, and that he would do so despite *not* being under the influence of some accuracy compromiser, like a time pressure. John should therefore appear to know that the bank will be open. However, it seems that, given his stakes, John does need to inquire further. Thus, his utterance of ‘I know that the bank will be open, but I need to inquire further’ should presumably seem true.

Can the traditional invariantist respond to this problem? It is unclear how to appeal to the suggestion put forward in the previous section in accounting for this judgment. But the egocentric-bias explanation is fairly obvious. The explanation will begin from the claim that readers of the case feel a heightened sense of epistemic anxiety, as a result of reflecting on John’s high stakes. Let’s just suppose that we feel a level of epistemic anxiety appropriate to John’s stakes, in the sense that we exhibit the level that we would feel, were we actually in his situation. In virtue of the egocentric bias, it might then follow that even though we know that John does not feel the same high anxiety (it’s in the case description that he does not), we might plausibly take him to exhibit higher anxiety than he really does. However, it was argued in §5.2 that it’s plausible that the relevant egocentric effect is only *partial*. Thus, John’s level of

epistemic anxiety should seem to us to be lower than ours. It should therefore seem to us that John is able to form an outright belief (without his accuracy appearing compromised) on less evidence than we are able to, and, consequently, that he should be able to know that the bank will be open on less evidence than us.

However, despite the quirk in his psychology, the amount of evidence John requires in order to appropriately close inquiry, given his stakes, is presumably going to be something like the amount that would satisfy *our* levels of epistemic anxiety. It follows that even supposing that John does not, in the case depicted, appear to us both to know that the bank will be open, and to need to check further, it seems that there should be a way to improve his epistemic position (but not sufficiently to satisfy our levels of epistemic anxiety), such that it does. That is, there should be a way to improve his epistemic position such that an utterance, by John, of ‘I know, but I must inquire further’ seems true. However, a slight improvement in John’s epistemic position does not seem to render such an utterance felicitous – it still seems infelicitous. In consequence, it seems that, even allowing an appeal to epistemic egocentrism, Nagel’s account is unable to respect our judgments concerning the relationships between knowledge and inquiry.

What about our judgments concerning the relationship between knowledge and chance? Recall that on Nagel’s account, (CK) is false, and that knowledge is therefore compatible with a chance, for the subject, that she is in error. It was proposed earlier that the reason utterances of the form ‘I know that p, but there’s a chance that not-p’ seem infelicitous is because, in circumstances in which it is salient to the subject that there is a chance, for her, that not-p, she would not form an outright belief that p (unless she were in accuracy compromising conditions).

This proposal is also subject to problems based around individuals who are stipulated to have psychologies that are atypical. For example, suppose that due to feeling low levels of epistemic anxiety, John typically forms outright beliefs on moderate evidence, even when it is salient to him that there is a small chance that he’s in error. Now consider an utterance by John of ‘I know that the bank will be open, but there’s a chance that it won’t be’. Such an utterance seems just as problematic as it would were it uttered by DeRose (who we can stipulate to have a more normal psychology).

Once again, there is no need to suppose that John lacks outright belief, or that he is under the influence of some accuracy compromising factor, like a time pressure. His moderate evidence should therefore enable him to know that the bank will be open. But the account rejects (CK), and grants that it's true that there's a chance, for John, that the bank is not open. The utterance should therefore seem felicitous.

There are also problems with third-person cases. Suppose that DeRose is in the situation described in the LOW context of the Bank Case. He was at the bank two weeks ago on a Saturday, and he believes that the bank will be open as a result. DeRose also has little at stake, and is *not* reflecting upon the fact that there is a chance, for him, that the bank won't be open. Suppose that Jim is in a different context, and is aware of DeRose's situation. Notice that it seems infelicitous for Jim to utter 'DeRose knows that the bank will be open, but there's a chance, for him, that it won't be'. However, given that DeRose is just not reflecting upon the fact that there is a chance, for him, that the bank won't be open, there seems no reason to expect either that he will lack an outright belief that it will be, or that if he has such a belief, it will be as a result of his being under accuracy compromising conditions. Thus, Nagel's proposal cannot account for our judgment in this case.

Of course, the proponent of this view might once again suggest that egocentric effects (or the proposal centred around the need for testimony) have a role to play in explaining this judgment. But these appeals will fail for parallel reasons to those given above. However, there is an interesting alternative to the account for the infelicity of utterances of the form 'I know that p, but there's a chance that not-p' provided above. This account is grounded in the suggestion that when it is salient to an individual that there is a chance that not-p, it is not *rational* for them to form an outright belief that p. Some remarks of Nagel might seem to express sympathy with such a view:

[A]s long as I am thinking about lottery wins or traffic accidents, certain explicit judgments will be out of reach for me: given my evidence in the controlled mode of cognition [that mode of cognition activated by considering certain kinds of error possibilities], I *should* at most judge it to be very probable that I am meeting my friend for lunch tomorrow. (Nagel 2011, p. 35; my emphasis.)

If it is alleged that the potential irrationality here is incompatible with knowledge, we have an alternative explanation for the the chance knowledge conjunctions. The idea would be that in uttering ‘I know that p, but there’s a chance that not-p’ one makes salient that there is a chance that not-p, and so it becomes irrational for one to outright believe that p. Consequently, one does not know p.

This could presumably explain those judgments that concern the individual with the atypical psychology. If we grant that, despite the quirk in his psychology, it is not rational for John to form an outright belief that the bank will be open when it is salient to him that there’s a chance, for him, that he’s in error, it follows – on the present proposal – that he will lack knowledge that the bank will be open. This could account for the apparent infelicitousness of an utterance by John of ‘I know that the bank will be open, but there’s a chance that it won’t be’.

However, the proposal looks unable to address the third person cases – in the case given above, DeRose is just not considering the fact that there’s a chance, for him, that the bank won’t be open. The account therefore provides no reason to expect an utterance by Jim of ‘DeRose knows that the bank will be open, but there’s a chance, for him, that it won’t be’ to seem infelicitous.

There are further problems besides.⁴⁸ For example, it seems that various problematic counterfactuals will come out true on the proposal. Suppose Jeff was at the bank two weeks ago on a Saturday, and believes that the bank will be open on Saturday as a result. Suppose that Jeff has little at stake, but that he is currently considering the fact that there is a chance, for him, that the bank won’t be open. In that case, it should seem true and appropriate for Jeff to utter ‘I don’t know that the bank will be open. But if it weren’t salient to me that there’s a chance that the bank won’t be open, then I

⁴⁸ The reader will notice that the proposal is suspiciously like a form of anti-intellectualism (and therefore not a form traditional invariantism), since it is the *salience* of the fact there is a chance that not-p that is central in bringing about a lack of knowledge. However, it’s probably not very wise to engage in definitional disputes. It is enough to observe that the account is sufficiently like anti-intellectualist accounts that it will suffer from many of the same problems: for example, those featuring third-persons and problematic counterfactuals (given in the main text). See Chapter 4 for further problems with anti-intellectualism.

might know that it will be'.⁴⁹ However, such counterfactuals do not seem true and appropriate.

Such problems suggest that this alternative, more prescriptive proposal, is no more successful. In consequence, it appears that the traditional invariantist is unable to explain our judgments concerning the relationships between knowledge and inquiry, or knowledge and chance, by appeal to the Sophisticated Belief Account. It was shown in §4 and §4.1 that the traditional invariantist is also unable to account for those judgments by appeal to the psychological proposals considered earlier, or to some derivative sense of propriety. The concern therefore arises that, in regard to those judgments, the traditional invariantist is forced to attribute significant speaker error.

§6. Summary.

In this chapter, I have explored the possibility that the traditional invariantist can explain the HIGH-LOW contrast by appeal to features of our psychological processing. Proposals centred around the idea that we overestimate the likelihood of salient error possibilities obtaining, or take into account the existence of various error possibilities obtaining, were shown to be unsuccessful. A more sophisticated proposal was then considered, which sought to exploit the relationships between knowledge, epistemic anxiety, and outright belief in accounting for the HIGH-LOW contrast. This proposal also promised to assist in explaining some *prima facie* problematic judgments concerning the relationships between knowledge and inquiry, and knowledge and chance. However, it was shown that this proposal could not capture pertinent manifestations of the HIGH-LOW contrast, nor could it adequately capture our judgments concerning the relationships between knowledge and inquiry, and knowledge and chance. It was also shown that the account entrained problems of its own – in particular, it seems to have the consequence that utterances such as ‘DeRose does not know that the bank will be open. But if he had less at stake, he would’ should seem true and appropriate.

⁴⁹ The ‘might’ in the consequent of the counterfactual is included to fend off potential appeals to the idea that knowledge is required for appropriate assertion. Cf. fn. 41, above.

The leading concern therefore remains that in regard to the HIGH-LOW contrast, and *also* in regard to our judgments concerning utterances of the form ‘I know that p, but I must inquire further’, and ‘I know that p, but there’s a chance that not-p’, the traditional invariantist is forced to attribute significant speaker error. In the next chapter, I consider an alternative explanation of the HIGH-LOW contrast, one centred on pragmatic features of language. It will be argued that this account is no more effective in addressing the leading concern.

Chapter 2: Traditional Invariantism and Pragmatics.

In this chapter, I explore the suggestion that the HIGH-LOW contrast has its source in some pragmatic features of language. In §1, I sketch the most prominent proposal along these lines. In §2, I discuss a problem with that particular proposal. In §3-5, I widen the scope of the critical discussion, and argue that any traditional invariantist pragmatic proposal will prove unsuccessful.¹

§1. The Relevance View.

In what follows, I shall continue to discuss the HIGH-LOW contrast as found in the Bank Case. It will be useful to quickly review the relevant judgments.

DeRose's utterance of (1) in LOW seems true and appropriate.

(1) I know the bank will be open.

His utterance of (2) in HIGH also seems true and appropriate, as would an utterance of (3).

(2) I don't know the bank will be open.

(3) I don't know that the bank hasn't changed its hours.

Finally, – though such judgments have not been in focus thus far – an utterance of (1) or (4) by DeRose in HIGH seems false and inappropriate.

(4) I know that the bank hasn't changed its hours.

¹ In writing the present chapter, I gratefully acknowledge the benefit of discussions with Torfinn Huvenes. He and I are composing joint work – provisionally entitled 'A Pragmatic Defence of Invariantism?' – on many of the issues dealt with here.

On a traditional invariantist account, utterances of (1) and (4) by DeRose are true in both LOW and HIGH and utterances of (2) and (3) by DeRose are false in both LOW and HIGH. The task is thus to explain away the apparent falsity of uttering (1) or (4) in HIGH and the apparent truth of uttering (2) or (3) in HIGH.

The outline of any pragmatic explanation of these judgments is fairly straightforward. First, the traditional invariantist will claim that an utterance of either (1) or (4) by DeRose in HIGH pragmatically conveys something false, whilst an utterance of either (2) or (3) by DeRose in HIGH pragmatically conveys something true. Second, she will claim that our judgments about truth and appropriacy are driven by facts about what those utterances pragmatically convey, rather than what they semantically express.²

The problems with attempts to explain away the apparent falsity of an utterance of (1) by DeRose in HIGH and the apparent truth of an utterance of (2) by DeRose in HIGH readily carry over to attempts to explain away the apparent falsity of an utterance of (4) by DeRose in HIGH and the apparent truth of an utterance of (3) by DeRose in HIGH. In what follows, I shall therefore restrict the discussion to those pragmatic explanations targeted at utterances of (1) and (2) by DeRose in HIGH.

What sort of initial constraints are there on a pragmatic explanation of this kind? I shall assume that any pragmatic explanation of our Bank Case judgments should feature only *general* conversational rules and invoke only *respectable* categories of pragmatic implication. In order that her proposal remain methodologically respectable, the traditional invariantist cannot appeal to conversational rules tailored to knowledge ascriptions alone, nor can she invoke some *sui generis* species of pragmatic implication.³

² In what follows, I shall speak of what is ‘pragmatically conveyed’ or ‘pragmatically implicated’ by an utterance. This language should be interpreted in as theory-neutral a fashion as possible. I wish to remain neutral on such issues as whether the invariantist must claim that a Gricean conversational implicature is present (although note that Hazlett (2009, p. 609) does) rather than some other (respectable) species of pragmatic implication.

³ See DeRose (1999, 2002, 2009). DeRose (1999) proposes two further constraints on any pragmatic explanation of our Bank Case judgments. These are much more controversial. See Brown (2006) for relevant discussion.

Let's consider the most prominent of the traditional invariantist pragmatic proposals.⁴ The proposal is centred around an appeal to Grice's maxim of Relation – the conversational injunction to 'be relevant!'.⁵ In virtue of the operation of this general conversational rule, it has been suggested that knowledge ascriptions (and denials) pragmatically convey various things about the epistemic position of the subject beyond the fact that she knows (or doesn't know) the proposition at issue. The view can readily be presented in terms of any number of characterisations of epistemic position, but for present purposes I shall follow Rysiew (2001) and Brown (2005) and characterise epistemic position in terms of the range of alternatives or error possibilities that the subject can rule out. Nothing turns on this particular choice.

Call the set of alternatives a subject must be able to rule out in order to know some proposition *p* the 'knowledge alternatives' to *p*.⁶ In the Bank Case, we are told that DeRose knows that the bank will be open on Saturday. It follows that he is able to rule out the knowledge alternatives to the proposition that the bank will be open on Saturday. Nevertheless, it might seem that his evidence is such that he is unable to rule out various more distant or unlikely alternatives to that proposition, such as the possibility in which the bank has recently changed its hours.

Now consider the HIGH context. DeRose and his wife are under a pressing financial obligation and DeRose's wife has just raised the possibility that the bank has changed

⁴ The view is defended in Brown (2005, 2006), Hazlett (2009), and Rysiew (2001, 2007). Rysiew (2001, 2007) offers another pragmatic explanation suited specifically to explain judgments about first-person knowledge attributions. The proposal (roughly put) is that a speaker can convey that she knows *p* simply by asserting that *p*, and so if a speaker chooses to assert that she knows *p*, her audience will assume she is trying to convey something *more* than that she knows *p* – in particular, Rysiew proposes that it will be assumed she is trying to convey that she can rule out the alternatives to *p* that are salient in the conversation (or something along those lines). However, this explanation does not generalize to third-person cases. It is also thrown into doubt by the observation that a speaker seems to convey that she can rule out the alternatives to *p* that are salient in the conversation (or something along those lines) by simply asserting that *p* (cf. DeRose 2002). I therefore set it aside.

⁵ Defenders of this view have typically been explicit in their appeals to the Gricean conversational maxim. It is clearly possible to present the view in terms of some more theory-neutral notion of conversational relevance. Such manoeuvring won't affect the criticisms to follow.

⁶ I am using the term 'knowledge alternatives' rather than the more familiar 'relevant alternatives' (issuing from the Relevant Alternatives Theory of Knowledge) to avoid confusion with the notion of conversational relevance.

its hours. It might therefore seem plausible that the relevant issue in HIGH is whether or not DeRose can rule out the possibility that the bank has changed its hours. Relatedly, it might seem that this – or something like it – is what DeRose’s wife is *really* seeking to establish in asking DeRose whether or not he knows the bank will be open.⁷

It has been suggested that these facts might be crucial in explaining our Bank Case judgments. In particular, it has been proposed that, on the assumption that DeRose is being relevant, it follows that an utterance of (1) by DeRose in HIGH will convey that he can rule out the possibility that the bank has changed its hours – something that is false, given his evidence. Similarly, it has been proposed that, on the assumption that DeRose is being relevant, an utterance of (2) by DeRose in HIGH will convey that he cannot rule out the possibility that the bank has changed its hours – something that is true, given his evidence.

In order to complete the explanation of why an utterance of (1) by DeRose in HIGH seems inappropriate and *false* and why an utterance of (2) by DeRose in HIGH seems appropriate and *true*, it has then been suggested that we tend either to confuse what is semantically expressed by those utterances with what they pragmatically convey or to focus on what they pragmatically convey rather than on what they semantically express.⁸

An initial issue. The characterisation of the relevant issue in the conversation in HIGH, and what is pragmatically conveyed by utterances of (1) and (2) in HIGH as a result, may strike some readers as excessively simple-minded. For example, those sympathetic to a relevance-based view might suggest instead that the relevant issue in HIGH is really whether or not DeRose can rule out a *range* of possibilities that includes the possibility that the bank has changed its hours, and that what is therefore pragmatically conveyed by utterances of (1) and (2) DeRose in HIGH is that DeRose can and cannot (respectively) rule out such a range of possibilities. Or perhaps they

⁷ See esp. Brown (2005a, p. 281; 2006, p. 426).

⁸ For the former suggestion, see Brown (2006, p. 428), Hazlett (2009, p. 610), and Rysiew (2001, p. 496; 2007, p. 640). For the latter, see Rysiew (2001, p. 486-7).

might suggest that an utterance of (1) by DeRose in HIGH conveys that *either* he can rule out such a range of possibilities *or* that such possibilities are not worth considering (and an utterance of (2) by DeRose in HIGH conveys the negation of that disjunction).⁹ As in the case of the particular characterisation of epistemic position being deployed, these kinds of subtlety are not central to the discussion that follows. I shall therefore just continue to talk as though the central claims of the relevance-based account are that the relevant issue in HIGH is whether or not DeRose can rule out the possibility that the bank has changed its hours and that what is pragmatically conveyed by utterances of (1) and (2) by DeRose in HIGH is that DeRose can and cannot (respectively) rule out that the bank has changed its hours.

In the remaining sections, I first press an objection against the relevance-based pragmatic proposal outline above, before moving on to introduce some more general problems affecting any traditional invariantist pragmatic proposal.

§2. A Problem for The Relevance View.

A first concern – limited to the relevance-based account – is grounded in the observation that merely assuming that DeRose is being relevant seems insufficient to explain why an utterance of (1) by DeRose in HIGH would convey that he can rule out the possibility that the bank has changed its hours.¹⁰

In general, if the relevant issue in a conversation is whether or not *p* is the case, then contributions that convey that *p* is the case and contributions that convey that *p* is not the case are presumably equally relevant. It is proposed that the relevant issue in HIGH is whether or not DeRose can rule out the possibility that the bank has changed its hours. It would seem to follow that, in HIGH, a contribution that conveys that DeRose can rule out the possibility that the bank has changed its hours is just as relevant as a contribution that conveys that he cannot rule out that possibility.

⁹ See Brown (2006) and Rysiew (2007) for suggestions along these lines.

¹⁰ See also DeRose (2009, p. 121-123) for an objection to the relevance-based explanation of the apparent truth of (2) as uttered by DeRose in HIGH.

This strongly suggests that the mere assumption that DeRose is being relevant is insufficient to explain why an utterance of (1) by DeRose in HIGH would convey that he is able to rule out the possibility that the bank has changed its hours rather than that he is *unable* to rule out that possibility. If an utterance of (1) by DeRose in HIGH would indeed convey that he is able to rule out that possibility, it seems to follow that there must be some *additional factor* – some feature of an utterance of (1), some conversational presupposition in HIGH, ... – that accounts for why that is so.

Let's consider some proposals for what that factor might be. If DeRose were to utter (1) in HIGH, he would be responding affirmatively to his wife's question 'Do you know that the bank will be open tomorrow?'. Perhaps that is the further feature at issue. Recall that it was suggested that in asking 'Do you know that the bank will be open tomorrow?' in the HIGH context, DeRose's wife is really asking the question whether or not DeRose can rule out the possibility that the bank has recently changed its hours. If DeRose were to utter (1) in HIGH, he would be giving an affirmative answer to his wife's explicit question. It therefore seems reasonable to expect that he would also convey an affirmative answer to her implicit question – that is, it seems reasonable to expect that he would convey that he can rule out the possibility that the bank has changed its hours.

The problem with this response is that it is insufficiently general. Even if we modify the HIGH context in such a way that DeRose's wife no longer asks 'Do you know that the bank will be open?', an utterance of (1) by DeRose in HIGH still seems inappropriate and false. For example, Stanley's (2005a) presentation of the Bank Case features no such question, but that does not seem to affect the judgments at issue.

An alternative proposal is that the feature in question is some item of *background knowledge* possessed by DeRose and his wife.¹¹ Suppose that it's part of the background knowledge of DeRose and his wife in HIGH that if DeRose knows that the bank is open, then it's more likely than not that he is able to rule out the possibility that the bank has changed its hours. If that were the case, it might go some way

¹¹ Cf. Grice (1975, p. 50).

towards explaining why, on the assumption that DeRose is being relevant, an utterance of (1) by DeRose in HIGH would convey that he can rule out the possibility that the bank has changed its hours, rather than that he cannot rule out that possibility.

The central problem with this suggestion is also that it is insufficiently general. Let it be stipulated that it is part of DeRose and his wife's background knowledge in HIGH that if DeRose occupies an epistemic position of at least the minimum strength required for him to know (on the traditional invariantist account) that the bank will be open tomorrow, then he will more likely than not be *unable* to rule out the possibility that the bank has changed its hours. Perhaps that minimum strength, on the traditional invariantist account, is equivalent to something like the strength of epistemic position provided by the firm recollection of having been at the bank one month ago on a Saturday. In that case, let it be stipulated that it is part of DeRose and his wife's background knowledge in HIGH that if DeRose occupies an epistemic position of at least that strength, then he will more likely than not be unable to rule out the possibility that the bank has changed its hours.

Given that stipulation, it is clearly not going to be part of DeRose and his wife's background knowledge in HIGH that if DeRose knows the bank will be open, it is more likely than not that he *is* able to rule out the possibility that the bank has changed its hours. The problem is that even with the stipulation in place, an utterance of (1) by DeRose in HIGH still seems inappropriate and false – he still seems to convey that he can rule out the possibility that the bank has changed its hours.¹²

¹² Another problem. Assuming that DeRose and his wife are fairly typical subjects, it's probably not very natural to suppose that the proposed conditional will form part of their background knowledge in the Bank Case *as originally described*. The reason is that on a traditional invariantist account it's fairly easy to know that one's bank will be open on Saturday – it's sufficient to have been there two weeks ago on a Saturday, to have been informed by a friend who may or may not have checked very recently that it will be open, or such. It therefore seems that the typical subject who knows that their bank will be open on Saturday is probably more likely than not *unable* to rule out the possibility that their bank has changed its hours in the past two weeks. Consequently, it's probably not very natural to suppose even in the original case that it is part of DeRose and his wife's background knowledge that if DeRose knows that the bank is open, then it's more likely than not that he is able to rule out the possibility that it has changed its hours. Thus, this view may well be unable to account for why an utterance of (1) by DeRose in HIGH conveys that he can rule out

It is rather unclear what other feature of the HIGH case might be responsible for an utterance of (1) by DeRose in HIGH conveying, on the assumption that DeRose is being relevant, that he can rule out the possibility that the bank has changed its hours. It therefore seems that we have identified a genuine challenge for those traditional invariantists sympathetic to a relevance-based pragmatic explanation – a challenge that it is far from straightforward to meet successfully.

§3. Comparisons with other Pragmatic Effects.

The preceding discussion was focused on issues specific to a relevance-based explanation of our Bank Case judgments. In the remaining sections, I move on to discuss several more general problems, which cast serious doubt on the viability of *any* traditional invariantist pragmatic proposal.

However, it will be useful to first make a concession. Henceforth, I shall just assume that whatever it is that the traditional invariantist wishes to claim is pragmatically conveyed by knowledge ascriptions and denials is pragmatically conveyed by those ascriptions and denials. It will nevertheless be helpful to have a slightly less roundabout way to talk about the pragmatically conveyances at issue. It seems fair to assume that whatever it is that is pragmatically conveyed by knowledge ascriptions and denials it is something broadly tied to the subject's epistemic position. In the next two sections (it potentially makes a difference in §5), I shall therefore just continue to talk as though an utterance of (1) by DeRose in HIGH and an utterance of (2) by DeRose in HIGH would convey that DeRose can and that he cannot (respectively) rule out that the bank has changed its hours. But once again, nothing turns on this.

In this section, I compare the pragmatic effect allegedly associated with knowledge ascriptions with other pragmatic effects. The revealed similarities and contrasts will provide some reason to be suspicious of a traditional invariantist pragmatic proposal. They will also set the stage for the discussion in the sections that follow.

that bank has changed its hours even in the case originally described (i.e. without any additional stipulations).

Consider an uncontroversial case of pragmatic implication:

(7) Jane has good handwriting.

Suppose Mary utters (7) in response to a query about the quality of Jane's philosophical work. Mary's utterance conveys that Jane is not a very good philosopher.

Suppose that Jane has good handwriting but that Jane is in fact a very good philosopher. In that case, it seems plausible that Mary's true utterance pragmatically conveys a falsehood. In this respect, her utterance would seem to be just like an utterance of (1) by DeRose in HIGH. However, whilst Mary's utterance of (7) seems inappropriate, it does not seem false – after all, Jane does have good handwriting. There is therefore a salient contrast with an utterance of (1) by DeRose in HIGH, which seems both inappropriate and false.

What about an utterance of (2) by DeRose in HIGH? Suppose that Mary utters (7) in response to a query about the quality of Jane's philosophical work. But suppose that Jane does not have good handwriting and is a poor philosopher. In that case, it seems plausible that Mary's false utterance pragmatically conveys something true. In this respect, her utterance would seem to be just like an utterance of (2) by DeRose in HIGH. However, whilst Mary's utterance of (7) *may* seem appropriate, it certainly does not seem true. There is therefore a salient contrast with an utterance of (2) by DeRose in HIGH, which seems both appropriate and true.

It was noted earlier that traditional invariantists have proposed that the reason we see the kinds of judgments regarding truth and appropriacy in the Bank Case that we do is that we tend either to confuse what is semantically expressed by knowledge ascriptions with what they pragmatically convey, or to focus on what they pragmatically convey rather than what they semantically express. As we have just seen, we clearly do not exhibit this kind of tendency in uncontroversial cases of pragmatic phenomena like those associated with (7). It is therefore pertinent to ask to how plausible it is to propose that we exhibit this tendency in the case of knowledge

ascriptions. In particular, it is pertinent to ask whether we exhibit this tendency with regard to any other cases of pragmatic phenomena.

DeRose (1999, 2002, 2009) has, in effect, suggested that as applied to the case of a false utterance of (2) by DeRose in HIGH the confusion/focusing proposal is not very plausible. However, as Brown (2006) points out, there are various more controversial cases of pragmatic phenomena in which we find judgments like those we find in the Bank Case – including our judgments regarding an utterance of (2) by DeRose in HIGH.¹³

(8) Peter has eaten

Suppose that Mary is discussing whether to get Peter something to eat and that Peter has had nothing to eat all day. In that case, an utterance by Mary of (8) seems inappropriate and false. And an utterance by Mary of (9) seems appropriate and true.

(9) Peter has not eaten.

On one proposal, what is semantically expressed by an utterance of (8) is that Peter has eaten at some past time, and what is semantically expressed by an utterance of (9) is that Peter has never eaten. For reasons broadly connected with relevance, it has been suggested that an utterance by Mary of (8) pragmatically conveys that Peter has eaten recently, whilst an utterance by Mary of (9) pragmatically conveys that Peter hasn't eaten recently.

Let's assume this pragmatic proposal is correct. In that case, a true utterance of (8) by Mary seems inappropriate and false because it pragmatically conveys something false. Conversely, a false utterance of (9) by Mary seems appropriate and true because it

¹³ Cases featuring utterances of (8) and (9) are putative cases of what Bach terms 'implicature' (see Bach (1994)). Brown (2006) highlights another (controversial) case, featuring referential uses of definite descriptions, in which we also find similar judgments to those found in the Bank Case. There are plausibly other (still controversial) cases besides – pragmatic treatments of quantifier domain restriction, for example. The 'implicature' example is, however, representative of these cases in all the respects that will concern us here.

pragmatically conveys something true. It therefore seems that, in this case, we exhibit a tendency to confuse what is semantically expressed by utterances of (8) and (9) with what they pragmatically convey, or perhaps to focus on what they pragmatically convey rather than on what they semantically express. This is clearly a welcome finding for the traditional invariantist – it seems that she is not forced to postulate some *sui generis* kind of semantics-pragmatics confusion.

It should be stressed, however, that it is far from uncontroversial that the kind of pragmatic treatment of utterances of (8) and (9) that we have considered is the right treatment. Nevertheless, in what follows I shall simply grant the traditional invariantist that cases like those involving utterances (8) and (9) warrant a pragmatic treatment. A central purpose of the remaining sections is to show that even supposing that such cases represent genuine cases of pragmatic phenomena, there are serious problems with the traditional invariantist's pragmatic manoeuvre.

§3.1 Disagreement.

One potential problem concerns the interplay between our judgments about truth and appropriacy and our judgments about disagreement. Suppose Emily is in a context similar to LOW in which the stakes are low and relatively distant error possibilities are not being considered. Suppose Emily overhears the conversation between DeRose and his wife in LOW. In Mary's context, an utterance of (10) seems true and appropriate.

(10) Keith knows that the bank will be open.

Now suppose that Tom is apprised of the details of DeRose's epistemic position in regard to the bank's being open, but that Tom occupies a different context, one that is similar to HIGH. In Tom's context, an utterance of (11) seems true and appropriate.

(11) Keith doesn't know that the bank will be open.

It might seem that if Emily were to utter (10) and Tom were to utter (11), they would disagree.¹⁴ In the same vein, suppose that Kurt is acquainted with Emily and Tom's respective situations and overhears both of their utterances. In that case, it might seem true and appropriate for Kurt to utter a report like (12).

(12) Emily and Tom disagree about whether Keith knows that the bank will be open.

What is the significance of these judgments? It might seem as though these judgments are just what we should expect on an invariantist picture. After all, on such a picture, what is semantically expressed by Tom's utterance of (11) is the negation of what is semantically expressed by Emily's utterance of (10). However, there are in fact good reasons to think that these are not the judgments we should expect to find.

In the example, Emily is in a context like LOW. On the views under consideration, her utterance of (10) will therefore convey what it semantically expresses – viz. that DeRose knows that the bank will be open. That explains why her utterance seems true and appropriate. But what does Tom's utterance of (11) convey? Tom is in a context like HIGH. His utterance of (11) will therefore *not* convey what it semantically expresses – viz. the falsehood that DeRose does not know that the bank will be open. If Tom's utterance conveyed that, then presumably his utterance would seem not seem true and appropriate. Instead, Tom's utterance conveys that DeRose cannot rule out the possibility that the bank has changed its hours. That explains why his utterance seems true and appropriate. But given that what Emily and Tom convey by their respective utterances is perfectly compatible, it is unclear why it would seem as though they disagree.

That argument is very suggestive. But the situation becomes more serious when we observe that in those cases of (allegedly) pragmatic phenomena in which we find judgments similar to those that we find in the Bank Case with respect to truth and

¹⁴ See Feldman (2001), Richard (2004), and esp. MacFarlane (2007, p. 19-21; *f/c*) for discussion pertaining to our disagreement judgments. There seems to be some dispute in the literature regarding our disagreement judgments in cases like these; this dispute – and its consequences for the current critique – is discussed in Chapter 5.

appropriacy, we do not find the kinds of judgments about disagreement reported in the case of (10) and (11).

Let's return to the case introduced in the previous section. Suppose that Mary utters (8), but this time in the course of discussing whether Peter has ever eaten.

(8) Peter has eaten.

Suppose that Peter has eaten at some point. In that case, Mary's utterance of (8) resembles Emily's utterance of (10) in a context like LOW – it conveys what it semantically expresses and what it semantically expresses is true. That accounts for why it seems true and appropriate. Now suppose that at the same time that Mary utters (8), John is discussing whether to get Peter something to eat. John utters (9).

(9) Peter has not eaten.

Suppose that although Peter has eaten at some point, he has not eaten all day. In that case, John's utterance of (9) resembles Tom's utterance of (11) in a context like HIGH – his utterance is false, but pragmatically conveys something true. That accounts for why it seems true and appropriate. However, unlike in the case of Tom and Emily, there is little temptation to judge that Mary and John disagree. In this case, it is far from clear that a report like (13) would be true and appropriate.

(13) Mary and John disagree about whether Peter has eaten.

On the kind of pragmatic treatment introduced in the previous section, what is semantically expressed by (8) is incompatible with what is semantically expressed by (9). Nevertheless, there is little or no appearance of disagreement in the relevant cases.

It seems fair to conclude that our judgments concerning disagreement represent a *prima facie* problem for a traditional invariantist pragmatic treatment of our Bank Case judgments.¹⁵ I now turn to introduce another, perhaps more serious, problem.

§4. Retraction.

In the previous section, I argued that given the kinds of judgments about truth and appropriacy that we find in the Bank Case, we should expect certain kinds of judgments about disagreement (on the assumption that a pragmatic view is correct). In this section, I argue that the traditional invariantist faces difficulty with respect to our judgments about retraction. In this case, not only are our judgments somewhat mysterious, they also seem to contrast with both non-controversial and more controversial pragmatic effects.

Suppose that DeRose utters (1) in LOW.

(1) I know the bank will be open.

And suppose that at some later time, DeRose and his wife enter a context with features like HIGH. For example, suppose DeRose and his wife find themselves in conversation with a friend who explains that he has an important cheque coming due and who raises the possibility that the bank has changed its hours. Suppose that in this new context DeRose utters (2).

(2) I don't know the bank will be open.

In that context, DeRose's utterance of (2) seems true and appropriate. On the account under consideration, that is because his utterance pragmatically conveys the truth that he cannot rule out the possibility that the bank has changed its hours.

¹⁵ See also Montminy (2009).

Suppose that DeRose's wife then questions DeRose over his earlier utterance of (1). In that case, it has been suggested that an utterance by DeRose of (14) seems true and appropriate.

(14) I spoke falsely

Similarly, an utterance of (15), in reference to his earlier utterance of (1), seems appropriate.¹⁶

(15) I take that back.

These judgments are a cause for concern. DeRose uttered (1) in LOW. On the proposals we are considering, it follows that his utterance was true and conveyed nothing false. But given that his utterance of (1) in LOW was true and conveyed nothing false, why would it seem appropriate for him to retract that assertion, or seem true and appropriate for him to claim that it was false?¹⁷

This problem is made more pressing when we observe that this is not what we find in either the uncontroversial cases of pragmatic phenomena, like those involving utterances of (7), or in the more controversial cases, like those involving utterances of (8).

(7) Jane has good handwriting

(8) Peter has eaten

Let's take the less controversial cases first. Suppose Peter asks Emily about the quality of Jane's handwriting. In response, Emily utters (16).

¹⁶ For reports of such judgments, see e.g. Hawthorne (2004a, p. 163), Williamson (2005a, §II), MacFarlane (2005, §2.3), Stanley (2005a, p. 115), and Davis (2007, p. 399-400).

¹⁷ The traditional invariantist has no special reason to think that knowledge of knowledge is hard to come by. Thus we can even assume that in his current context (resembling HIGH) DeRose knows that he knew (when in LOW) that the bank will be open.

(16) Jane does not have good handwriting.

Suppose that Jane does not have good handwriting. In that case, Emily's true utterance of (16) straightforwardly conveys what it semantically expresses and does not convey anything false. In this respect, her utterance resembles DeRose's true utterance of (1) in LOW. Now suppose that at some later time, Emily and Peter find themselves in a different context, in which a friend has just asked about the quality of Jane's philosophy. In this context, Emily replies by uttering (7).

(7) Jane has good handwriting

Suppose that Jane is a poor philosopher. In that case, Emily's false utterance of (7) conveys the truth that Jane is not a very good philosopher. In this respect, her utterance resembles DeRose's false utterance of (2) in the context like HIGH, which conveys the truth that DeRose cannot rule out the possibility that the bank has changed its hours.

Suppose that Peter then questions Emily about her earlier utterance of (16). In that case, it seems *false* for Emily to utter (14).

(14) I spoke falsely

Similarly, it seems very strange for Emily to utter (15).

(15) I take that back.

Once questioned, it also seems appropriate for Emily to retract her most recent assertion of (7) rather than to retract her original assertion of (16). This clearly contrasts sharply with the case of knowledge ascriptions.

Let's turn now to the more controversial cases of pragmatic phenomena introduced in §3. Suppose Mary and John are discussing whether Peter has ever eaten, and Mary utters (8).

(8) Peter has eaten

Suppose that Peter has eaten at some past time. In that case, Mary's true utterance of (8) straightforwardly conveys what it semantically expresses and does not convey anything false. In this respect, it resembles DeRose's true utterance of (1) in LOW. Now suppose that a short time later Mary and John are approached by a friend, who asks whether they should get Peter something to eat. In that new context, Mary utters (9).

(9) Peter has not eaten

Suppose that Peter has had nothing to eat all day. In that case, Mary's false utterance of (9) conveys the truth that Peter has not eaten recently. In this respect, her utterance resembles DeRose's false utterance of (2) in the context resembling HIGH, which conveys the truth that DeRose cannot rule out the possibility that the bank has changed its hours.

Suppose that John then questions Mary over her earlier utterance of (8). In that case, it once again seems *false* for Mary to utter (14).

(14) I spoke falsely

Similarly, it seems inappropriate for Mary to utter (15).

(15) I take that back.

Once questioned then, it seems inappropriate for Mary to retract her previous assertion of (8). Our judgments in this case clearly also contrast sharply with the case of knowledge ascriptions.

How might the traditional invariantist respond to the disagreement and retraction problems? One response to our judgments about retraction is to attempt a further pragmatic manoeuvre. It might be proposed that in retracting his earlier assertion of

(1) in LOW, DeRose would pragmatically convey something true, whilst in standing by that assertion, he would pragmatically convey something false.

There are perhaps two main issues with this kind of response. The first issue concerns the contrast with our judgments in other cases of pragmatic phenomena. The natural way to cash out this pragmatic response is to propose that the considerations – perhaps having to do with relevance – that lead a false utterance of (2) by DeRose in the context resembling HIGH to seem true and appropriate *also* lead a false utterance of (14) by DeRose in the context resembling HIGH – that is, a claim to the effect that his earlier utterance of (1) was false – to seem true and appropriate. But if that were the case, then why (for instance) wouldn't the considerations – those having broadly to do with relevance – that lead a false utterance of (9) by Mary, in the context in which she is discussing whether to get Peter something to eat, to seem true and appropriate *also* lead a false utterance (14) by Mary in the same context – that is, a claim to the effect that her earlier utterance of (8) was false – to seem true and appropriate? As we have seen, they do not. This kind of contrast is surely troubling.

The second issue is that one can set up the context resembling HIGH in such a way that the truth (or otherwise) of the speaker's earlier assertions is central to the current discussion (MacFarlane (2005)). Suppose that after having uttered (1) in LOW, DeRose later finds himself on the witness stand. This courtroom context is similar to HIGH – the presiding judge has just raised the possibility that the bank has changed its hours and there is a great deal at stake. Suppose the judge asks DeRose 'Do you know that the bank will be open tomorrow?' and that DeRose responds by uttering (2). The judge then questions DeRose about his earlier utterance of (1), perhaps implying that DeRose might be rather an inconstant witness. In this case, it still seems true for DeRose to admit that his earlier assertion of (1) was false.¹⁸

¹⁸ In this kind of case, it might seem that although it would be true for DeRose claim that his earlier assertion was false, it might nevertheless be appropriate, in such circumstances, for him to add 'but obviously I was speaking loosely earlier' and thereby excuse his earlier assertion (cf. Schiffer 1996, p. 328). This hardly alleviates the problematic situation for the traditional invariantist; indeed, it may only serve to deepen the mystery. On the kind of account under discussion, DeRose's earlier knowledge ascription in LOW will have conveyed what it semantically expressed – that DeRose knows that the bank will be open – and will not

The problem is that given the nature of the circumstances, it seems that if DeRose's earlier assertion of (1) really was true (and conveyed nothing false) it seems like it really ought not to seem true for him to admit that that assertion was false. It is much harder to envisage a plausible pragmatic account of our judgments in this case.

An alternative response for the traditional invariantist is to postulate an additional kind of semantics-pragmatics confusion. It might be suggested that when in the context resembling HIGH, DeRose thinks that what he managed to convey by uttering (1) in LOW is the same as what he would convey by uttering (1) in the context resembling HIGH. If that were the case, it might explain why DeRose would retract his earlier assertion. A similar proposal could be put forward in response to the potential problem surrounding disagreement introduced in the previous section. Recall Kurt, the third-party who overhears Emily and Tom's utterances of (10) and (11) and reports them as disagreeing by uttering (12).

(10) Keith knows that the bank will be open

(11) Keith doesn't know that the bank will be open

(12) Emily and Tom disagree about whether Keith knows that the bank will be open

If Kurt assumed that Emily and Tom had both conveyed by their utterances what he would convey by uttering (10) and (11) in his own context, that might explain why he would judge that Emily and Tom disagree.

Perhaps the account can be extended to explain why it seems appropriate *to us* for DeRose to retract his assertion and why it might seem true and appropriate *to us* for Kurt to utter (12). Even supposing it can, the suggestion is fairly radical. The proposal is not merely that we confuse what is semantically expressed by an utterance with

have conveyed anything false. Presumably then, DeRose was speaking *strictly*, not loosely. See Chapter 5, §4 for further discussion.

what it pragmatically conveys. The proposal is instead that we confuse what is semantically expressed by an utterance (not with what it pragmatically conveys but) with what is pragmatically conveyed by an utterance of the same sentence in some *other* context.

However, given DeRose is fully aware of the relevant differences between his former circumstances in LOW and his current circumstances in the context resembling HIGH, why would he make such a mistake? Similarly, given Kurt is aware of the differences between Emily and Tom's respective circumstances, why would he make that kind of mistake? It certainly seems rather mysterious. Furthermore, as we have seen in the cases presented above, we do not appear to exhibit this kind of mistake in other cases of pragmatic phenomena – both of the uncontroversial and the more controversial variety. It is therefore not only rather mysterious why we would make such a mistake, but it also seems a mistake unique to the pragmatic effects associated with knowledge ascriptions. This is clearly an unhappy result.

I return to the issues of retraction and disagreement – and the recent dispute concerning the true character of those judgments – in Chapter 5. I argue that whilst our disagreement judgments may not in fact represent a serious problem for the traditional invariantist pragmatic account, our retraction judgments do represent such a problem. (However, it is suggested that it is nevertheless important to observe that the kind of judgments concerning disagreement that are to be expected given a pragmatic defence of invariantism, are not those that might be expected given other defences.) In the next section, I present one final – and particularly pressing – problem, which once again reveals a sharp contrast between the pragmatic effect (allegedly) associated with knowledge ascriptions and both uncontroversial and more controversial pragmatic effects alike.

§5. Cancellability

Another pressing concern for the traditional invariantist focuses on the issue of cancellability. Consider the following example.¹⁹

(6) Peter: I've run out of petrol. Is there a garage nearby?

Mary: There's one around the corner.

Mary's response seems to convey that the garage is open. On a popular view, that is because her response pragmatically implicates that the garage is open. However, Mary can cancel that implication. For example, by an utterance of (17).

(17) There's one around the corner, but it's not open.

Following Grice, we can understand what it is for a pragmatic implication to be cancellable as follows.²⁰ Suppose an utterance pragmatically implicates that *p*. That implication is cancellable just in case the speaker can felicitously add something along the lines of 'but not *p*' and thereby avoid any commitment to *p*.

The worry is that the pragmatic implication invoked to explain the apparent falsity of an utterance of (1) by DeRose in HIGH does *not* seem to be cancellable. For example, as remarked in the previous chapter, an utterance of (18) by DeRose in HIGH does not seem felicitous.²¹

(18) I know that the bank will be open, but we need to investigate further.

And the situation is similar for each of the following cancellation attempts:

(19) a. I know that the bank will be open, but I cannot rule out that it has changed its hours.

¹⁹ The example is drawn from Grice (1975, p. 51).

²⁰ See Grice (1975, p. 57). In Grice's terminology, what follows is a characterisation of what it is to be explicitly (rather than contextually) cancellable.

²¹ Cf. Cohen (1999, p. 60).

- b. I know that the bank will be open, but I'm not in a strong enough position for us to rely on it being open.
- c. I know that the bank will be open, but it's possible that it has changed its hours.

There are two strategies the invariantist can pursue at this point. The first is to argue that, *contra* initial appearances, the relevant pragmatic implications are in fact cancellable. The second is to suggest that even though they are not cancellable, this is not a problem. I evaluate each in turn.

§5.1 Option 1: The Implication is Cancellable

In addressing the cancellability objection, Brown and Rysiew suggest that DeRose might be able to cancel the implications at issue by other means than uttering (18) or (19a-c). Let's review their suggestions.

Rysiew (2001, p. 495) suggests that an utterance of (20) by DeRose in HIGH would be felicitous.²²

(20) I know that the bank will be open, but of course I cannot rule out the *bizarre* alternatives to the bank's being open.

However, an utterance of (20) can only cancel the implication generated by an utterance (1) by DeRose in HIGH if the latter implicates that DeRose can rule out the *bizarre* alternatives to the bank's being open (or something that entails that he can rule out the *bizarre* alternatives). But it doesn't.

If DeRose's utterance of (1) in HIGH implicated that he could rule out the *bizarre* alternatives to the bank's being open, then his utterance would presumably seem inappropriate and false if he could not rule those alternatives out. But suppose that the

²² Rysiew (2001, p. 495) actually makes the more general claim that utterances of the form 'I know that p, but I cannot rule out the *bizarre* alternatives to p' seem felicitous.

bank's manager informed DeRose earlier in the day that the bank would be open on Saturday. DeRose's epistemic position is now such that were he to utter (1) in HIGH his utterance would seem appropriate and true. However, although improved, his epistemic position is nevertheless still such that he cannot rule out the bizarre alternatives to the bank's being open – he still cannot rule out possibilities like the possibility that the bank will quantum tunnel to Alaska. The invariantist must do better.

Brown (2006, p. 428) suggests that utterances of (21a) and (21b) by DeRose in HIGH would be felicitous.²³

(21) a. I know that the bank will be open, but I am not in a really strong epistemic position with respect to the proposition that the bank will be open.

b. I know that the bank will be open, but my belief that it will be open doesn't match the facts in a really distant possible world

The first thing to note about (21a) and (21b) is that they feature both technical vocabulary and highly context sensitive vocabulary. It is very hard to judge what expressions like 'really strong epistemic position' or 'really distant possible world' pick out in HIGH. One should therefore be sceptical that even supposing utterances of (21a) and (21b) are felicitous, they serve to cancel the implication generated by an utterance of (1) by DeRose in HIGH.

Perhaps more importantly, when one tries to specify a particular reading for 'really strong epistemic position' or 'really distant possible world', it seems that the only readings on which utterances of (21a) and (21b) seem felicitous are readings on which they do *not* cancel the relevant implication. For example, if 'really strong epistemic position' picks out the kind of epistemic position such that those who occupy it can

²³ Brown (2006, p. 428) actually makes the more general claim that utterances of the form 'S knows that p, but S is not in a really strong epistemic position with respect to p' and of the form 'S knows that p, but S's belief that p does not match the facts of the matter in a really distant possible world' seem felicitous.

rule out the possibility in which the bank quantum tunnels to Alaska, then an utterance of (21a) by DeRose in HIGH might seem felicitous. But for reasons just given in the discussion of (20), that utterance would not serve to cancel the implication generated by an utterance of (1) by DeRose in HIGH. To cancel that implication, ‘really strong epistemic position’ needs to pick out the kind of epistemic position that would enable those who occupy it to rule out that the bank has changed its hours. Or it needs to pick out the kind of epistemic position such that, for those who occupy it, further inquiry in circumstances like those of DeRose and his wife would be unnecessary. But when ‘really strong epistemic position’ picks out this kind of property, an utterance (21a) by DeRose in HIGH no longer seems felicitous. As a result, it does not seem that an utterance of (21a) by DeRose in HIGH can be used to cancel the implication generated by an utterance of (1) by DeRose in HIGH. Similar remarks apply to (21b).²⁴

§5.2 Option 2: The Implication is not Cancellable

The alternative is to accept that the pragmatic effect associated with knowledge ascriptions is not cancellable. Is this plausible?

In the literature on conversational implicature, a failure of cancellability is taken to be a very strong indication that a conversational implicature is not present.²⁵ For example, Grice (1978, p. 44) remarks that he believes that ‘all conversational implicatures are cancelable’; Sadock (1978, p. 372) claims that cancellability is ‘the

²⁴ There is a further problem. As noted, both (21a) and (21b) feature technical vocabulary. This raises the concern that, assuming utterances of (21a) or (21b) represent the only way to cancel the pragmatic effect associated with an utterance of (1) by DeRose in HIGH, the only way *to talk about* what is pragmatically conveyed by an utterance of (1) by DeRose in HIGH is by employing technical vocabulary. Whilst the issues are complicated, it would be at least somewhat surprising if this were the case. After all, ordinary speakers are supposed to be able to work out what is pragmatically conveyed by an utterance of (1) by DeRose in HIGH. Thanks to Torfinn Huvenes here.

²⁵ As is widely known, *conventional* implicatures are not easily cancelled. There is, however, significant dispute over whether there even is such a category as conventional implicature (see e.g. Bach (1999)). Nor has anyone sought to explain our Bank Case judgments by appeal to conventional implicature. I shall not embark on a discussion of the merits of such an explanation here, other than to note that the disagreement and retraction problems of §3.1 and §4 would apply to it. For some additional objections to a conventional implicature story along roughly the lines required to explain the Bank Case phenomena, see the appendix of Fantl & McGrath (2007).

best of the tests' for conversational implicatures; and Levinson (1983, p. 114) suggests that of the properties of conversational implicatures 'perhaps the most important, is that they are cancellable'.

It is probably too quick to assume that the pragmatic effect associated with knowledge ascriptions must be a conversational implicature. Nevertheless, it seems that the considerations that lead us to expect that conversational implicatures will be cancellable should also lead us to expect that the pragmatic effect associated with knowledge ascriptions will be cancellable.

For instance, Levinson (1983, p. 114) supports the claim that conversational implicatures are cancellable by pointing out that they represent the conclusions of *defeasible inferences*. Suppose that an utterance conversationally implicates that p. In that case, the proposition that p will presumably represent the conclusion of an inference from features of the contextual situation (including the participants' background knowledge), the general rules governing conversation, and what is semantically expressed by the utterance at issue. However, such an inference is surely a good candidate for a defeasible inference. That is, it is surely possible to block the inference to p by the introduction of further information. It follows that the implication that p should be cancellable because it should be possible for a speaker to add something along the lines of 'but not-p' to her utterance and thereby introduce the information that not-p and block the relevant inference to p.

Now – on the assumption that the pragmatic implications at issue are respectable – it seems that the pragmatic implication associated with a knowledge ascription is also going to represent the conclusion of an inference from features of the contextual situation (including the participants' background knowledge), the general rules governing conversation, and what is semantically expressed by the knowledge ascription at issue. It therefore seems that it should *also* represent the conclusion of defeasible inference, whether or not it turns out to be a conversational implicature. Consequently, there is an expectation that the implication should be cancellable.

This is clearly troubling for the traditional invariantist. However, Matthew Weiner (2006) has recently observed that there appear to be some pragmatic implications that

are not cancellable.²⁶ Suppose Mary utters (22) to Peter, who is taking up several seats on a crowded train.

(22) I'm curious as to whether it would be physically possible for you to make room for someone else to sit down.

Mary's utterance pragmatically implicates that Peter should make room for someone else to sit down. Suppose Mary attempts to cancel the implication by uttering something along the lines of (23).

(23) I'm curious as to whether it would be physically possible for you to make room for someone else to sit down. Not that you should move, I'm just curious.

In uttering (23), Mary *still* implicates that Peter should make room for someone else to sit down. She therefore does not cancel the implication that would be generated by an utterance of (22).

The problem – if I may so term it – is that (23) is naturally interpreted as ironic as well. In fact, this feature is common to all of Weiner's cases – the attempts to cancel the implication are naturally interpreted as attempts at further irony. In general, it seems that the phenomenon that Weiner is bringing to our attention is that there can be pragmatic effects associated with cancellation attempts that render them unsuccessful.

How does this impact the traditional invariantist proposals under consideration? The following seems plausible. If the traditional invariantist can explain *why* the pragmatic effect associated with an utterance of (1) by DeRose in HIGH is not cancellable – perhaps by arguing that there is some pragmatic effect associated with the relevant cancellation attempts – then the failure of cancellability does not significantly impugn her view. However, if she cannot supply such an explanation, the problem remains in force.

²⁶ Hazlett (2009, p. 608) draws readers' attention to Weiner's paper in his brief discussion of the cancellability objection.

In what follows, I consider and reject what seem to be the most promising ways to carry out that explanatory task. The explanations I shall consider could account for the failure of any of (18) or (19a-c) to cancel the implicature associated with an utterance of (1) by DeRose in HIGH, but the explanations suffer from exactly parallel problems in each case. I shall therefore restrict my focus to traditional invariantist attempts to explain the failure to cancel the relevant implicature by means of an utterance of either (18) or (19a) by DeRose in HIGH.

(18) I know that the bank will be open, but we need to check further

(19a) I know that the bank will be open, but I cannot rule out that it has changed its hours

A first suggestion. Rysiew (2001, p. 496) proposes that when a pragmatic implication is nearly universal we should expect that it will be more difficult to cancel. In support of this contention, he highlights the following case.

(24) They fell in love and got married

Utterances of (24) nearly universally implicate that the relevant individuals fell in love and *then* got married. Rysiew suggests that attempts to cancel that implication are uncomfortable. In particular, he suggests that utterances of (25) are uncomfortable.

(25) They fell in love and got married, but not in that order.

Rysiew then goes on to propose that the reason cancellation attempts are uncomfortable in the case of knowledge ascriptions is that such ascriptions nearly universally convey that the subject can rule out the salient possibilities of error. (On this account, a true utterance of (1) by DeRose in HIGH seems inappropriate and false because it pragmatically conveys that DeRose can rule out the (salient) possibility that the bank has recently changed its hours.)

The proposal is, however, problematic. There may be *some* discomfort associated with utterances of (25). Nevertheless, any discomfort associated with utterances of (25) seems of an entirely different order to that associated with utterances of (18) or (19a) by DeRose in HIGH. Unlike in the case of an utterance of either (18) or (19a) by DeRose in HIGH, utterances of (25) do not seem straightforwardly infelicitous. This suggests that the discomfort associated with the former utterances has a different source to the discomfort associated with the latter.

More importantly, it is far from clear the universal character of the implication is what explains any discomfort associated with utterances of (25). There are a great many implications with such universal character that are comfortably cancelled. For example, utterances of (26) nearly universally implicate that not all of the students passed.

(26) Some of the students passed

However, this implication can be readily cancelled – for example, by an utterance of (27).

(27) Some of the students passed. In fact, they all did.

Indeed, a more plausible explanation of any discomfort with an utterance of (25) is that it constitutes a violation of something like the maxim of Manner – the injunction that one ‘be perspicuous’. To utter (25) is surely a somewhat periphrastic way of conveying the relevant order of love and marriage. There is a much more straightforward way of conveying that order – viz. an utterance of (28).

(28) They got married and then fell in love.

Given the ready cancellation of other ‘universal implications’ this seems a more promising account of the discomfort associated with utterances of (25).²⁷ The prospects for Rysiew’s proposal are not good.²⁸

Let’s pursue a slightly different suggestion. Notice that an utterance of either (29a) or (29b) by DeRose in HIGH seems infelicitous.

- (29) a. The bank will be open, but we need to check further
- b. The bank will be open, but I cannot rule out that it has changed its hours

On the assumption that knowledge is factive – that S knows p entails p – the truth of (18) entails the truth of (29a) and the truth of (19a) entails the truth of (29b). It therefore seems natural to propose that the infelicity of utterances of (18) and (19a) is rooted in the infelicity of utterances of (29a) and (29b).²⁹

It is not necessary to explore the details of the explanation of the infelicity of (29a) and (29b) – such an account will presumably invoke the epistemic constraints on proper assertion – to see that the explanation is insufficiently general. The problem arises when we turn to cases featuring third person knowledge ascriptions.

²⁷ Notice also that were DeRose to respond to his wife’s question ‘Do you know that the bank will be open?’ in HIGH with an utterance of either (18) or (19a) he would not plausibly be violating the maxim of Manner. So the traditional invariantist cannot explain the infelicity of those utterances by appeal to that maxim instead.

²⁸ A further concern. Notice that if we keep the reference of ‘they’ fixed, utterances of (24) convey the same thing about the order in which the relevant individuals fell in love and got married in nearly all contexts. However, even if we keep the reference of ‘I’ fixed, there is an important sense in which what is pragmatically conveyed by utterances of (1) will not exhibit this sort of stability. In some contexts, utterances of (1) will convey that the speaker can rule out one set of alternatives, but in other contexts they will convey that the speaker can rule out wholly different sets of alternatives. From that point of view, it is not clear what is conveyed by an utterance of (1) by DeRose in HIGH counts as universal in the relevant sense. Thanks to Torfinn Huvenes here.

²⁹ Cf. DeRose (2002, 2009) on the ‘Generality Objection’ and Pritchard (2005).

Suppose that at the same time that DeRose and his wife are in HIGH, Peter and Mary are in a separate context, HIGH*. Like DeRose and his wife in HIGH, Peter and Mary in HIGH* have an important cheque coming due and they need to deposit their paycheques before Monday. However, not only was Mary at the bank two weeks ago on a Saturday, but the bank manager has also recently informed Mary that the bank will be open on Saturday. Her epistemic position is therefore such that she can rule out the possibility that the bank has changed its hours and is such that she does not need to engage in further inquiry.

Suppose that Peter raises the possibility that the bank has changed its hours and then asks Mary ‘Does DeRose know that the bank will be open on Saturday?’. Mary is acquainted with the character of DeRose’s evidence and his circumstances in HIGH. She replies by uttering (10).

(10) Keith knows that the bank will be open.

Just like a (true) utterance of (1) by DeRose in HIGH, Mary’s (true) utterance of (10) in HIGH* seems inappropriate and false. Furthermore, it does not seem as though Mary can cancel the pragmatic implication that would be required to explain why her utterance of (10) seems inappropriate and false. For example, an utterance of (30a) or (30b) by Mary in HIGH* seems straightforwardly infelicitous.

(30) a. Keith knows that the bank will be open, but he ought to check further
b. Keith knows that the bank will be open, but he cannot rule out that the bank has changed its hours

How should the traditional invariantist handle this apparent failure of cancellability? A parallel explanatory strategy to that proposed in the case of an utterance of (1) by DeRose in HIGH would seek to root the infelicity of utterances of (30a) and (30b) in HIGH* in the infelicity of utterances of (31a) and (31b) in HIGH*.

(31) a. The bank will be open, but Keith ought to check further

b. The bank will be open, but Keith cannot rule out that the bank has changed its hours

The problem is that utterances of either (31a) or (31b) by Mary in HIGH* do *not* seem infelicitous. The traditional invariantist therefore cannot pull a similar manoeuvre in this kind of case.

A final suggestion. Brown (2006, p. 428) proposes that it might prove more difficult to cancel a pragmatic effect in cases in which we tend to confuse what is semantically expressed by the utterance in question with what it pragmatically conveys.

However, this is not what we find in other cases of putative pragmatic phenomena apparently involving such confusions. For example, speakers can cancel the pragmatic effect commonly associated with utterances of (7) by uttering (32).

(7) Peter has eaten

(32) Peter has eaten, but not recently

The invariantist could of course claim that we exhibit a different kind of semantics-pragmatics confusion to that present in other cases of pragmatic phenomena. It might be proposed that this sort of confusion leads to the failure of cancellability we have observed. But in the absence of some independent motivation, such a proposal seems unacceptably *ad hoc*.

The traditional invariantist does not seem to have an adequate explanation for why the implication required to explain the apparent falsity of an utterance of (1) by DeRose in HIGH is not cancellable. I conclude that this failure of cancellability represents a serious objection to any traditional invariantist pragmatic explanation.

§8. Conclusion.

In this chapter, I have explored an alternative strategy for the traditional invariantist, one that roots the judgments comprising the HIGH-LOW contrast in some pragmatic effect. It was first argued that the most prominent such explanation – featuring an appeal to conversational relevance – is in important respects inadequate. It was then shown that any traditional invariantist pragmatic explanation suffers from potential problems associated with our retraction and disagreement judgments, and a particularly serious problem in regard to the issue of cancellability. (In Chapter 5, §4, I shall address a recent dispute regarding the real character of our retraction and disagreement judgments. I argue that, despite that dispute, the traditional invariantist faces a genuine problem with respect to our retraction judgments, though she may be able to avoid potential problems associated with our disagreement judgments.) In the course of the discussion, we have also seen the extent to which our judgments concerning knowledge ascriptions and denials seem to contrast with our judgments concerning even the more controversial kinds of pragmatic phenomena found in the literature.

I conclude that the arguments and considerations introduced over the course of the preceding two chapters throw serious doubt on the prospects for a traditional invariantist explanation of the HIGH-LOW contrast based around an appeal to *either* the character of our psychological processing *or* to pragmatic features of natural language. This strongly suggests – and the case will be reinforced in Chapter 5, §4 – that in regard to that contrast in our judgments, the traditional invariantist is forced to attribute significant error to speakers.

In the previous chapter (§4), it was also argued that the traditional invariantist faces problems in regard to utterances of the form ‘I know that p, but we need to inquire further’, and ‘I know that p, but I can’t rely on p in my practical reasoning’. The apparent infelicitousness of such utterances seems to suggest that knowledge closes inquiry, and that some principle like (1a) is true:

(1a) If S knows that p, then it is appropriate for S to rely on p in her practical reasoning.

In the previous chapter, it was argued that the traditional invariantist faces some pressing problems if she seeks to embrace the idea that knowledge closes inquiry, or a principle like (1a). Perhaps the most virulent problems arose in regard to judgments concerning high stakes subjects, and judgments concerning such counterfactual claims as ‘If there were more at stake, he would need to inquire further’. It was shown that attempts to invoke derivative senses of propriety, or to invoke the psychological proposals considered in the previous chapter, offered little assistance in alleviating those problems.

It is also far from straightforward to see how the traditional invariantist might invoke some aspect of the pragmatic explanation of the HIGH-LOW contrast – or some other pragmatic mechanism – to alleviate those problems. For example, it is not clear what pragmatic effect could plausibly account for the fact that, although false, an utterance, in reference to a subject like DeRose in LOW, of ‘If he had more at stake, he would need to inquire further’ seems (clearly) true. In order to avoid the attribution of significant error to speakers, it therefore seems that the traditional invariantist should reject the claim that knowledge closes inquiry, and principle (1a). Of course, this leaves her in need of an explanation for why – despite being true – utterances of the form ‘I know that p, but we need to inquire further’ and ‘I know that p, but I can’t rely on p in my practical reasoning’ typically seem infelicitous.

In the previous chapter, it was argued that the psychological proposals considered therein provided little assistance in explaining the fact that, although true, utterances of the form ‘I know that p, but we need to inquire further’ seem infelicitous. Could the traditional invariantist invoke some aspects of the pragmatic account of the HIGH-LOW contrast – or perhaps some other pragmatic mechanism – to explain such judgments?

In the course of the discussion of cancellability in the previous section, we rejected various proposals – either invoking some pragmatic effect associated with knowledge ascriptions, or some other pragmatic mechanism – for why an utterance of ‘I know that p, but we need to inquire further’ might fail to cancel the pragmatic effect allegedly associated with knowledge ascriptions (§5.2). Of course, those proposals just *are* potential explanations – either invoking some pragmatic effect associated with knowledge ascriptions, or some other pragmatic mechanism – for the fact that,

despite being true, utterances such as those of the form ‘I know that p, but we need to inquire further’ seem infelicitous. Their failure therefore suggests, in the apparent absence of some more promising proposal, that appeal to pragmatic effects will also provide little assistance to the traditional invariantist in alleviating the problems associated with rejecting the claim that knowledge closes inquiry, or principle (1a). Another apparent upshot of the argument of the previous two chapters is therefore that whether the traditional invariantist accepts *or* rejects the claim that knowledge closes inquiry (or principle (1a)), she is forced to attribute significant error to speakers.

The failure to adequately account for the HIGH-LOW contrast, or for our judgments concerning the relationships between knowledge and inquiry, and knowledge and appropriate practical reasoning, supplies us with some strong reason to reject the traditional invariantist position. In the next chapter, I discuss our judgments pertaining to the relationships between knowledge and chance (some of the relevant judgments were discussed in Chapter 1, §4.1), and our judgments pertaining to the closure of knowledge under multi-premise inference. It is argued that those judgments supply further (strong) reasons to reject the traditional invariantist position.

Chapter 3: Chance and Possibility.

We often use the language of chance and possibility to talk about our epistemic predicament. If I say to you ‘Peter might be in Athens’ or ‘There’s a chance that I left your watch at the store’, there is a natural reading of my utterances on which I am indicating something about my (or perhaps our) epistemic predicament.

This talk encourages a number of questions. What are the relationships between our talk of epistemic chance and possibility and our talk of knowledge? What are the relationships between these epistemic notions of chance and possibility and other notions of chance and possibility?

In the present chapter, I explore the impact of our talk of chance and possibility on our theorising about knowledge. The upshot will be that once again the traditional invariantist faces a number of serious problems. In subsequent chapters, I move on to evaluate the success of various non-traditional forms of invariantism in accounting for the arguments and phenomena introduced here.

§1. Initial Issues.

A first set of issues concerns the infelicity of various conjunctive claims. In cases in which the kind possibility invoked is epistemic, utterances of the form ‘I know that p, but it might be that not-p’ typically seem infelicitous. For example, an utterance of ‘I know that Peter is in his office, but he might not be’ seems infelicitous. In a similar vein, in cases in which the kind of chance invoked is epistemic, utterances of the form ‘I know that p, but there’s a chance that not-p’ typically seem infelicitous. For example, an utterance of ‘I know that Peter is in his office, but there’s a chance he isn’t’ seems infelicitous.¹

A natural explanation of the oddity of such claims is that knowledge is incompatible with this kind of chance and possibility:

¹ Cf. Chapter 1, §4.1. Such judgments have been reported by *inter alia* Hawthorne (2004a, p. 24) and Stanley (2005b).

(PK) If not-p is epistemically possible for a subject S, then S does not know that p.

(CK) If there is a (non-zero) epistemic chance, for a subject S, that not-p, then S does not know that p.

For example, consider an utterance of ‘I know that Peter is in his office, but there’s a chance he isn’t’. The speaker is asserting both that she knows that Peter is in his office and that there is an epistemic chance, for her, that Peter is not in his office.² Given principle (CK), the speaker is straightforwardly asserting a contradiction.³

§2. Epistemic Chance and Epistemic Possibility.

The principles linking chance and possibility offer a neat explanation of the infelicity of the conjunctive claims just considered. They are also intuitively attractive. However, there are grave problems associated with a traditional invariantist endorsement of (PK) and (CK). In the next two sections, I shall explore a number of these problems and argue that various traditional invariantist responses to them fail. In later sections, the prospects for an alternative explanation of the infelicity of utterances of the form ‘I know that p, but there’s a chance/possibility that not-p’ will be considered. The upshot will be that *either way*, traditional invariantism faces serious difficulties.

² In the discussion that follows, I shall largely just assume that utterances of the form ‘It might be that p’ and ‘There’s a chance that p’ semantically express (respectively) that it’s epistemically possible, for the speaker, that p, and that there’s an epistemic chance, for the speaker, that p. In a number of respects, this represents an oversimplification (and treads on relativists’ toes), but it is a harmless one for present purposes.

In recent years, there has been significant dispute over the correct semantics for expressions of epistemic possibility and chance (see, for example, Egan, Hawthorne, and Weatherson (2005), von Fintel and Gillies (2008), and MacFarlane (2011)). This dispute has chiefly been centered on how best to capture the influence, on the truth/falsity of our talk of epistemic chance and possibility, of *the individual or group* whose epistemic predicament is being indicated. This dispute is not central to the discussion to follow, which is more concerned with questions about which features of the *epistemic predicament* of the individual or group indicated in our ordinary talk of epistemic chance and possibility are relevant to the truth/falsity of that talk. (Are the relevant features, for example, their knowledge, their evidence, some proper subset of their knowledge or evidence? Does the relevant feature vary depending upon the context?)

³ See, *inter alia*, Hawthorne (2004a, esp. p. 24-28) and Stanley (2005b).

An immediate problem for the traditional invariantist seeking to endorse (CK) is that, for almost all of the beliefs we hold, it seems that there's a chance we are mistaken. To illustrate the point, suppose that Mary truly believes that Peter is in the office. Mary's belief is based on the fact that she has seen Peter's coat hanging in the office hallway and on the fact that a reliable colleague has just informed Mary that Peter is in the office. Ordinarily, we would take Mary to know that Peter is in the office. Nevertheless, it seems that there's a chance, for Mary, that Peter is not in the office. There's a chance, for Mary, that Peter has just walked out, there's a chance, for Mary, that her colleague was lying, and so on. Similarly, it seems true for Mary to utter 'There's a chance that Peter isn't in the office'.

Judgments of this sort suggest that for much of what we take ourselves to know, there is a chance that we are wrong. Given (CK), the upshot appears to be that we know very little. But that clearly contradicts one of the central tenets of traditional invariantism.

There is a more pressing variant of this problem.⁴ Suppose Peter writes down a large number of claims that he takes himself to know. Peter writes that his car is parked in the driveway, that his daughter will be in California next month, and so on. Suppose Peter occupies a moderate epistemic position with respect to each of the corresponding propositions, p_1 to p_n – he remembers leaving his car in the driveway a few hours ago, his (reliable) daughter assured him she was going to California next month, and so on – and that they are all true.

Now suppose that Peter competently deduces the conjunction of those propositions – $p_1 \& p_2 \dots \& p_n$. Given the myriad potential sources of error, it seems that there is certainly an epistemic chance, for Peter, that the conjunction is false. And utterances by Peter of 'There's a chance that $p_1 \& p_2 \dots \& p_n$ is false' or 'There's a chance that one of p_1 to p_n is false' seem clearly true.⁵ In a related vein, it seems that Peter would

⁴ Cf. Hawthorne (2004a, esp. p. 46-50 & p. 181-183). Note that Hawthorne's focus is chiefly on whether the subject knows various lengthy conjunctions, rather than (as here) on their epistemic chance/possibility.

⁵ Note that it is implausible to suggest that the (sole) source of the apparent chance, for Peter, that the conjunction is false lies in the fact that there is a chance, for Peter, that he made a mistake in the deduction of the conjunction from the conjuncts. Even if we stipulate that there

be justified in placing a high level of confidence in the negation of the conjunction. This obviously coheres with the claim that there is a (high) epistemic chance, for Peter, that the conjunction is false, and is difficult to understand on the assumption that there's no epistemic chance, for Peter, that the conjunction is false.

So it seems compelling that there is a chance, for Peter, that the conjunction is false. However, Peter occupies at least a modest epistemic position with respect to each of p_1 to p_n and those propositions are all true. The traditional invariantist therefore allows that Peter knows each proposition p_i . Let's suppose he does. Given principle (CK), it follows that for each proposition p_i , there is no chance, for Peter, that p_i is false.

The invariantist who endorses (CK) is therefore faced with a choice. She can claim that for each proposition p_i , there is no chance, for Peter, that it is false, but nevertheless claim that there *is* a chance, for Peter, that the conjunction $p_1 \& p_2 \dots \& p_n$ is false. Alternatively, she can claim that the judgment that there is a chance, for Peter, that the conjunction is false is mistaken. The problem is that neither option seems palatable. It seems as though Peter would be taking on contradictory attitudes were he to accept the proposition that there is a chance, for him, that the conjunction $p_1 \& p_2 \dots \& p_n$ is false but then reject, for each p_i , the proposition that there is a chance, for him, that p_i is false. Relatedly, it is difficult to understand where the chance that the conjunction is false could possibly issue from if not from some chance of falsehood attached to one or more of the conjuncts.⁶ Nevertheless, the alternative – viz. claiming that there is in fact no chance, for Peter, that the conjunction is false – flies in the face of the compelling thought that there is such a chance (indeed, a significant one).

I have illustrated both of the above arguments by appeal to our judgments concerning epistemic chance. These raise problems for a traditional invariantist endorsement of (CK). What about (PK)? One way to put pressure on that principle is to observe that

is no such chance, it still seems as though there is a high chance, for Peter, that the conjunction is false. For relevant discussion, see Lasonen-Aarnio (2008) and Fantl & McGrath (2009, p. 12-13).

⁶ To allay concerns rooted in the idea that a chance of error is traceable to the performance of the deduction, let it be stipulated that there is no chance, for Peter, that he made a mistake in the deduction of the conjunction from the conjuncts. Cf. fn. 5.

our notions of epistemic chance and epistemic possibility seem inextricably linked. That is, that the following principle seems true:

(PC) It is epistemically possible, for a subject S, that p if and only if there is a (non-zero) epistemic chance, for S, that p.

If further evidence for (PC) is required, just consider utterances of the form ‘There’s a chance that p, but it’s not possible that p’ and ‘It’s possible that p, but there’s no chance that p’. For example, suppose Mary utters ‘There’s a chance that Peter is in the office, but it’s not possible that Peter is in the office’ or ‘It’s possible that Peter is in the office, but there’s no chance that he is’. Insofar as Mary’s utterances are read as invoking epistemic chance and possibility, both are clearly infelicitous. This lends strong support to principle (PC).⁷

Given (PC), any argument that puts pressure on (CK) also puts pressure on (PK), and *vice versa*. It follows that the arguments given above, presented in terms of epistemic chance and placing pressure on traditional invariantist endorsement of (CK), also put pressure on traditional invariantist endorsement of (PK). But in any case – and perhaps unsurprisingly – an exactly similar pair of arguments to those given above could also be constructed by appeal to our judgments concerning epistemic possibility. (I leave this task to the reader.) These raise an even more direct challenge to a traditional invariantist endorsement of (PK).

In this section, we have considered two (related) arguments that invoke some simple judgments concerning epistemic chance, and that put pressure on traditional invariantist endorsement of (CK) and (PK). In the next section, I raise further problems for such endorsement. These problems invoke various judgments concerning the relationships between, on the one hand, epistemic chance and possibility, and, on the other hand, *objective chance*.

§2.1 Objective Chance.

⁷ Note that (PC) will plausibly have to be restricted to finite probability spaces. Such a restriction is not central to the arguments that follow.

The attempt to forge the kind of link between knowledge and epistemic chance/possibility expressed in principles (CK) and (PK) becomes even more challenging given the plausibility of various principles linking our notions of epistemic chance and possibility to the notion of objective chance. Our best physics appears to tell us that our world is irreducibly chancy – that there is a non-zero objective chance that nearly all contingent propositions about the future are false. What are the relationships between this objective notion of chance and our epistemic notions of chance and possibility? And how does this impact on the traditional invariantist position?⁸

Suppose that Peter has just dropped his pen. Peter utters ‘I know that there’s a small objective chance that the pen will fly off sideways (rather than fall to the floor), but it’s not possible that it will fly off sideways’. It is difficult to make sense of such an utterance – at least on the supposition that Peter lacks clairvoyant powers, or God-like access to the future.

In a related vein, on the supposition that Peter knows that there’s a small objective chance that his pen will fly off sideways, it seems true for him to utter ‘The pen might fly off sideways’. It also seems false for him to utter ‘It’s not the case that the pen might fly off sideways/It’s not possible that the pen will fly off sideways’. Such judgments strongly suggest that the following principle is true:

(KOP) If S knows that there’s a (non-zero) objective chance that e occurs at t, then it’s epistemically possible, for S, that e occurs at t.⁹

In addition to explaining the judgments presented above, principle (KOP) is also intuitively attractive. However, supposing that Peter is acquainted with contemporary

⁸ This section clearly owes a debt to John Hawthorne’s (2004a, 2005a) (see also Hawthorne and Lasonen-Aarnio (2009)) excellent work on the relationship between objective chance and knowledge. Though note that the focus here is rather different – being centered chiefly on the relationship between objective chance and epistemic chance/possibility, rather than on the relationship objective chance and knowledge.

⁹ Cf. Hawthorne (2004a, p. 91). Assume the domain of subjects is limited to those lacking clairvoyant powers, or God-like access to the future. This principle may also require slight adjustment in light of the issues involving reference-fixing uses of definite descriptions raised in Hawthorne & Lasonen-Aarnio (2009). Such adjustments are irrelevant here.

physics, there seems little ground for the traditional invariantist to challenge the claim that Peter knows there is a small objective chance that his pen will fly off sideways. Thus, if (KOP) is true, it seems that it must be epistemically possible, for Peter, that his pen will fly off sideways. Given principle (PK), it follows that Peter doesn't know his pen won't fly off sideways.

Our best physics seems to suggest that there is an objective chance that nearly all contingent propositions concerning the future are false. Parallel reasoning will therefore quickly lead to the conclusion that, given traditional invariantism and (PK), we – or at least those of us better acquainted with contemporary physics – have very little knowledge of the future. For the traditional invariantist, this is clearly a very unhappy result.

Can the traditional invariantist who seeks to endorse (PK) reply to such concerns? It might be proposed that the sense of possibility invoked in Peter's utterances is not epistemic in character. The obvious suggestion is rather that it is *physical* possibility – that is, possibility given the laws of nature.¹⁰ If that were so, then the judgments presented above, such as our judgment that an utterance by Peter of 'I know that there's a small objective chance that the pen will fly off sideways, but it's not possible that it will' seems infelicitous, would not lend support to principle (KOP). Rather, those judgments would lend support to principle (KOPP):

(KOPP) If S knows that there's an objective chance that e occurs at t, then it's physically possible that e occurs at t.

Unlike in the case of epistemic possibility, it is rather less plausible that – at least in general – the mere physical possibility of some event threatens our knowledge that it does not occur. For example, it's physically possible that I don't exist. But it doesn't seem like that fact alone constitutes much threat to my knowing that I do exist. Similarly, if physical possibility is being invoked, an utterance of 'It's possible that I don't exist, but I know that I do' seems true (and not infelicitous). It therefore seems that if Peter's utterance of, for example, 'I know that there's an objective chance that

¹⁰ For a similar suggestion in the case of a related problem concerning 'might'-counterfactuals, see Hawthorne (2005b).

the pen will fly off sideways, but it's not possible that it will fly off sideways' invokes only physical possibility, the problematic character of that utterance cannot be so readily worked up into a threat to traditional invariantism. The traditional invariantist who endorses (PK) could therefore suggest that given the utterances reviewed above concern only physical possibility, our judgments concerning them do not threaten her endorsement of (PK), nor do they embody any independent challenge to her position.

However, the proposal is problematic. First, if it's correct, it seems that there should be a clearly *false* reading of, for example, Peter's utterance of 'The pen might fly off sideways', one on which it invokes our ordinary sense of epistemic possibility. But there doesn't seem to be one. Compare the situation to that in which I utter 'I might not exist'. There is a reading of that utterance on which the possibility invoked is physical possibility, and on which the utterance seems true. But there is also an epistemic reading of that claim, which (almost invariably) seems false. However, this is not what we find in the case of Peter's pen – there does *not* seem to be a false reading of his utterance of 'The pen might fly off sideways', on which it invokes our ordinary sense of epistemic possibility.

Second, once Peter has uttered 'The pen might fly off sideways' it seems incorrect for him to claim to know that his pen won't fly off sideways. But if his utterance of 'The pen might fly off sideways' expressed merely physical possibility, it's not clear why this would be so. This problem is made more pressing when we observe that were I to utter 'I might not exist', thereby invoking merely physical possibility, it does *not* seem incorrect for me to then claim that I nevertheless know that I do exist.

It therefore seems that an appeal to physical possibility provides little assistance to the traditional invariantist who seeks to endorse (PK). I conclude that our judgments concerning utterances such as those of the form 'I know that there's a small objective chance that p, but it's not possible that p' lend strong support to principle (KOP), and that this puts significant additional pressure on a traditional invariantist endorsement of (PK).

There is a more pressing variant of this problem too. Consider a large number of contingent propositions about the future, q_1 to q_n , that the traditional invariantist

allows that Peter knows. Suppose that for each proposition q_i , there is a small objective chance that it is false. On the assumption that n is sufficiently large, the conjunction $q_1 \& q_2 \dots \& q_n$ will therefore have a very high objective chance of being false.¹¹ As before, there seems no obstacle (on the traditional invariantist account) to Peter knowing the deliverances of our best physics, so let's also suppose that Peter knows that there is a very high objective chance that the conjunction is false.

Now consider an utterance by Peter of 'I know that there's a very high objective chance that the conjunction $q_1 \& q_2 \dots \& q_n$ is false, but it's not possible that it's false'. On the assumption that Peter lacks God-like access to the future, Peter's utterance is very difficult to make sense of. In a similar vein, given Peter knows that there is a high objective chance that the conjunction is false, an utterance by him of ' $q_1 \& q_2 \dots \& q_n$ might be false' would seem true. Finally, given Peter knows there is a very high objective chance that the conjunction is false, it seems that Peter would be justified in placing a high degree of confidence in the falsity of the conjunction. This coheres with the claim that it's epistemically possible, for Peter, that the conjunction is false but is difficult to understand on the assumption that it is epistemically impossible, for Peter, that the conjunction is false.

These judgments suggest that the following, highly plausible, principle is true:

(KHOP) If S knows that there's a high objective chance that p , then it's epistemically possible, for S , that p .¹²

(KHOP) delivers the result that it's epistemically possible, for Peter, that the conjunction $q_1 \& q_2 \dots \& q_n$ is false. Is this a problem for the traditional invariantist? The traditional invariantist allows that Peter knows each of q_1 to q_n . Suppose he has that knowledge. Given (PK), it follows that, for each proposition q_i , it is not

¹¹ Two assumptions: (a) It should be assumed that q_1 to q_n are probabilistically independent. I shall leave this assumption implicit in subsequent discussion; (b) I also assume that it makes sense to talk of the objective chance of a proposition. If necessary, the discussion could readily be recast in terms of the objective chance of various events.

¹² Let the domain of subjects be limited to those without clairvoyant powers or God-like access to the future. The principle may also require subtle emendations in light of the issues associated with referential uses of definite descriptions raised in Hawthorne & Lasonen-Aarnio (2009). Such emendations are irrelevant to the matters at hand.

epistemically possible, for Peter, that q_i is false. In the previous section, we saw that maintaining that there is an epistemic chance, for a subject, that a conjunction is false, but no chance, for that subject, that any given conjunct is false, seems very problematic. A similar result holds for epistemic possibility. It seems that Peter would be taking on contradictory attitudes were he to accept the proposition that it's epistemically possible, for him, that the conjunction $q_1 \& q_2 \dots \& q_n$ is false but reject, for each q_i , the proposition that it's possible, for him, that q_i is false.¹³ It seems that the traditional invariantist who endorses (PK) must therefore maintain that it is not epistemically possible, for Peter, that the conjunction is false.¹⁴ She must therefore reject the compelling connection expressed in (KHOP). This is surely another, significant cost.

Let me present one final – perhaps more theoretically grounded – objection, this time to a traditional invariantist endorsement of principle (CK).¹⁵ There are some plausible relationships between objective chance and rational credence or confidence. These also appear to have consequences for our notion of epistemic chance. Here's a quote from David Lewis:

Whatever makes it true that the chance of decay is 50% must also, if known, make it rational to believe to degree 50% that the decay will occur. (Lewis 1994, p. 476)

Perhaps even more compelling: if S knows that the objective chance that e occurs at t is greater than 0, then S's credence that e will occur at t should be greater than 0. This certainly seems like a compelling principle.¹⁶ Let's apply the principle to the case of Peter and his pen. Suppose Peter knows that the objective chance that his pen will fly off sideways is greater than 0. It follows that Peter's credence that the pen will fly off sideways should be greater than 0.

¹³ Ignoring the possibility of error traceable to the performance of the deduction of the conjunction from the conjuncts. Cf. fn. 5.

¹⁴ An alternative argument for this conclusion would invoke our judgments about epistemic chance and principle (PC), tying epistemic possibility and epistemic chance.

¹⁵ This argument differs somewhat from, but was clearly inspired by, an argument presented in Hawthorne (2004a, p. 91-94).

¹⁶ Once again, on the assumption that the subject does not have special access to the future.

What role does our notion of epistemic chance play in our cognitive economy? One central role for our notion of epistemic chance arguably lies in its relationship to rational credence. One very plausible facet of that relationship is captured in the following principle:

(CC) If it is not rational for S to assign a credence of 0 to p, then the epistemic chance, for S, that p is greater than 0.

This principle is intuitively very attractive. It seems that if it's the case that you ought not to be fully confident that some proposition is false, that must be because there's a chance, for you, that it's true.¹⁷ Peter knows that the objective chance that the pen will fly off sideways is greater than 0, so Peter's credence that the pen will fly off sideways should be greater than 0. Assuming (CC), it follows that the epistemic chance, for Peter, that the pen will fly off sideways is greater than 0.

This clearly has repercussions for the traditional invariantist who seeks to endorse (CK). Given that principle, it will obviously follow that Peter does not know that the pen will not fly off sideways. This argument will quickly generalise to yield the conclusion that we – or at least those of us acquainted with contemporary physics – know very little about the future. In order to maintain (CK), the traditional invariantist must therefore *either* deny the compelling idea that if a subject knows that the objective chance of some event e occurring is greater than 0, she should have a credence that e will occur that is greater than 0, *or* she must deny the compelling relationship between rational credence and epistemic chance captured by (CC). Neither option is attractive.

As remarked at the close of the previous section, problems for a traditional invariantist endorsement of (CK) also represent problems for a traditional invariantist endorsement of (PK), and *vice versa*. The reason is that the following principle seems to hold:

¹⁷ The sense of rationality in play here is epistemic, rather than practical, in nature. Suppose that Peter is threatened with a horrible death if he places full confidence in the truth of some proposition p. In that case, even supposing that there's no chance, for Peter, that not-p, he nonetheless ought not to place full confidence in p. This is clearly not the kind of rationality at play in (CC).

(PC) It is epistemically possible, for a subject S, that p if and only if there is a (non-zero) epistemic chance, for S, that p.

Given (PC), it follows that the final objection of this section, centred around an appeal to rational credence/confidence, also constitutes an objection to traditional invariantist endorsement of (PK), and that the first two objections of this section also represent objections to traditional invariantist endorsement of (CK).

Over the course of the last two sections (§2 and §2.1), we have seen a variety of arguments that lend strong support to the conclusion that for much of what we take ourselves to know, there is an epistemic chance/possibility that we are wrong. These arguments make significant trouble for a traditional invariantist endorsement of principles (CK) and (PK). In the following section, I consider and reject some possible responses to these arguments.

§3. Chance Contextualism and Other Replies.

How can the traditional invariantist who wishes to endorse (CK) and (PK) respond to the objections raised in §2 and §2.1? One general strategy would exploit the idea that considering questions about epistemic chance and possibility tends to make salient certain possibilities of error. It might be proposed that this has a distorting effect on our judgments about epistemic chance and possibility.

Recall the case of Mary, discussed in §2. Mary truly believes that Peter is in the office. Mary's belief is based on the fact that she saw Peter's coat hanging in the office hallway, and on the fact that a reliable colleague just told her that Peter is in the office. On the assumption that traditional invariantism is true, Mary knows that Peter is not in the office. Given (CK), it follows that there is no chance, for Mary, that Peter is in not in the office. The problem was that there *does* seem to be a chance, for Mary, that Peter is not in the office. Similarly, it seems true for Mary to utter 'There's a chance that Peter isn't in the office' and false for Mary to utter 'There's no chance that Peter isn't in the office'.

Notice that when we address the question of whether there is a chance that we are wrong about some proposition, certain error possibilities typically become salient. In addressing the question of whether there is a chance, for her, that Peter is not in the office, relatively distant or unlikely possibilities, like the possibility in which Peter has just walked out of the office, will plausibly become salient to Mary. In the previous two chapters, we explored traditional invariantist proposals to the effect that the salience of such error possibilities leads, for either pragmatic or psychological reasons, to mistaken judgments about knowledge. If something like principle (CK)* is true, those proposals might be extended to explain our judgments concerning epistemic chance.

(CK)* There is a (non-zero) epistemic chance, for S, that not-p if and only if S does not know that p.¹⁸

Let's consider the pragmatic account first. Given (CK)*, an utterance by Mary of 'There's a chance Peter is not in the office' is true just in case Mary does not know that Peter is in the office.¹⁹ It is therefore reasonable to expect that the pragmatic effects associated with an utterance by Mary of 'There's a chance that Peter is not in the office' will be similar to those associated with an utterance by Mary of 'I don't know that Peter is in the office'.

Let's assume that by considering the question of whether there's a chance, for her, that Peter is not in the office, the possibility that Peter has just walked out of the office becomes salient to Mary. On the proposals considered in Chapter 2, in contexts in which that possibility becomes salient, a false utterance by Mary of 'I don't know that Peter is in the office' would typically pragmatically convey that Mary cannot rule out the (relatively distant) possibility that Peter has just left the office (or something along those lines). Mary's epistemic position might seem to be such that she cannot

¹⁸ An alternative to (CK)* is a *compatibility* principle: (CK)** There is a (non-zero) epistemic chance, for S, that not-p if and only if p is compatible with what S knows.

¹⁹ Recall that I am making the simplifying assumption that utterances of the form 'There's a chance that p' and 'It might be that p' express that there's an epistemic chance, for the speaker, that p, and that it's epistemically possible, for the speaker, that p, respectively. See fn. 2.

rule out such a (relatively distant) possibility, so her utterance would convey something true.

It might therefore be proposed that the reason Mary's false utterance of 'There's a chance that Peter is not in the office' seems true is because it pragmatically conveys the truth that Mary cannot rule out the possibility that Peter has just walked out of the office. Conversely, it might be proposed that the reason Mary's true utterance of 'There's no chance that Peter is not in the office' seems false is because it pragmatically conveys the falsehood that Peter *can* rule out the possibility that Peter has just walked out of the office.

Alternatively, the traditional invariantist might propose a psychological account. In particular, she might propose that the kind of psychological processes that lead us, when distant or unlikely error possibilities are raised, to (falsely) judge that Mary does not know that Peter is in the office, also lead us – given (CK)* – to (falsely) judge that there's a chance, for Mary, that Peter is not in the office.

Either the pragmatic or the psychological story could then be extended to cover our judgments concerning epistemic possibility, given something in the neighbourhood of the following principle:

(PK)* It is epistemically possible, for a subject S, that not-p if and only if S does not know that p.²⁰

Are these responses likely to be successful? It is not clear whether these strategies can be suitably extended to provide a response to all of the problems highlighted with the traditional invariantist endorsement of (CK) and (PK). But for the sake of argument, let's just assume that they can. In Chapters 1 and 2, I argued that there are serious difficulties associated with these kinds of traditional invariantist proposals when applied to our judgments concerning knowledge ascriptions, such as those embodied by the HIGH-LOW contrast. Perhaps unsurprisingly, those same difficulties extend to

²⁰ An alternative is another compatibility principle: (PK)** It is epistemically possible, for S, that p if and only if p is compatible with what S knows. Cf. Hawthorne (2004a, p. 26) and Stanley (2005b).

the kind of pragmatic and psychological account of our talk of epistemic chance and possibility being considered here. For example, one serious difficulty encountered with the pragmatic proposal in the case of knowledge ascriptions was that the relevant effect does not seem to be cancellable. This objection has a simple analogue to the case at hand. Mary does not seem able to cancel the pragmatic effect allegedly associated with an utterance of ‘There’s no chance that Peter is in the office’. For instance, an utterance by Mary of ‘There’s no chance that Peter is in the office, but I cannot rule out that Peter has just walked out of the office’ does not seem felicitous. The situation is much the same for the other objections pushed in Chapters 1 and 2. The weight of those objections suggest that the traditional invariantist would be unwise to rely on such manoeuvring here.

However, Hill and Schechter (2007) offer the traditional invariantist an interesting alternative. On the most natural reading of their proposal, they offer a way for the traditional invariantist to endorse principles (CK) and (PK) and *also* endorse the arguments and judgments that seem to support the conclusion that, for much of what we know, there is both a chance and possibility that we are wrong. The grounding idea is that there is some kind of ambiguity or context sensitivity in our epistemic talk of ‘chance’ and ‘possibility’.

In some circumstances, they suggest our ‘chance’ talk expresses chance *on our evidence* and our ‘possibility’ talk expresses *compatibility with our evidence*. This is what happens when we judge, for example, that there’s a chance, for Mary, that Peter is not in the office (see §2, above). On their account, we are correctly judging that there is a chance, on Mary’s evidence, that Peter is not in the office. Similarly, the reason an utterance by Mary of ‘There’s a chance that Peter isn’t in the office’ seems true is because she is expressing the truth that there’s a chance, on her evidence, that Peter is not in the office.

The account may also be extended to explain why Peter’s utterance of ‘I know that there’s an objective chance that the pen will fly off sideways, but it’s not possible that it will fly off sideways’ seems infelicitous (see §2.1, above). Suppose Peter is expressing that he knows that there is an objective chance that the pen will fly off sideways but that it is not compatible with his evidence that the pen flies off

sideways. Hill and Schechter might propose that the following connection between objective chance and evidence holds:

(KOEV) If S knows that there is a (non-zero) objective chance that e occurs at t, then it is compatible with S's evidence that e occurs at t.²¹

If principle (KOEV) is true, and Peter is expressing that it is compatible with his evidence that his pen will fly off sideways in the second conjunct of 'I know that there's an objective chance that the pen will fly off sideways, but it's not possible that it will fly off sideways' then it is clear why that utterance would seem infelicitous – it expresses a contradiction. A similar strategy could be deployed to account for the other arguments and phenomena presented in §2 and §2.1.

But what about the infelicity of utterances of the form 'I know that p, but there's a chance that not-p' and 'I know that p, but it might be that not-p' discussed in §1, above? Hill and Schechter suggest that in these kinds of circumstances our 'chance' talk expresses chance *on our knowledge* and our 'possibility' talk expresses *compatibility with our knowledge*. The reason we judge, for example, that an utterance by Mary of 'I know that Peter is in the office, but he might not be' seems infelicitous, is because Mary is expressing the conjunction that she knows that Peter is in the office and that it is compatible with what she knows that Peter is not in the office. In a similar vein, insofar as 'epistemic chance' and 'epistemic possibility' are understood to be capturing the sense of epistemic chance and possibility expressed in cases like these, the traditional invariantist can endorse principles (CK) and (PK).²²

Neat though this solution may seem, it suffers from a variety of problems. First, given there is an ambiguity or context-sensitivity in our ordinary talk of 'chance' and 'possibility', why isn't there an obviously true reading of 'I know that Peter is in the office, but there's a chance he isn't' – one on which it expresses chance on Mary's evidence? Similarly, why isn't there an obviously true reading of 'I know that Peter is

²¹ Once again, this principle would probably require subtle emendations, such as a limitation to those subjects lacking clairvoyant powers. Such emendations won't concern us here.

²² This aspect of the proposal can also be extended to cover the phenomena that will be introduced in §4.

in the office, but he might not be' – one on which it expresses compatibility with Mary's evidence?

Second, the view seems unable to account for the effect of our admission that there's a chance or possibility that we're wrong on our subsequent claims to know. The reason an utterance by Mary of 'There is a chance that Peter is not in the office' seems true is, we are told, because she is expressing the truth that there is a chance on her evidence that Peter is not in the office. But notice that once Mary has uttered 'There's a chance that Peter is not in the office' it seems very uncomfortable for her to then claim to know that Peter is in the office (as observed by Hawthorne, 2004a, p. 24). But why? Given that Mary's 'chance' claim expressed chance on her evidence, the truth of her claim is *perfectly compatible* with her knowing that Peter is in the office.

Third – and perhaps most pressing – even when the speaker tries to be explicit about the notion of chance or possibility she is considering, the relevant conjunction still seems infelicitous. For example, an utterance by Mary of 'I know that Peter is in the office, but it's compatible with my evidence that he isn't' still seems infelicitous. Indeed, it was just this sort of utterance that Lewis (1996) thought was so troublesome:

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 possibilities of error, just *sounds* contradictory. (Lewis, 1996, p. 549)

For Lewis, an uneliminated possibility is stipulated to be one that is compatible with the subject's *evidence*.

A possible response to these concerns is for the traditional invariantist to maintain that when we consider a chance claim and a knowledge claim together, our uses of 'chance' invariably pick out chance on our knowledge. (A similar proposal can be tabled for the relevant uses of 'might'.) If that were so, it's clear why there would not seem to be a true reading of Mary's utterance of 'I know that Peter is in the office, but

there's a chance he isn't'. However, the proposal is suspiciously *ad hoc*. Furthermore, whilst it is unclear whether or not the proposal will assist in responding to the second concern, it certainly seems helpless in addressing the third. It seems deeply implausible that Mary is really expressing compatibility with her knowledge when she utters 'I know that Peter is in the office, but it's compatible with my evidence that he isn't'.²³

Finally, the response actually threatens to undercut some of the central explanatory advantages of the Hill and Schechter view. One of those advantages is that it explains why we judge that, for much of what we ordinarily take ourselves to know, there is a chance we are mistaken. But when we judge that, for much of what we ordinarily take ourselves to know, there is a chance we are mistaken, we are plausibly considering both knowledge claims and chance claims. Our 'chance' talk should therefore pick out chance on our knowledge. However, if that were so, then presumably we would *not* judge that for much of what we ordinarily take ourselves to know, there is a chance we are mistaken.

These considerations throw serious doubt on the prospects of successfully responding to the problems brought out in the previous section by postulating contextualism about (or ambiguity in) our ordinary talk of 'chance' and 'possibility'. In the apparent absence of any other adequate response to those problems, it seems that the traditional invariantist would be wise to reject (CK) and (PK).

Let's therefore suppose that the traditional invariantist rejects (CK) and (PK). In order to safeguard her position, she will require an explanation of the infelicity of utterances of the form 'I know that p, but there's a chance that not-p' and 'I know that p, but it might be that not-p' that does not involve an appeal to those principles. In the next section, I argue that the most promising attempts to carry out that explanatory task fail.

§4. Rejecting (CK) and (PK).

²³ Note that Schechter and Hill reject the kind of equivalence between knowledge and evidence defended in Williamson (2000, ch. 9).

Given the apparent severity of the problems for the traditional invariantist who wishes to accommodate principles (PK) and (CK), it is worthwhile considering the possibility that there is an alternative explanation for the infelicity of utterances of the form ‘I know that p, but there’s a chance that not-p’ and ‘I know that p, but it might be that not-p’. This strategy may garner some support from putative counterexamples to principles like (PK) and (CK). For example, utterances of the following conjunctions seem felicitous.²⁴

- (1) a. I know they’re going to lose, but I’m going to carry on watching just in case.
- b. I know it’s not going to rain, but I’m going to take an umbrella just in case.
- c. There’s a chance that all the golfers get a hole-in-one on the thirteenth, but we know that won’t happen.

In Chapter 1, we explored not only the viability of various psychological explanations of our judgments to the effect that we don’t know various propositions, but also the viability of a (traditional invariantist) psychological explanation of the infelicity of utterances of the form ‘I know that p, but there’s a chance that not-p’. Those explanations were found to be inadequate (see Chapter 1, §4.1). There is, however, an important alternative.

Dougherty and Rysiew (2009, 2011) provide the traditional invariantist with a *pragmatic* account of the infelicity of utterances of the form ‘I know that p, but it might be that not-p’ and ‘I know that p, but there’s a chance that not-p’.²⁵ For present purposes, I shall just focus on utterances of the form ‘I know that p, but there’s a chance that not-p’. Dougherty and Rysiew suggest that our epistemic uses of ‘chance’ express chance *on our evidence*. In addition, they claim that for the vast majority of those propositions that we know, there is a small chance on our evidence that they are false. This enables them to readily accommodate the arguments and judgments

²⁴ (1a) and (1b) are from Hawthorne (2004a, p. 24). Hawthorne credits the examples to Brian Weatherson and Tamar Gendler. (1c) is inspired by a case in Vogel (1999, p. 165).

²⁵ See also Rysiew (2001).

presented in §2 and §2.1 that seemed to suggest that, for much of what we take ourselves to know, there's a chance that we are wrong. On their account, an utterance by Mary of 'I know that Peter is in the office, but there's a chance he isn't' comes out true. Their proposal is that the infelicity of such utterances has a pragmatic source.

The proposal is centered around an appeal to the notion of conversational relevance. It is suggested that the reason utterances of the form 'I know that p, but there's a chance that not-p' seem infelicitous is because (a) we assume that a speaker would only mention the chance, on their evidence, that not-p if that chance were relevant to the conversation, and (b) if that chance were sufficiently large to be relevant to the conversation, then it would be sufficiently large to preclude knowledge. For example, consider an utterance by Mary of 'I know that Peter is in the office, but there's a chance that he isn't in the office'. On Dougherty and Rysiew's proposal, Mary's utterance of 'there's a chance that he isn't in the office' would only be relevant to the conversation if her utterance of 'I know that Peter is in the office' were false. That accounts for why her utterance of the conjunction seems infelicitous.

Is this proposal viable? A first objection stems from the observation that, in a variety of circumstances, it seems that a chance can be sufficiently large to be relevant to the conversation *without* being sufficiently large to preclude knowledge (Fantl & McGrath, 2009, p. 22-23). For example, suppose a colleague asks Mary whether or not Peter is in the office, and stresses that the matter is extremely important. In such circumstances, even a very small chance on Mary's evidence that Peter is not in the office may well be relevant to the conversation. It therefore seems that an utterance by Mary of 'there's a chance he's not in the office' could be relevant without the chance on her evidence being large enough to preclude knowledge. The problem is that even in such 'high-stakes' circumstances, an utterance by Mary of 'I know that Peter is in the office, but there's a chance he isn't' seems infelicitous.²⁶

It is possible to expand the pragmatic account to respond to this species of objection. To do so, the traditional invariantist first needs to suggest that the epistemic

²⁶ For another example to the same effect, envisage an utterance by DeRose of 'I know that the bank's hours have not changed, but there's a chance they have' in the HIGH context of the Bank Case. Such an utterance seems infelicitous.

requirements on proper assertion are context variable. In particular, she needs to suggest that in high stakes situations, speakers need to occupy stronger epistemic positions in order for their assertions to be appropriate. This proposal is lent support by the observation that, in the LOW context of the Bank Case, it seems appropriate for DeRose to assert that the bank will be open, whilst in the HIGH context, a similar assertion does not seem appropriate.²⁷ A natural explanation for this particular contrast between HIGH and LOW is that the epistemic requirements on appropriate assertion are influenced by how much is at stake.

Dougherty and Rysiew could then propose the following. In *all* circumstances, Mary's utterance of 'there's a chance he isn't' will be relevant only if the chance, on Mary's evidence, that Peter isn't in the office is sufficiently large to be relevant to the conversation. In circumstances in which *little is at stake*, the chance on Mary's evidence that Peter is not in the office will be sufficiently large to be relevant to the conversation only if it is sufficiently large to preclude knowledge. In circumstances in which little is at stake, it is therefore clear why an utterance by Mary of 'I know that Peter is in the office, but there's a chance he isn't' would seem infelicitous.

On the other hand, in circumstances in which *a great deal is at stake*, the chance on Mary's evidence that Peter is in the office can be sufficiently large to be relevant to the conversation even supposing it's very small (and so insufficiently large to preclude knowledge). However, in such circumstances, the conditions on proper assertion will be such that even a very small chance, on Mary's evidence, that Peter is not in the office will preclude her from being able to appropriately assert that Peter is in the office. It follows that even a very small chance, on her evidence, that Peter is not in the office will preclude Mary from being able to appropriately assert the logically stronger proposition that she knows that Peter is in the office.²⁸ In short, in circumstances in which a great deal is at stake, Mary's utterance of 'I know that Peter is in the office, but there's a chance he isn't' seems infelicitous because in order for her utterance of the second conjunct to be relevant, she must fail to meet the standards for appropriate assertion of the first conjunct.

²⁷ See DeRose (2002, p. 177; 2009, p. 90).

²⁸ I assume that knowledge is factive – that is, that (necessarily) if S knows p, then p.

However, even with these emendations, the proposed explanation for our judgments to the effect that utterances of the form ‘I know that p, but there’s a chance that not-p’ seem infelicitous suffers from a number of serious problems. First, it seems like the (original) pragmatic effect proposed to explain the infelicity of uttering the conjunctive claim in low stakes scenarios should be *cancellable*. The proposed pragmatic effect is a straightforward relevance-based pragmatic effect. As Grice himself noted, in the case of relevance-based pragmatic effects, one can make it clear that one does not intend one’s contribution to be relevant, and thereby cancel the pragmatic effect at issue.²⁹ It should therefore seem felicitous for Mary to utter something along the lines of ‘The chance that Peter is not the office is not relevant right now, but bear in mind that even though I know that he’s in the office, there’s a chance he’s not’. However, such an utterance does not seem felicitous.

In a related vein, Dodd (2010) notes that such utterances as ‘Strictly speaking, I know that Peter is the office, but there’s a chance he’s not’ seem infelicitous. One central function of ‘strictly speaking’ seems to be to encourage our focus on to what is semantically expressed by the embedded clause – on what is ‘literally said’. Given that what is semantically expressed by Mary’s utterance of ‘I know that Peter is in the office, but there’s a chance he isn’t’ is – on the view under discussion – straightforwardly true, it therefore seems that there should be a straightforwardly true reading of her utterance of ‘Strictly speaking, I know that Peter is in the office but there’s a chance he isn’t’. However, there doesn’t seem to be one.

Second, there are a variety of utterances in the vicinity of utterances of the form ‘I know that p, but there’s a chance that not-p’ that would seem to come out false on the account under discussion, but that we judge to be true. For example, it seems true for Mary to utter ‘If there’s a chance that Peter is not in the office, then I don’t know that he is’ (Hawthorne, 2004a, p. 25). However, on the account under discussion, it is perfectly compatible with the Mary knowing that Peter is in the office that there is an epistemic chance, for Mary, that Peter is not in the office. The apparent truth of such conditionals is therefore somewhat puzzling.

²⁹ Grice (1989, p. 39). Grice actually makes a more general claim, one not limited to his maxim of Relation.

It may be possible to extend the original pragmatic account to explain the apparent truth of these conditional utterances (Dougherty and Rysiew, 2009, p. 129-130). It might be proposed that an utterance by Mary of ‘If there’s a chance that Peter is not in the office, then I don’t know that Peter is in the office’ pragmatically conveys that if there’s a *relevant* chance on Mary’s evidence that Peter is not in the office, then Mary doesn’t know that Peter is in the office. This may be able to explain why the conditional seems true as uttered in low stakes circumstances. However, it does not seem that it can be extended to explain why the conditional would seem true as uttered in high stakes circumstances. In high stakes circumstances, the chance on Mary’s evidence that Peter is not in the office can be sufficiently large to be relevant without being large enough to preclude knowledge.

Third, one useful test for determining whether the infelicity of asserting a conjunction is rooted in some pragmatic effect associated with asserting both conjuncts – rather than being rooted in the fact that the conjuncts cannot both be true – is to embed the conjunction under ‘Suppose’.³⁰ For example, utterances of ‘Dogs bark, but I don’t know that they do’ typically seem infelicitous. Nevertheless, utterances of ‘Suppose dogs bark but I don’t know that they do’ are entirely comfortable. This suggests that the two conjuncts are consistent, and that the explanation for why utterances of ‘Dogs bark but I don’t know that they do’ seem infelicitous is broadly pragmatic in nature. One natural proposal is that knowledge is a condition on appropriate assertion.

Let’s therefore consider an utterance by Mary of ‘Suppose that I know that Peter is in the office, but there’s a chance that he isn’t’. That utterance does not seem felicitous (insofar as the chance in question is held to be epistemic chance for the speaker). This suggests that an utterance by Mary of ‘I know that Peter is in the office, but there’s a chance he isn’t’ expresses a genuine contradiction, and that a pragmatic explanation of its infelicity is therefore *not* appropriate.

In reply to this concern, Dougherty and Rysiew might try and argue that what we might call ‘The ‘Suppose’ Test for Contradiction’ delivers some potentially pertinent false positives. In particular, notice that an utterance by Mary of the more simple

³⁰ Thanks to Paolo Santorio.

‘Suppose that Peter is in the office, but there’s a chance he isn’t’ is not altogether comfortable *either*.³¹ In this case, it does not seem especially plausible that the reason Mary’s utterance seems uncomfortable is because the relevant supposition cannot be consistently made. On any standard view of epistemic chance claims, an utterance by Mary of ‘there’s a chance Peter isn’t in the office’ is true just in case there is a certain kind of imperfection in Mary’s epistemic predicament. For example, the imperfection could be that Mary does not know that Peter is in the office, or it could be that it is consistent with Mary’s evidence that Peter is not in the office, or It seems rather implausible to claim that, necessarily, Mary’s epistemic predicament can exhibit the relevant imperfection only if Peter is in the office. Certainly that is not the case on the two leading accounts of the relevant imperfection – that it is a lack of knowledge or a lack of entailing evidence. It therefore seems that an utterance by Mary of ‘Peter is in the office, but there’s a chance he isn’t’ does *not* express a contradiction. In consequence, it seems plausible that in this case, The ‘Suppose’ Test for Contradiction delivers a false positive, and that there is some broadly pragmatic explanation of the uncomfortableness of an utterance by Mary of ‘Suppose that Peter is in the office, but there’s a chance he isn’t’.³²

These observations are potentially very important for the case at hand. Assuming that Mary knows that Peter is in the office entails that Peter is in the office, it might well be possible to extend the pragmatic explanation of the infelicity of ‘Suppose that Peter is in the office, but there’s a chance he isn’t’ to account for the infelicity of ‘Suppose that I know that Peter is in the office, but there’s a chance he isn’t’. The traditional invariantist could therefore account for the infelicity of the ‘Suppose’ claim by *other means* than being forced to postulate that the knowledge claim and the epistemic chance claim are inconsistent.

³¹ Cf. Yalcin (2007).

³² Dorr and Hawthorne (ms.) attempt such an explanation.

Note that Yalcin (2007) actually suggests that the infelicity of such ‘Suppose’ claims implies that utterances of the form ‘p but there’s a chance that not-p’ express contradictions (though his focus is mainly on claims of the form ‘p but it might be that not-p’). His proposal is therefore that the result is not a false positive for The ‘Suppose’ Test after all. Clearly, if Yalcin is right, then Dougherty and Rysiew cannot pursue the line of response suggested in the main text.

Let us just grant that this is so. It does not provide much solace to the traditional invariantist. The reason is that there are other pertinent examples in which we witness failures of The ‘Suppose’ Test but in which the traditional invariantist cannot attempt a similar manoeuvre. For instance, not only does an utterance by Mary of ‘I know that Peter is in the office, but there’s a chance he isn’t’ seem infelicitous, an utterance by some third-party, call him ‘Kurt’, of ‘Mary knows that Peter is in the office, but there’s a chance, for her, that he isn’t’ also seems infelicitous.³³ On Dougherty and Rysiew’s view, epistemic chance is equivalent to chance on the subject’s evidence. It is therefore the case, on their view, both that Mary knows that Peter is in the office and that there is an epistemic chance, for her, that Peter isn’t in the office.³⁴ An utterance by Kurt of ‘Mary knows that Peter is in the office, but there’s a chance, for her, that he isn’t’ will therefore come out true.

Presumably, Dougherty and Rysiew will seek to extend their relevance-based pragmatic account to account for the infelicity of such utterances. Let’s therefore apply The ‘Suppose’ Test in this case. Consider an utterance by Kurt of ‘Suppose Mary knows that Peter is in the office, but there’s a chance, for her, that he isn’t’. Such an utterance seems infelicitous. However, observe that an utterance by Kurt of ‘Suppose that Peter is in the office, but there’s a chance, for Mary, that he isn’t’ seems entirely *felicitous*. The infelicity of ‘Suppose Mary knows that Peter is in the office, but there’s a chance, for her, that he isn’t’ cannot therefore be rooted in the infelicity of ‘Suppose that Peter is in the office, but there’s a chance, for Mary, that he isn’t’, in the way in which the infelicity of ‘Suppose I know that Peter is in the office, but there’s a chance he isn’t’ might be rooted in the infelicity of ‘Suppose that Peter is in the office, but there’s a chance he isn’t’.

In the absence of any *other* pragmatic explanation for the infelicity of an utterance by Kurt of ‘Suppose Mary knows that Peter is in the office, but there’s a chance, for her, that he isn’t’, there is strong pressure towards the conclusion that such an utterance expresses a genuine contradiction.³⁵ In other words, there is strong pressure towards

³³ See, for example, Stanley (2005b).

³⁴ Note that Dougherty and Rysiew are assuming that an equation like $E=K$ (Evidence = Knowledge) is false. $E=K$ is defended in Williamson (2000, ch 9).

³⁵ Dougherty and Rysiew might suggest that the reason the ‘Suppose’ claim seems infelicitous in this case is because it pragmatically conveys that the hearer should suppose that Mary

the conclusion that principle (CK) is true, and that consequently the kind of pragmatic explanation that Dougherty and Rysiew are proposing is not appropriate.

Dougherty and Rysiew's pragmatic explanation appears to face a number of serious difficulties. It is also worth stressing that the objections pressed here represent challenges to *any* pragmatic account of the infelicity of utterances of the form 'I know that p, but there's a chance that not-p' and 'I know that p, but it might be that not-p'. So where does the traditional invariantist stand? Over the course of this section, I have argued that explanations of the infelicity of utterances of the form 'I know that p, but there's a chance that not-p' and 'I know that p, but it might be that not-p', which do not rest on an endorsement of (CK) and (PK), are unsuccessful. This suggests that the traditional invariantist must accept (CK) and (PK). However, in the opening sections of the chapter, we saw that the traditional invariantist encounters extremely serious problems if she endorses those principles. Thus: *either* in regard to the various judgments suggesting that, for much of what we take ourselves to know, there's a chance/possibility that we are wrong, *or* in regard to our judgments concerning the relationship between knowledge and epistemic chance/possibility, such as that utterances of the form 'I know that p, but there's a chance that not-p' and 'I know that p, but it might be that not-p' seem infelicitous, the traditional invariantist appears forced to attribute significant error to speakers.

In the remainder of this chapter, I explore some aspects of the relationship between knowledge and deductive inference. The discussion will reveal further difficulties for the traditional invariantist. It will also provide an opportunity to explore the relationship, for the traditional invariantist, between endorsing principles like (CK) and (PK) and the preservation (or otherwise) of knowledge by deductive inference.

§5. Multi-Premise Closure.

knows that Peter is in the office and that there's a *relevant* epistemic chance, for Mary, that he isn't. This *may* succeed in explaining the infelicity of the 'Suppose' claim as uttered in low stakes circumstances, but for reasons already introduced, it will certainly not assist in explaining the infelicity of the 'Suppose' claim as uttered in high stakes circumstances.

What is the relationship between knowledge and deductive inference? A natural proposal is that competent deduction *extends* knowledge. If one knows the premises of a valid argument, and competently deduces the conclusion of that argument from the premises, then one also knows the conclusion. This is indeed a compelling thought.

(MPC) If a subject S knows propositions p_1 to p_n and competently deduces q from p_1 to p_n – retaining throughout her knowledge of p_1 to p_n – then S knows q .³⁶

Can the traditional invariantist endorse this principle? In what follows, I shall argue that whether or not the traditional invariantist endorses the principle, she faces serious difficulties.

Suppose Peter writes down a large number of (contingent) claims that he takes himself to know about the future. He writes that his daughter will be in California next month, that the temperature in London tomorrow won't exceed forty degrees centigrade, and so on. Suppose that Peter occupies a moderate epistemic position with respect to each of the corresponding propositions, r_1 to r_n – his (reliable) daughter told him that she would be in California next month, he knows that temperatures above forty degrees in London are very rare, and so on – and that they are all true. It seems that the traditional invariantist must allow that Peter knows each of r_1 to r_n .

Now suppose Peter competently deduces the conjunction of those propositions, r_1 & r_2 ... & r_n . Given (MPC), Peter knows that conjunction. But that seems absurd. Given the sheer number of beliefs – and myriad potential sources of error – it seems that Peter does not know the conjunction. In a related vein, it seems that Peter is justified in being confident that the conjunction is *false*. But if one is justified in being confident that a proposition is false, it hardly seems that one can know it to be true.

³⁶ See Williamson (2000, p. 117) & Hawthorne (2004a, p. 33). Providing precise formulations of closure principles is notoriously difficult. Even if something in the vicinity of (MPC) is correct, it is doubtful that it will look in all details exactly like the principle stated. (See Lasonen-Aarnio (2008) for some concerns with this particular kind of 'competent deduction' principle.)

Traditional invariantism is therefore *prima facie* incompatible with (MPC). There are two broad options for the traditional invariantist here. She can argue that, contra initial appearances, traditional invariantism is compatible with the principle. Alternatively, she can reject the principle. In the next section (§5.1), I argue that the traditional invariantist encounters significant problems if she rejects (MPC). In the following section (§5.2), I argue that she also encounters significant problems if she endorses that principle.

§5.1 Rejecting MPC.

Before bringing to light the problems with a traditional invariantist rejection of (MPC), it will first be useful to explore the relationship, for the traditional invariantist, between endorsing or rejecting (MPC), and endorsing or rejecting (PK) and (CK). Recall (PK) and (CK):

(PK) If not- p is epistemically possible for a subject S , then S does not know that p .

(CK) If there is a (non-zero) epistemic chance, for a subject S , that not- p , then S does not know that p .

If the traditional invariantist rejects (CK), it seems that she must also reject (MPC). There are a number of considerations that support this claim. First, in rejecting (CK), the traditional invariantist allows that, for each proposition r_i that Peter knows, there is a small chance, for Peter, that r_i is false. Assume that the number of propositions that Peter knows, n , is very large. It is a consequence of any plausible notion of epistemic chance that there will then be a *high* epistemic chance, for Peter, that the conjunction, $r_1 \& r_2 \dots \& r_n$, is false. However, it seems very implausible that Peter can know some proposition even though there is a high epistemic chance, for Peter, that that proposition is false.³⁷ In this vein, notice that utterances of the form ‘I know that p , but there’s a high chance that not- p ’ seem altogether bizarre.

³⁷ Hawthorne (2004a, p. 182-183) draws this conclusion.

Second, if there is a high epistemic chance, for a subject, that some proposition p is false, then it seems that the subject should place only a low degree of confidence in the proposition that p . It follows that if there is a high epistemic chance, for Peter, that the conjunction $r_1 \& r_2 \dots \& r_n$ is false, it seems that Peter ought to place no more than a low degree of confidence in the conjunction. That obviously sits very uncomfortably with the idea that Peter knows the conjunction.

Finally, it seems that if there is a high epistemic chance, for Peter, that a proposition is false, Peter's belief in that proposition is not safe (at all) from error. For those sympathetic to the idea that knowledge requires safety from error, this supplies a further reason to believe that rejection (CK) entails a rejection of (MPC).³⁸

Rejection of (CK) seems to require rejection of (MPC). What about the other way around? Can the traditional invariantist endorse (PK) and (CK) but reject (MPC)? Suppose the invariantist accepts (CK). It follows that, for each proposition r_i that Peter knows, there is no epistemic chance, for Peter, that r_i is false. But if, for each conjunct, there's no chance, for Peter, that it's false, what is to prevent him from knowing the conjunction?

This supplies some reason to think that the traditional invariantist who endorses (PK) and (CK) must also endorse (MPC). There are further reasons. First, suppose that for each proposition r_i that the traditional invariantist allows that he knows, Peter claims that there is no chance that r_i is false, but when faced with their conjunction, $r_1 \& r_2 \dots \& r_n$, claims that he does not know it to be true.

It seems that by claiming not to know the conjunction, Peter would be going back on one of his earlier chance claims. But on the assumption that (CK) holds but (MPC) fails, Peter's claims are perfectly consistent. In a similar vein, on the assumption that (CK) holds but (MPC) fails, we can suppose that it is true that Peter believes the conjunction $r_1 \& r_2 \dots \& r_n$ on the basis of competent deduction from propositions which are each such that there is no chance, for Peter, that they are false, but that

³⁸ Safety principles are defended in Sainsbury (1997), Sosa (1999), Williamson (2000), and Pritchard (2005). For relevant further discussion, see Hawthorne & Lasonen-Aarnio (2009).

nevertheless Peter does not know the conjunction is true. That claim seems very uncomfortable.

One final consideration. As remarked earlier, it seems that Peter would be endorsing a contradiction were he to claim, of each conjunct, that there is no chance that it is false, but then claim that there *is* a chance that the conjunction is false. Relatedly, it is difficult to understand where a chance that the conjunction is false could issue from if not a chance of falsehood associated with one of the conjuncts.³⁹ It therefore seems that the traditional invariantist who endorses (CK) must also accept that there is no epistemic chance, for Peter, that the conjunction is false. But it would then be true that Peter believes the conjunction on the basis of competent deduction from propositions that are each such that there is no chance, for him, that they are false *and* that there is no chance, for him, that the conjunction is false but – *nevertheless* – Peter does not know the conjunction. That seems bizarre.

If the traditional invariantist endorses (PK) and (CK), it seems that she must also endorse (MPC). Similarly, if the traditional invariantist rejects (CK) and (PK), that seems to entrain a rejection of (MPC). How does this bear on the conclusions of the present chapter?

It seems that the traditional invariantist who *rejects* (MPC) is strongly pressured to find some explanation of the infelicity of utterances of conjunctions like ‘I know that p, but there’s a chance that not-p’ that does not feature an appeal to (PK) or (CK). However, as we have seen, there does not seem to be a plausible explanation of that kind.

However, the problems do not stop there. There are also concerns specifically attached to a rejection of (MPC). In the remainder of this section, I will outline some of those concerns, and consider some possible responses.

³⁹ Set aside issues associated with a chance of error traceable to the deduction of the conjunction from the conjuncts. They aren’t relevant here. Cf. fn. 5.

First, (MPC) itself seems *prima facie* true. That principle is also simply a statement of the intuitive idea that deduction extends knowledge.⁴⁰ More importantly, consider our reaction were Peter to affirm of each of r_1 to r_n that he knows it to be true, but then claim *not* to know that the conjunction $r_1 \& r_2 \dots \& r_n$ is true:

Emily: Do you know that you will go to the theatre tomorrow?

Peter: Yes

Emily: Do you know that your daughter will be in California next month?

Peter: Yes

Emily: Do you know that the temperature in London will not exceed forty degrees centigrade tomorrow?

Peter: Yes

...

Emily: Do you know that you will go to the theatre and that your daughter will be in California, and ... ?

Peter: No, of course I don't know that.

In his final response, it seems as though Peter is going back on at least one of his earlier claims to know (Hawthorne, 2004a, p. 49). In a similar vein, it seems that once Peter has admitted that he does not know the conjunction, it would be appropriate for him to admit that he does not know at least one of the conjuncts and inappropriate for him to reaffirm, of each conjunct, that he knows it.⁴¹ In other words, it seems appropriate to reason in accord with (MPC). This is further evidence for its truth.

These observations present a further challenge for the traditional invariantist seeking to reject (MPC). If she cannot explain these observations away, it seems that she must attribute further significant error to speakers.

Hill and Schechter (2007) offer the following explanatory proposal: the reason that (MPC) seems true (even though it isn't) is because the principle holds in the most

⁴⁰ See, in particular, Williamson (2000, p. 117) and Hawthorne (2004a, p. 33).

⁴¹ In such circumstances, it may be appropriate for Peter to retreat to the claim that he nevertheless knows, of all of the conjuncts, that they are *very likely* to be true. This is clearly compatible with (MPC).

salient cases of competent deduction from known premises. On Hill and Schechter's account, knowledge that p is compatible with a small (but not a large) chance, on the subject's evidence, that not- p . In considering cases of competent deduction from known premises, Hill and Schechter suggest that we tend to only consider deductions involving 'paradigm cases' of knowledge – for example, knowledge of simple mathematical truths, or knowledge of our immediate perceptual environment – and only a small number of premises. Thus, in the most salient cases of competent deduction from known premises, the premises will be few and the evidential chance of error associated with each premise will plausibly be *very* small. In such cases, the chance on the subject's evidence that the conclusion is false will therefore also be small. Plausibly, sufficiently small that the subject will know the conclusion.

Hill and Schechter suggest that this proposal can also be invoked to explain the attractiveness of the idea that competent deduction extends knowledge, since competent deduction extends knowledge in the most salient cases of competent deduction from known premises.

However, it is not clear how the strategy can be successfully extended to account for all of the relevant judgments. In the dialogue between Emily and Peter presented above, we can introduce many premises – a hundred, for example – or we can introduce premises that, by traditional invariantist lights, only just amount to knowledge. It does not make the dialogue any less strange – in claiming not to know the conjunction Peter still seems to be going back on his earlier claims to know the conjuncts. Similarly, it seems that even in such a case, it would be appropriate for Peter to infer that at least one of his earlier claims to know was false. It may be plausible that such cases are not the most salient cases of competent deduction from known premises, but why don't we automatically recognise the failure of (MPC) when confronted with them? Why can't we happily judge that Peter knows each (high chance) conjunct, but not the (lower chance) conjunction?

This point is made even more pressing when we compare (MPC) with other epistemological theses. Consider, for example, the principle stating that justified true belief is sufficient for knowledge. Philosophers initially found this principle attractive. Plausibly, this was (at least in part) because in the most salient cases of

justified true belief, it seemed that the subject has knowledge. Nevertheless, when presented with cases in which the principle failed – those with certain ‘Gettier features’ – it immediately seemed true for the subject to claim that he had a justified true belief, but not true for him to claim that he had knowledge.⁴² Similarly, it seemed that were some third-party acquainted with the particulars of a subject in a Gettier situation to utter a claim of the form ‘S has a justified true belief that p, but S doesn’t know that p’, her claim would be true. And finally, despite the fact that in all the most salient cases of justified true belief the subject also had knowledge, there was little temptation to infer that the subject in the Gettier case *does* have knowledge on the basis of the fact that she has a justified true belief.

The contrast with the case of (MPC) is stark. In that case, when presented with cases in which (MPC) fails – for example, those with the feature of many premises – it does not seem true for the subject to claim to know the conjuncts but not the conjunction. For example, it does not seem true for Peter to claim to know each r_i but then claim not to know their conjunction. Similarly, it does not seem true for some third party acquainted with Peter’s situation to utter ‘Peter knows r_1 , Peter knows r_2 , ... Peter knows r_n , but Peter doesn’t know $r_1 \& r_2 \dots \& r_n$ ’. In such a ‘Many Premise’ case, there *is* significant temptation to infer from the fact that Peter doesn’t know the conjunction that he doesn’t know one of the conjuncts. These kinds of contrast put further pressure on any attempt to extend Hill and Schechter’s simple error-theoretic account.

However, perhaps the traditional invariantist could combine Hill and Schechter’s error-theoretic account with another account. Suppose that the Hill and Schechter story adequately explains why we find (MPC), and the claim that deduction extends knowledge, *prima facie* attractive. A different explanation of our judgments when confronted with cases in which (MPC) (allegedly) fails still seems to be required. How might such an explanation go?

It seems plausible that considering the conjunction of a large number of propositions, each of which the subject has fairly good evidence to believe, encourages to salience

⁴² See Gettier (1963).

the thought that the subject has made an error, which she failed to notice, with respect to one of the conjuncts. One might therefore propose that when the question of whether Peter knows the conjunction $r_1 \& r_2 \dots \& r_n$ is considered, possibilities in which Peter appears to have just the same evidence for each conjunct, but in which one or more of the conjuncts is false, will become salient.

The traditional invariantist who rejects (MPC) might therefore propose that in claiming that he does not know the relevant conjunction, Peter makes salient error possibilities that lead, for either pragmatic or psychological reasons, to our (incorrectly) judging that Peter does not know some of the conjuncts. This might account for our reactions to the dialogue between Peter and Emily.

In Chapters 1 and 2, I argued that the kind of pragmatic or psychological accounts that might go some way towards explaining why, in circumstances in which error possibilities of this kind are made salient, we would judge that the subject does not have knowledge are unsuccessful. The traditional invariantist would therefore be unwise to rely on such proposals here.

The upshot of this section is three-fold. First, that a rejection of (MPC) seems to entail a rejection of (CK) and (PK), and *vice versa*. Second, that the traditional invariantist who rejects (MPC) will inherit the problems – highlighted in §1 and §4 – associated with a rejection of (CK) and (PK). Third, that there are specific problems attached to the rejection of (MPC) to which it seems that the traditional invariantist does not have an adequate response.

§5.2 Endorsing (MPC).

Can the traditional invariantist endorse (MPC)? For reasons presented at the beginning of the previous section, to do so will seemingly entail a commitment to (CK) and (PK) and the problems – highlighted in §2 and §2.1 – that such a commitment brings. But there are further, grave difficulties. Indeed, these difficulties are perhaps the most pronounced of all those we have encountered so far. It is therefore also important to stress that insofar as a commitment to (CK) and (PK) seems to entail a commitment to (MPC), the difficulties associated with a traditional

invariantist endorsement of (MPC) also present a very serious objection to her endorsement of (CK) and (PK).

Suppose Peter competently deduces the conjunction $r_1 \& r_2 \dots \& r_n$ from its conjuncts, and that Peter knows each of the conjuncts. By (MPC), Peter knows the conjunction. However, supposing that the conjuncts are all contingent propositions about the future, and that there are sufficiently many, it follows that there will be a very high objective chance that the conjunction $r_1 \& r_2 \dots \& r_n$ is false. Moreover, it seems that Peter can know this fact. It follows that Peter knows that $r_1 \& r_2 \dots \& r_n$ but also knows that there is a very high objective chance that $r_1 \& r_2 \dots \& r_n$ is false (Hawthorne 2005a).

On the assumption that Peter lacks any clairvoyant powers, or God-like access to the future, that result seems absurd. Two further observations compound the problem. First, it seems that given that Peter knows that the objective chance that conjunction is false is very high, Peter is only justified in placing a low degree of confidence in $r_1 \& r_2 \dots \& r_n$. That seems incompatible with Peter knowing $r_1 \& r_2 \dots \& r_n$. Second, it seems that, given Peter's belief in the conjunction is objectively very likely to be false, his belief is not (at all) safe from error (Hawthorne & Lasonen-Aarnio, 2009). For those sympathetic to the idea that knowledge requires such safety, this supplies a further reason to think that Peter does not know the conjunction.

This problem is particularly stark when presented in terms of objective chance. But it is powerful even when presented without reliance on that notion. For example, suppose Peter writes down a large number of claims that he takes himself to know about the past.⁴³ Peter writes that Mary was in the office yesterday, that his daughter was in California last month, and so on. Suppose Peter occupies a moderate epistemic position with respect to each of the corresponding propositions, s_1 to s_n , – Mary told Peter that she was in the office yesterday, Peter's (reliable) daughter told him she was in California last month, and so on – and that they are all true. It seems that the traditional invariantist must allow that Peter knows each of s_1 to s_n . Suppose Peter

⁴³ On standard accounts of objective chance, true propositions concerning the past are assigned an objective chance equal to 1 (and false propositions concerning the past are assigned objective chance equal to 0).

competently deduces the conjunction of those propositions, $s_1 \& s_2 \dots \& s_n$. Given (MPC), Peter knows $s_1 \& s_2 \dots \& s_n$.

However, given the myriad potential sources of error – that Mary was lying about being in the office, that his daughter was confused about the timing of her California trip – it seems that Peter does not know that conjunction. Furthermore, it seems he would be justified only in placing a low degree of confidence in that conjunction, and that his belief in the conjunction is not safe from error. The traditional invariantist who endorses (MPC) seems to be in trouble.

Can the traditional invariantist sympathetic to (MPC) respond? Williamson (2009) proposes that (MPC) holds, and that our (mistaken) judgment to the effect that Peter does not know such long conjunctions can be explained by appeal to putative failures of knowledge of knowledge.⁴⁴

Suppose that, for each proposition r_i , Peter knows that r_i but does not know that he knows r_i . In such a case, even the traditional invariantist who endorses (CK) is free to maintain that, for each proposition r_i , there is a small epistemic chance, for Peter, that he does not know that r_i . Let's assume that to be the case.⁴⁵ For example, let's assume that there is a small epistemic chance, for Peter, that he does not know that his daughter will be in California next month. Suppose Peter believes the conjunction $r_1 \& r_2 \dots \& r_n$ on the basis of competent deduction from its conjuncts. In that case, it is plausible that if Peter knows the conjunction $r_1 \& r_2 \dots \& r_n$, then it is in virtue of having deduced it from the conjuncts. Let's therefore also assume that it's epistemically certain (i.e. epistemic chance equal to 1), for Peter, that he knows the conjunction only if he knows all of the conjuncts. Finally, suppose that the number of conjuncts, n , is very large and that epistemic chances obey the axioms of standard probability theory.

⁴⁴ See Williamson (2009, p. 314-19). For ease of exposition, I shall simplify Williamson's discussion to some extent.

⁴⁵ Indeed, given certain assumptions, it follows that this will be the case. See Williamson (2009, p. 314-315).

It is being assumed that, for each conjunct r_i , there's a small chance, for Peter, that he doesn't know that r_i . Given n is very large, there will therefore be a large epistemic chance, for Peter, that he does not know all of the conjuncts. Correspondingly, there will be a small epistemic chance, for Peter, that he does know all of the conjuncts. Given that it's epistemically certain, for Peter, that he knows the conjunction only if he knows all of the conjuncts, it follows that the epistemic chance, for Peter, that he knows the conjunction can be no greater than the epistemic chance, for Peter, that he knows all of the conjuncts. Consequently, there will only be a small epistemic chance, for Peter, that he knows the conjunction $r_1 \& r_2 \dots \& r_n$.

Williamson suggests that this last fact is crucial – at least, if we assume that a subject's judgment about whether she knows some proposition p is driven by facts about how epistemically likely or unlikely it is, for her, that she knows p . In the above example, the epistemic chance, for Peter, that he knows $r_1 \& r_2 \dots \& r_n$ is small. We might therefore expect Peter to be reluctant to judge that he knows the conjunction $r_1 \& r_2 \dots \& r_n$. Similarly, the epistemic chance, for Peter, that he does *not* know the conjunction $r_1 \& r_2 \dots \& r_n$ is large. We might therefore expect Peter to be inclined to judge that he does not know the conjunction. These facts are nevertheless perfectly compatible with Peter typically being inclined to judge that he knows a given proposition r_i , since for each proposition r_i , there is a large epistemic chance, for Peter, that he knows r_i . For example, there is a large epistemic chance, for Peter, that he knows his daughter will be in California next month. We can therefore expect Peter to be inclined to judge that he knows that his daughter will be in California next month.

How successful is this proposal? There are perhaps two central limitations. The first is that it is not clear how to extend the suggestion to account not only for the fact that *Peter* may be inclined to judge that he doesn't know the lengthy conjunction, but also for the fact that it seems *to us* that Peter does not know the lengthy conjunction. In the presentation of the case, it is stipulated (a) that r_1 to r_n are all true, (b) that Peter occupies a moderate epistemic position with respect to each of r_1 to r_n , and (c) that Peter has competently deduced the conjunction of r_1 to r_n . Given that these facts are simply stipulated, it is far from clear how to make a parallel case that the epistemic

chance, *for us*, that Peter knows the conjunction is very small. So why would we (mistakenly) judge that Peter doesn't know the conjunction?

The second limitation is that the traditional invariantist will presumably have to allow that much of our knowledge is knowledge that we know we have. The typical case will therefore not have the profile of the one just given. Let's therefore suppose instead that, for each proposition r_i , Peter not only knows that r_i but also knows that he knows that r_i . On this supposition, Williamson's argument for the conclusion that there is only a small epistemic chance, for Peter, that he knows the conjunction does not go through. The reason is that that argument rested on the assumption that, for each proposition r_i , there is a small epistemic chance, for Peter, that he does not know r_i . But given the assumption that, for each proposition r_i , Peter knows that he knows r_i , it follows from (CK) that, for each proposition r_i , there is *no* epistemic chance, for Peter, that he does not know r_i .⁴⁶

To handle this – presumably much more common – kind of case, Williamson offers an alternative proposal. He suggests that we can explain Peter's reluctance to judge that he knows the conjunction $r_1 \& r_2 \dots \& r_n$ in cases of this sort by appeal to the fact that there will be *some* level of higher-order knowledge of the conjunction that Peter lacks. And:

[T]he apparent absence of a given number of iterations of knowledge can cause doubts about all lower numbers of iterations, by a domino effect, since lack of knowledge that one has $n + 1$ iterations implies lack of warrant to assert that one has n iterations. (2009, p. 318-319)

An initial problem with this proposal stems from the fact that if Peter's higher-order knowledge of the conjunction gives out at some level, then presumably his higher-order knowledge of (at least some of) the conjuncts does too. Despite this, Peter is not typically reluctant to judge that he knows any given conjunct. For example, Peter will typically be happy to judge that he knows that his daughter will be in California next month. So why is the situation different with respect to the conjunction?

⁴⁶ Recall that the traditional invariantist who accepts (MPC) must accept (CK).

It might be proposed that the relevant difference is that when Peter considers whether or not he knows some long conjunction, it typically becomes *salient* to him that he lacks some higher-order knowledge of that conjunction. Indeed, this may be necessary for the ‘domino-effect’ Williamson talks of to come into play. But why would it typically become salient to Peter that there is such an absence? And why would it typically become salient that there is an absence of relevant higher-order knowledge when Peter considers whether or not he knows the conjunction but not when he considers whether or not he knows that his daughter will be in California next month?

More importantly, it’s not clear how sensitive it is reasonable to expect subjects to be to facts about the presence or absence of significantly higher-order knowledge. Simply coming to recognise, for example, that I don’t know that I know that I know that I know that Peter is in Athens may constitute a rather majestic cognitive feat. Is our sensitivity to *this* kind of fact plausibly responsible for our judgments? And although I don’t fully understand how Williamson’s ‘domino-effect’ reasoning is supposed to work, it seems that the most that such considerations should be able to achieve is getting the subject to entertain *some level* of doubt about whether she has the relevant first-order knowledge. It’s not clear that this is sufficient to explain not only why Peter would refrain from judging that he knows the (objectively unlikely) conjunction $r_1 \& r_2 \dots \& r_n$, but also why he would go so far as to judge that he does not know that conjunction.⁴⁷

Let’s set such objections aside for a moment. In general, providing an explanation for why, on the assumption that a particular theory is true, certain *prima facie* problematic judgments might arise may do little to assuage the sense that the consequences of that theory are still deeply implausible. In that case at hand, – and even setting aside the objections just pressed – providing an explanation, grounded in our putative sensitivity to facts about higher order knowledge failure, for why we would judge that subjects like Peter do not know such (objectively unlikely) conjunctions as $r_1 \& r_2 \dots \& r_n$ seems to do little to assuage the sense that the

⁴⁷ And even if such reasoning might lead *Peter* to entertain sufficient doubt to lead him to judge that he doesn’t know the conjunction, it’s another step still to argue that it would lead *us* to entertain sufficient doubt about whether Peter knows the conjunction to be inclined to judge that he doesn’t know it.

consequences of endorsing both traditional invariantism and (MPC) render that combination of positions deeply implausible. For example, it seems to do little to assuage the deep implausibility in the claim that a subject can know that some proposition p has an objective chance of being true that is close to 0, and yet – absent clairvoyant powers – nevertheless know that p is true.

I conclude that traditional invariantist endorsement of (MPC) comes at an enormous cost. It seems deeply implausible to maintain that a subject can know that some proposition p has an objective chance of being true that is close to 0, and yet – absent clairvoyant powers – nevertheless know that p is true. Such a claim is lent further support given some very plausible relationships between (i) knowledge of objective chance and (ii) rational levels of confidence, and perhaps also by the idea that knowledge of some proposition p requires a kind of safety from error that is incompatible with there being a (known) high objective chance that p is false. Finally, in the absence of some more promising explanation than one grounded in our putative sensitivity to failures of higher order knowledge, the traditional invariantist position is further imperilled by the fact that she seems unable to account for why, on the assumption that (MPC) is true, we would (incorrectly) judge that subjects like Peter do not know such (objectively unlikely) conjunctions as $r_1 \ \& \ r_2 \ \dots \ \& \ r_n$.

§6. Summary.

It will be useful to summarise the discussion of the chapter and the interplay between the various principles introduced. First, the infelicity of conjunctions of the form ‘I know that p , but there’s a chance that not- p ’ and ‘I know that p , but it might be that not- p ’ suggests that (one aspect of) the relationship between knowledge and epistemic chance/possibility is captured by principles (CK) and (PK).

(PK) If not- p is epistemically possible for a subject S , then S does not know that p .

(CK) If there is a (non-zero) epistemic chance, for a subject S , that not- p , then S does not know that p .

In §2 and §2.1, I argued that acceptance of these principles appears extremely problematic for the traditional invariantist, given the wealth of considerations apparently showing that, for much of what we know, there's a chance/possibility that we are wrong. In §3, I argued that traditional invariantist attempts to respond to these concerns appear unsuccessful. In §4, it was then argued that traditional invariantist explanations of the infelicity of the original conjunctions that do not involve appeal to (CK) and (PK) are unsuccessful, and further problematic phenomena associated with the traditional invariantist rejection of (CK) and (PK) were introduced. The upshot was that whether or not the traditional invariantist accepts (CK) and (PK), she must attribute significant error to speakers.

In the final section of the chapter, I discussed the relationship, for the traditional invariantist, between (PK) and (CK) and the extension of knowledge by deductive inference – in particular, the relationship between (PK) and (CK) and (MPC).

(MPC) If a subject S knows propositions p_1 to p_n and competently deduces q from p_1 to p_n – retaining throughout her knowledge of p_1 to p_n – then S knows q .

It was argued that a traditional invariantist commitment to (CK) and (PK) appears to require her commitment to (MPC), and *vice versa*. The remaining discussion compounded the problems with traditional invariantism. The traditional invariantist who endorses (CK) and (PK) is faced, not only with the problems brought to light in §2 and §2.1, but also with the very troubling conclusion that a subject can know a proposition despite knowing that it is extremely objectively unlikely to be true. The traditional invariantist who rejects (CK) and (PK) is faced, not only with the prospect of no adequate explanation of the infelicity of conjunctive claims such as 'I know that Peter is in the office, but he might not be', and related phenomena introduced in §4, but also with the fact that ordinary speakers seem to be committed to (MPC) in a way that cannot be readily explained away.

Thus, the argument of the last three chapters suggests that in respect to (i) the HIGH-LOW contrast, (ii) our talk concerning the relationship between knowledge and further inquiry, (iii) our talk concerning the relationship between knowledge and epistemic chance and possibility, and (iv) our judgments pertaining to the closure of knowledge

under (multi-premise) deductive inference, the traditional invariantist is forced to posit significant speaker error. The cumulative weight of these problems seems to warrant a move away from traditional invariantism. In the final two chapters, I shall explore the merits of some non-traditional alternatives.

Chapter 4: Interest and Salience Relativity.

The traditional invariantist is an *intellectualist*. She asserts that whether a subject's true belief amounts to knowledge is a matter only of such 'truth-relevant' factors as the nature of the subject's evidence or the reliability of her belief forming mechanisms. One popular response to the kind of phenomena we have been studying is to suggest that intellectualism is false and that whether a subject's true belief amounts to knowledge is *also* a matter of her practical environment or perhaps the error possibilities salient to her.

The anti-intellectualist can come in many stripes. In particular, she can join the traditional invariantist in maintaining both that 'knows' is not context-sensitive and that we know much of what we take (ordinarily) ourselves to know, or she can make *additional* departures from traditional invariantism with respect to these commitments as well. The focus of the present chapter is on the more conservative form of anti-intellectualism – commonly termed Subject Sensitive Invariantism (SSI) – that features both the claim that 'knows' is not context sensitive and the claim that we know much of what we take ourselves to know. I leave discussion of those proposals that involve rejecting the latter claims to the final chapter.

In §1, I outline some of the central motivations for SSI. In §2-4, I introduce some standard objections to the most common form of SSI and consider some possible responses. Finally, in §4-7, I consider the relationships between knowledge and chance in the light of an endorsement of SSI. The upshot will be that SSI faces a number of very serious difficulties and, in consequence, does not seem to represent a significant improvement on traditional invariantism.

§1. Motivations.

An initial motivation for SSI is that it enables us to respect the HIGH-LOW contrast, at least as that contrast is manifested in the Bank Case. Recall that in the LOW context of the Bank Case, it seems true and appropriate for DeRose to claim that he knows that the bank will be open on Saturday, but that in the HIGH context, it seems true and appropriate for DeRose to deny that he knows that same proposition.

DeRose's belief that the bank will be open seems to bear the same 'truth-relevant' properties in LOW and HIGH – it is based on the same evidence, formed via an equally reliable method, and so on. If that's right, then given intellectualism, it follows that DeRose in HIGH knows that the bank will be open if and only if DeRose in LOW knows that the bank will be open. But the proponent of SSI is not bound by a commitment to intellectualism. She therefore proposes that the difference in DeRose's practical environment between LOW and HIGH, or perhaps the difference in the error possibilities salient to him, has the result that DeRose knows in LOW but not in HIGH. This straightforwardly accounts for why it seems true and appropriate for DeRose to self-ascribe knowledge in LOW, but deny himself knowledge in HIGH.

It is *far* more common to suggest that the crucial difference between LOW and HIGH lies in DeRose's practical environment. It is therefore important to ask what might motivate that kind of position in particular. One prominent motivation for a version of SSI that ties knowledge to the subject's practical environment begins with the claim that one of the following principles (or something relevantly similar) is true:¹

- (1) a. If S knows that p, then S can appropriately rely on p in her practical reasoning.
- b. If S knows that p, then S can appropriately treat p as a reason for acting.
- c. If S knows that p, then S can appropriately act as if p.

In what follows, I shall largely work with (1a), but the discussion could readily be recast in terms of a principle like (1b) or (1c). Let's review some of the most important considerations that speak in favour of (1a).

¹ See, in particular, Fantl & McGrath (2002, 2007, 2009), Hawthorne (2004a, ch. 4), Stanley (2007, p. 202), Hawthorne & Stanley (2008), and (though the focus is more on justification than on knowledge) Weatherson (2005).

One important source of support for (1a) issues from the conjunctions linking knowledge and inquiry that we encountered in Chapters 1 and 2.² Recall that utterances such as those of the form ‘I know that p, but I must inquire further’ typically seem infelicitous. For example, it seems infelicitous for DeRose (in HIGH or LOW) to utter ‘I know that the bank will be open, but I must inquire further’. If a subject is able to appropriately rely on some proposition p in her practical reasoning, it seems at least *prima facie* plausible that she need not inquire further into whether or not p is true. (If a subject can rely on some proposition in her practical reasoning, why would she need to inquire further into its truth?) Consequently, on the assumption that (1a) is true, it is entirely unsurprising that utterances of conjunctions like ‘I know that the bank will be open, but I must inquire further’ typically seem infelicitous. In a similar vein, notice that utterances of the form ‘I know that p, but I cannot rely on p in my practical reasoning’ also typically seem infelicitous. For example, an utterance of ‘I know that the bank will be open, but that the bank will be open is not something I can rely on in my practical reasoning’ by DeRose in HIGH seems infelicitous. This lends even more direct support to (1a).

Another set of considerations that lend at least *prima facie* support to (1a) – which draw on phenomena we have not encountered heretofore – concerns how we praise and criticise action.³ Let’s consider an example of the latter sort. Suppose that Mary goes back into the house for a third time to check that the stove is turned off. In that case, it will typically seem appropriate for her husband, Peter, to criticise Mary’s action by saying ‘You know that the stove is turned off; stop checking!’.

In the example, it is plausible that Mary is engaging in further inquiry that she need not engage in and that, as a result, she is wasting time. Peter’s ability to appropriately criticise Mary by citing the fact that she knows that the stove is turned off can be neatly accounted for on the assumption that (1a) is true. If Mary knows that the stove is off, then, on the assumption that (1a) is true, it follows that Mary can rely on the proposition that the stove is off in her practical reasoning. But if Mary can rely on the proposition that the stove is off in her practical reasoning, then it’s plausible that she need not inquire further into its truth and that, as a result, she is wasting her time by

² See, in particular, Hawthorne (2004a, ch. 4).

³ See, in particular, Fantl & McGrath (2007) and Hawthorne & Stanley (2008).

going to check that the stove is turned off for a third time. Given (1a), the force and appropriateness of Peter's criticism are readily accounted for.

These considerations provide some strong support for (1a). However, it can seem that the only plausible way to allow that (1a) is true and *also* to allow that we know much of what we take ourselves to know is to embrace the kind of anti-intellectualism the proponent of SSI councils us to endorse.

The reason is that in circumstances in which a subject has a great deal at stake in whether or not some proposition is true, it seems that the subject needs to occupy a very strong epistemic position in order to rely on that proposition in her practical reasoning. For example, if Mary is a person of modest means and is offered a bet that pays a penny if Peter is in the office, but that costs Mary fifty thousand pounds if Peter is not in the office, it does not seem that Mary can rely on the proposition that Peter is in the office in her practical reasoning unless she occupies a very strong epistemic position with respect to that proposition. To illustrate: in the circumstances in which Mary is offered the bet, the following reasoning is acceptable only if Mary occupies a very strong epistemic position with respect to the first premise:⁴

- (1) Peter is in the office.
- (2) If Peter is in the office, then if I accept the bet, I gain a penny.
- (3) If Peter is in the office, then if I accept the bet, I lose nothing.

Therefore,

- (4) If I accept the bet, then I gain a penny and lose nothing.

Therefore,

- (C) I should accept the bet.

⁴ Cf. Hawthorne (2004a, p. 29).

It seems appropriate for Mary to reason in this way only if she has very strong grounds for the first premise. For instance, whilst it is arguably appropriate for Mary to reason in this way if she is in the office and engaged in a face-to-face conversation with Peter, it does not seem appropriate if she is simply relying on the testimony of a (usually very reliable) friend.

Now consider a subject, S, who has a true belief that p. Suppose – as just argued – that a subject who has a great deal at stake in whether or not some proposition is true needs to occupy a very strong epistemic position in order to rely on that proposition in her practical reasoning. In that case, it seems that it will be (metaphysically) possible that a subject, S*, has a true belief that p that matches S's belief that p with respect to 'truth-relevant' properties, but – because S* has a great deal at stake in the matter of whether p is true – S* needs to occupy a very strong epistemic position with respect to p in order to appropriately rely on p in her practical reasoning,

S and S*'s true beliefs that p are stipulated to match as regards 'truth-relevant' properties. On the assumption that intellectualism is true, it follows that S knows that p if and only if S* knows that p. However, given (1a), S* knows that p only if she occupies a very strong epistemic position with respect to p. It follows that S knows that p only if she occupies a very strong epistemic position with respect to p.

That result holds for arbitrary S and p – i.e. it represents a general necessary condition on knowledge. However, it seems that we do not occupy a very strong epistemic position with respect to a great many of those propositions that we ordinarily take ourselves to know. It follows that given intellectualism and (1a), we do not know much of what we ordinarily take ourselves to know.

Notice that the truth of (1a) does not, by itself, bring about such sceptical consequences. In more typical practical environments – those characterised by lower stakes – it seems that subjects can appropriately rely on propositions in their practical reasoning even though they do not occupy a very strong epistemic position with respect to those propositions. The anti-intellectualist suggests that, in consequence, she can embrace (1a) *and* maintain that we know much of what we ordinarily take ourselves to know. If knowledge comes and goes simply with changes in the subject's

practical environment, then even though S's and S*'s beliefs match with respect to truth-relevant properties, it does not follow that S knows that p if and only if S* knows that p, unless S and S* *also* match with respect to (relevant features of) their practical environments. S* is in an unusual practical environment, characterised by high stakes. Just so long as S does not occupy that kind of (unusual) practical environment, she can therefore know that p even though she does *not* occupy a very strong epistemic position with respect to p. In accepting a form of SSI that ties knowledge to the practical environment of the subject, the proponent of SSI can therefore embrace (1a) and still maintain that we know much of what we ordinarily take ourselves to know. It has been alleged that this represents an important advantage of such a position.⁵

In the remainder of this chapter, I first consider some standard objections to SSI, and evaluate what seems to be the most promising responses to those objections. I then explore the standing of SSI with respect to the relationships between knowledge and chance discussed in the previous chapter. The upshot of the discussion will be that SSI is deeply problematic, and does not appear to represent a significant improvement over more traditional forms of invariantism.

§2. Initial Objections.

As remarked above, it is much more common to suggest that a version of SSI is true on which knowledge is tied to the practical environment of the subject, rather than the error possibilities salient to her. Our initial focus will therefore be on the former proposal. There are a number of well-known objections to those forms of SSI that tie knowledge to the practical environment of the subject. Let's quickly review the most troubling – each drawing on phenomena associated with variations on the Bank Case – before evaluating some possible responses.

A first problem arises with various third-person cases.⁶ Suppose that DeRose and his wife are in LOW, but that Peter and Mary are in a situation similar to HIGH. Peter and Mary have a very important cheque coming due and cannot wait until Monday to

⁵ See, in particular, Fantl & McGrath (2007).

⁶ See, in particular, DeRose (2004 & 2009, p. 230-234).

deposit their paycheques. Like DeRose, Mary was at the bank two weeks ago on a Saturday. But unlike DeRose, Mary spoke with the bank's manager earlier today (Friday), and he assured her that the bank will be open on Saturday.

Suppose Peter and Mary are considering whether to stop at the bank today (Friday) or wait until Saturday to deposit their paycheques. Peter is aware that Mary was at the bank two weeks ago on a Saturday, but is not aware that Mary spoke to the manager earlier in the day. Upon seeing the long queues at the bank, Peter suggests that they could call up DeRose to see if he can confirm whether the bank will be open tomorrow (Saturday) and Peter asks Mary (who is aware of DeRose's circumstances and epistemic position as described in LOW) whether or not DeRose might know that the bank will be open on Saturday.

In that case, it seems that it would be true and appropriate for Mary to deny that DeRose knows that the bank will be open on Saturday. It also seems false and inappropriate for Mary to claim that DeRose does know that the bank will be open on Saturday. (It nevertheless would seem true and appropriate for Mary to reassure Peter by saying that she called up the bank earlier in the day and that therefore *she* knows that the bank will be open tomorrow.)⁷

The problem is that, on SSI, DeRose knows that the bank will be open on Saturday in LOW. (This accounts for why it seems true and appropriate for DeRose to self-ascribe knowledge in LOW.) It follows that an utterance by Mary of 'DeRose does not know that the bank will be open' in the context resembling HIGH is false, and that an utterance of 'DeRose does know that the bank will be open' by Mary in that same context is true. However, the former utterance seems true and appropriate and the latter seems false and inappropriate.

A second problem concerns various counterfactual judgments.⁸ In a nutshell, the problem is that it seems that we do not take a counterfactual supposition of a change

⁷ If you do not share the judgments concerning this particular case, there are other examples that can be used to illustrate the point – for example, variations on the Thelma-Louise-Lena case featured throughout DeRose (2009).

⁸ See *inter alia* Hawthorne (2004a, p. 177, n. 40), Stanley (2005a, p. 106), and Schaffer (2006, p. 104). Hawthorne credits the objection to Timothy Williamson.

in stakes to influence whether or not a subject has knowledge. Suppose that Mary and Peter are in the situation described in the previous example – the stakes are high for them, but Mary has undertaken additional checks to establish that the bank will be open on Saturday. However, this time, suppose that DeRose and his wife are in HIGH, rather than LOW.

Mary utters ‘DeRose does not know that the bank will be open. But if there were less at stake for DeRose, he would know that the bank will be open’. That utterance seems clearly false. The problem is that – on SSI – Mary’s utterance is true. In his current practical circumstances, SSI delivers the result that DeRose does not know that the bank will be open. Yet it also seems to be a straightforward consequence of SSI that if there were there less at stake for DeRose in the matter of the bank’s being open – for instance, if he were in a practical situation more closely resembling the one depicted in LOW – then he would know that the bank will be open on Saturday.

A third problem arises when we consider claims concerning the knowledge of pairs of subjects who appear to occupy the same epistemic position with respect to some proposition, but whose stakes differ.⁹ For example, suppose that DeRose is in HIGH and that Peter is in a situation similar to LOW. DeRose believes that the bank will be open on Saturday on the basis of being there two weeks ago on a Saturday and the stakes are high for DeRose. Peter also believes that the bank will be open on Saturday on the basis of being there two weeks ago on a Saturday, but the stakes are low for Peter. On SSI, it seems that Peter will know that the bank will be open, but DeRose will not. Nevertheless, it would seem false for some third-party acquainted with Peter and DeRose’s respective situations to utter ‘Peter knows that the bank will be open, but DeRose does not’. In the same vein, it would seem true for the third-party to utter conditionals such as ‘If DeRose does not know that the bank will be open, then neither does Peter’. But on SSI, such conditionals are not true, since it is a feature of the view that whilst DeRose does not know that the bank will be open, Peter does.¹⁰

⁹ See, for example, Neta (2007, p. 182-183) and Blome-Tillman (2009, p. 328-330).

¹⁰ There are some further conditional judgments that make trouble for SSI. For example, suppose that the third-party is aware of the character of Peter’s evidence (or epistemic position), but not aware of Peter’s practical environment. In that case, it seems that SSI predicts that an utterance by the third-party of ‘If Peter has a lot at stake, he doesn’t know that the bank will be open, but if Peter has little at stake, he does know that the bank will be open’

These problems are made especially vivid when we reflect that in the case of other notions that are plausibly tied to the subject's practical environment, we do not see judgments like these. For instance, one such notion is the need for a subject to inquire further. It is plausible that this is a stakes sensitive notion: that a change (merely) in how much is at stake for a subject can influence whether further inquiry on the part of that subject is necessary. However, notice that it does *not* seem infelicitous to utter, for example, 'DeRose needs to inquire further, but Peter does not' in circumstances in which DeRose and Peter differ simply in features of their practical environment. Nor do such claims as 'DeRose needs to inquire further. But if he had less at stake, he would not need to inquire further', in circumstances in which DeRose occupies a practical environment characterised by high stakes (and occupies a modest epistemic position with respect to the proposition(s) at issue), seem false.

These contrasts bring into focus the severity of the problems reviewed above. It seems that those forms of SSI that tie knowledge to the practical environment of the subject are compromised unless something can be done to assuage the force of these problems. In the following two sections, I consider and reject the more promising proposals for how to do this.

§3. Stanley's Response and Pragmatic Responses.

In regard to our first problem, Stanley (2005a) attempts to explain away our *prima facie* problematic judgments. He suggests that in Peter and Mary's present context, Peter is really concerned with the question of whether DeRose would make a good informant for them. In consequence, in asking whether or not might DeRose know that the bank is open, Peter is really concerned with the *counterfactual* question of whether if DeRose were to share their practical situation, he would know.

When *we* are asked for our intuitions about the case, we intuitively recognise that what [Peter and Mary] *really* care about is whether [DeRose] would know, were he in [Peter and Mary's] practical situation. We recognise that this answer is negative. That is, we recognise that the proposition that

should seem true and appropriate. But the utterance seems infelicitous. Hawthorne (2004a, p. 166) credits similar objections to Stewart Cohen, Troy Cross, and Cian Dorr.

[Peter and Mary] *really* want answered – *would [DeRose] know, were he in our practical situation* – is in fact false. (Stanley, 2005, p. 102-103; emphasis in original.)

As a result, we judge that were Mary to deny that DeRose knows that the bank will be open on Saturday, her utterance would seem true and appropriate. And we judge that were she to claim that DeRose does know that the bank will be open on Saturday, her utterance would seem false and inappropriate.

An initial concern. It is far from clear that it is legitimate to appeal to our judgments concerning counterfactuals like ‘If DeRose were to share our practical environment, he would not know that the bank will be open’ in order to explain the judgment that an utterance by Mary of ‘DeRose does not know that the bank will be open’ seems true (and the corresponding positive knowledge ascription false). The reason is that, as transpired the previous section, it appears that we do not take a counterfactual supposition of a change in stakes to influence whether or not the subject has knowledge. (Unlike counterfactual suppositions of a change in ‘truth-relevant’ factors – the extent of the subject’s evidence, the reliability of the process of belief formation, and so on –, that do seem to have that influence.) The concern therefore arises that insofar as we are inclined to judge that the counterfactual ‘If DeRose were to share Peter and Mary’s practical environment, he would not know that the bank will be open’ is true, we do so *because* we judge that ‘DeRose does not know the bank will be open tomorrow’ is true and that a change in practical environment would not affect the truth of that claim. By instead invoking our judgment concerning the counterfactual to explain our judgment concerning the apparent truth of the knowledge denial (and the apparent falsity of the corresponding knowledge ascription), Stanley is open to the objection that his proposed order of explanation is implausible.

The most natural response to this concern might be to suggest that Peter and Mary are really interested in some closely related question – perhaps whether DeRose can rule out the possibility that the bank has changed its hours, rather than whether, were DeRose to share Peter and Mary’s practical environment, he would know that the bank will be open. However, even supposing such a response is successful, there are more pressing troubles.

Given the talk of ‘the proposition that [Peter and Mary] *really* want answered’, the most natural interpretation of Stanley’s proposal is that he is postulating some kind of pragmatic effect. Suppose Peter and Mary are really concerned, not with the question of whether or not DeRose knows that the bank will be open on Saturday, but instead with the question of whether or not, were DeRose to share their practical environment, he would have that knowledge. In that case, it might seem plausible that a false utterance by Mary of ‘DeRose does not know that the bank will be open’ would pragmatically convey the truth that were DeRose to share her practical environment, he would not know, and that a true utterance by Mary of ‘DeRose does know that the bank will be open’ would pragmatically convey the falsehood that were DeRose to share her practical environment, he would know. So perhaps it is these facts about what is pragmatically conveyed that are supposed to account for our judgments concerning truth and appropriacy.

However, such a pragmatic manoeuvre suffers from simple parallels of the problems presented for a traditional invariantist pragmatic manoeuvre in Chapter 2. Perhaps most pertinently, it seems that we should expect the pragmatic effect in question to be *cancellable*. But it is not. For example, it does not seem felicitous for Mary to respond to Peter by uttering ‘DeRose knows that the bank will be open, but if he were to share our practical environment, he would not know’.

It might be replied that the infelicity of this particular utterance just reflects the fact, already hinted at, that the counterfactual claim was poorly chosen in the first place. Perhaps it should instead be suggested that what is pragmatically conveyed by an utterance by Mary of ‘DeRose knows that the bank will be open’ is that DeRose can rule out the possibility that the bank has changed its hours (or something more along those lines). But this move is to no avail, since this purported pragmatic effect is not cancellable either. It does not seem felicitous for Mary to respond to Peter by uttering ‘DeRose knows that the bank will be open, but he cannot rule out that the bank has changed its hours’.

Interpreted as a pragmatic manoeuvre then, Stanley’s proposal seems implausible. But it is of course open to him to reply that he did not intend to be invoking any pragmatic

effect. In that case, there is some obvious pressure to give an account of just what kind of effect *is* being invoked. However, even in the absence of a fuller account, it is deeply unclear why the effect that Stanley is postulating would remain in place once Mary has clarified her claim.

In essence, Stanley's proposal is that, in evaluating the truth and appropriacy of an utterance by Mary of 'DeRose knows that the bank will be open', we focus on the counterfactual claim that if DeRose were to share Peter and Mary's practical environment, he would know that the bank will be open on Saturday, rather than on the simple knowledge claim that DeRose knows that the bank will be open on Saturday. The reason we focus in that way is simply because we recognise that Peter and Mary are really interested in the question of whether the counterfactual claim is true, rather than the simple knowledge claim. However, if *that* is the foundation of our judgment then there's no discernible reason why, in those circumstances in which the simple knowledge claim and the counterfactual claim are explicitly contrasted, the effect would remain in place. In other words, it seems that an utterance by Mary of 'DeRose knows that the bank will be open, but if he were to share our practical environment, he would not know' should appear perfectly felicitous, whether or not it is a pragmatic effect properly so-called that is being invoked. But as we have seen, that utterance does not seem felicitous.

What about the other two objections raised in the previous section? It seems that Stanley's suggestion could be put to work in accounting for some aspects of the third problem highlighted in the previous section. One aspect of that problem concerned various conjunctive claims, such as those of the form 'Peter knows that the bank will be open on Saturday, but DeRose doesn't' where the only relevant difference between DeRose and Peter lies in their respective practical environments. Perhaps any third-party uttering such conjunctive claims would *really* be interested in the question of whether, were DeRose and Peter to occupy her (the third party's) practical environment, they would know. However, such a manoeuvre will suffer from parallel problems. In particular, it seems that the third-party should be able to clarify-away the effect. For example, if she is in a high stakes situation, like that of DeRose in HIGH, it should be felicitous for her to utter something along the following lines, 'Peter knows that the bank will be open on Saturday, but DeRose does not. However, were DeRose

and Peter to occupy my practical situation, neither of them would know that the bank will be open on Saturday.’ But such an utterance does not seem felicitous.

It is deeply unclear how to successfully extend Stanley’s suggestion to account for the second problem – concerning our judgments about counterfactuals. An alternative seems to be required.

§4. Psychological Accounts.

Hawthorne (2004a) envisages addressing some of the problems outlined in §2 by appeal to psychological mechanisms. In particular, Hawthorne proposes an appeal to the availability heuristic similar to that discussed in Chapter 1.¹¹ Consider the first problem of §2 once again – concerning third-person knowledge ascriptions. Perhaps past overexposure to banks (or other institutions) changing their hours leads Mary to overestimate the likelihood that the bank has recently changed its hours. That might explain why she would (falsely) claim that DeRose does not know that the bank will be open on Saturday. In a similar vein, perhaps we judge that Mary’s knowledge denial is true and appropriate because, due to an overexposure to banks changing their hours, we also overestimate the likelihood that the bank has recently changed its hours.

It is worth first considering a line of attack against such a proposal that does not seem particularly troubling. DeRose (2009) takes this kind of ‘psychological’ proposal to be the best defence of SSI on the market. He then objects to it on the grounds that, given she is in a very strong epistemic position with respect to the proposition that the bank has not changed its hours – the manager informed her earlier in the day that the bank would be open on Saturday –, Mary will in fact judge that there is very little (or perhaps no) chance that the bank has changed its hours. Consequently, DeRose argues, it is not very plausible to suggest that the explanation for why Mary would judge that DeRose does not know that the bank will be open on Saturday lies in her estimation of the chance that the bank has changed its hours.¹²

¹¹ See Hawthorne (2004a, p. 163-164).

¹² See DeRose (2009, p. 236-237).

DeRose's objection could also be extended to *our* judgments to the effect that Mary would truly and appropriately deny that DeRose (of the example) knows that the bank will be open. It is simply stipulated in the case description that the bank will be open on Saturday. So aren't we also going to judge that the likelihood that the bank has changed its hours is zero?

The reason this kind of concern isn't very troubling is that there seem to be other, closely related, (mis)judgments that the advocate of SSI can appeal to. Perhaps DeRose is correct to allege that Mary will be inclined to judge that there is little or no chance that the bank has changed its hours – at least insofar as this is interpreted as the claim that there is little or no epistemic chance, *for her*, that the bank has changed its hours. Nevertheless, it seems plausible that past overexposure to banks changing their hours would – for availability reasons – lead Mary to overestimate the frequency with which banks, *in general*, change their hours. But if an individual judges that an event of a certain kind occurs more frequently than it really does, it is plausible that this will lead her to require that a subject go to greater lengths – occupy a stronger epistemic position – than she really needs to in order to know that a particular event of that kind does not occur.¹³ It is therefore plausible that Mary will require DeRose to occupy a stronger epistemic position with respect to the proposition that the bank has not changed its hours than he really does in order to know that proposition. On the assumption that (single premise) closure holds, DeRose knows that the bank will be open only if he knows that the bank has not changed its hours.¹⁴ Given (single premise) closure, it is therefore also plausible that Mary will require DeRose to occupy a stronger epistemic position than he really does with respect to the proposition that the bank has changed its hours in order to know that the bank will be open on Saturday. This could explain why Mary would judge (falsely) that DeRose does not know that the bank will be open on Saturday.

¹³ For similar reasons, Mary will also plausibly overestimate the likelihood, *for DeRose*, that the bank has changed its hours. This fact could also be exploited to explain the judgment at issue.

¹⁴ I am assuming a simplified form of closure here, on which knowledge is closed under (single premise) entailment. But the argument could be recast in terms of some other formulation of the closure principle, such as that if S knows that p, and p entails q, then S is in a position to know q.

The situation is much the same for us, as readers of the cases. Our misjudgment of the frequency with which banks change their hours in general could lead us to judge that DeRose must occupy a stronger epistemic position than he really does in order know that the bank will be open on Saturday. In turn, this could explain why it seems true and appropriate for Mary to deny that DeRose knows that the bank will be open.

However, similar problems to those discussed in Chapter 1 will afflict the proposal. For example, it is really not clear to what extent we are problematically overexposed to such events as banks having changed their hours. But nor is it clear that there is another source of the purported excess availability.¹⁵ More pressingly, even if we (and Mary) are fully aware of the just how likely it is, for DeRose, that the bank has changed its hours, it *still* seems true and appropriate for Mary to deny that DeRose knows that the bank will be open on Saturday. Such concerns, and the others raised in Chapter 1, §2, put serious pressure on the adequacy of any such proposal.

There is, however, an alternative psychological proposal. It might seem that the advocate of SSI – perhaps unlike the traditional invariantist – can make successful appeal to the kind of egocentric bias discussed towards the end of Chapter 1 (§5.2). Suppose that as a result of some egocentric bias, Mary treated DeRose as though he shared her practical concerns. In that case, it would be unsurprising if Mary were to judge that DeRose does not know that the bank will be open on Saturday, since if DeRose shared Mary’s practical concerns, then, on SSI, he would not know that the bank will be open on Saturday.

The most plausible way to cash out this egocentric proposal is probably to suggest that, despite being aware that DeRose is in a low stakes scenario, Mary allows features of her own practical situation to colour her judgment regarding how much evidence DeRose requires in order to rely on the proposition that the bank will be open on Saturday in his practical reasoning. In particular, it might be plausible to suggest that Mary overestimates the amount of evidence DeRose requires to rely on

¹⁵ See Chapter 1, §2 for further discussion.

that proposition in his practical reasoning.¹⁶ Could this explain why she would judge (falsely) that DeRose does not know that the bank will be open?

Regrettably, similar problems to those brought to light in Chapter 1 affect this kind of manoeuvre as well. Perhaps most importantly, Mary is aware that DeRose occupies a different (low stakes) practical situation from her. It is consequently rather implausible to suggest that she would treat DeRose as though he shared a practical situation exactly like hers. The most that the relevant experimental work could plausibly be used to try and support is the claim that Mary will treat DeRose's practical situation as more like her own than it really is. (See Chapter 1, §5.2, for relevant discussion.)

The empirical work therefore (arguably) supports the claim that Mary will take it that DeRose needs to occupy a stronger epistemic position than he really does in order to know that the bank will be open. But it does not support the claim that she will take it that DeRose needs to occupy an epistemic position of the same strength that *she* needs to occupy in order to know that the bank will be open. It follows that, even taking into account the effects of the egocentric bias, Mary should allow that DeRose can know that the bank will be open on Saturday on lesser grounds than she requires to know that same proposition. There should therefore be cases in which DeRose's epistemic position is such that Mary would judge that DeRose knows that the bank will be open on Saturday, even though she would also judge that it's the case that were she to occupy an epistemic position of equivalent strength, she would not know that proposition. But there do not seem to be any such cases. Insofar as the third-person case is set up such in such a way that Mary would judge that, were she to occupy DeRose's epistemic position, she would not know that the bank will be open on Saturday, it seems that Mary would respond to Peter's inquiry by claiming that

¹⁶ It might be proposed that a kind of 'anchor and adjustment' underlies the egocentric effect here – see Chapter 1, §5.2. Perhaps Mary treats her judgment of how much evidence *she* needs to gather in order to rely on the proposition that the bank will be open in her practical reasoning as an anchor and then adjusts that judgment, on the basis of information concerning the difference in stakes between DeRose and herself, to yield a judgment concerning how much evidence DeRose needs to gather in order for him to rely on the proposition that the bank will be open on Saturday in his practical reasoning. As with other cases in which such 'anchor and adjustment' is at work, it might be plausible to suggest that Mary underadjusts, and so judges that DeRose needs to gather more evidence than he really does.

DeRose does not know that the bank will be open, and that, were she to do so, her utterance would seem true.

What about the other problems brought to light in §2? The deficiency in this kind of egocentric proposal also affects attempts to apply it in those cases. For example, one aspect of the third problem of §2 was that if Peter and DeRose occupy the same epistemic position with respect to the proposition that the bank will be open, but differ in how much they have at stake in the matter of the bank's being open, then an utterance by some third-party of 'Peter knows that the bank will be open, but DeRose does not' will be true. However, such a conjunctive claim seems false and inappropriate.

One egocentric explanation would be to suggest that the third-party treats DeRose and Peter as though they share an exactly similar practical environment to her own. That would seem to provide an explanation at least of why the third-party would disinclined to assert the conjunction. However, for reasons already given, it is plausible only that the third-party will treat Peter and DeRose's respective situations as *more* like her own than they really are, rather than as exactly similar to her own. Consequently, even taking into account the egocentric effect, the third-party should take DeRose and Peter to require differing epistemic positions in order to know that the bank will be open. There should therefore still be cases in which the third-party will be inclined to judge that 'Peter knows that the bank will be open but DeRose does not' even though Peter and DeRose differ only with respect to features of their practical environments and not with respect to the character of their epistemic positions. But it does not seem that there are such cases.

Where does this leave us? The extent of these problems suggests that perhaps the best option for the advocate of SSI is to argue that other views are afflicted by parallels of these problems (or perhaps by other, more serious problems) and that her view can still triumph on the grounds of 'best fit', even in the absence of a satisfying explanation of our mistaken judgments. One purpose of the next chapter will to argue that there is a form of invariantism that better fits our judgments than SSI, and that this line of response will therefore prove unsuccessful.

§5. Knowledge and Chance.

The critique of SSI thus far has been focussed on Bank Case-like judgments and on a formulation of SSI on which knowledge is tied to the practical environment of the subject. But as noted in the Introduction, cases in which there are shifts *merely* in what error possibilities are salient to the relevant attributors seem to offer similar HIGH-LOW contrasts.

It is obviously interesting to ask about the extent to which SSI might be able to address such cases and the extent to which such cases might represent an objection to SSI.¹⁷ There is also the question of the relationships between knowledge and epistemic chance/possibility explored in the previous chapter, which seemed to create such trouble for traditional invariantism. It is the latter issues that I shall focus on here, but the discussion to follow will also bear on the former.

Recall that utterances of the form ‘I know that p, but there’s a chance that not-p’ and ‘I know that p, but it might be that not-p’ typically seem infelicitous. This suggests that principles (CK) and (PK) are true.

(CK) If S knows that p, then there is no epistemic chance, for S, that not-p

(PK) If S knows that p, then it is not epistemically possible, for S, that not-p

Can the advocate of SSI endorse such principles?¹⁸ In the previous chapter, we reviewed a number of arguments that seemed to suggest that, for much of what the traditional invariantist claims that we know, there is an epistemic chance (and possibility), for us, that we are wrong.

¹⁷ Schiffer (2007, p. 190-191) notes that a stakes sensitive version of SSI cannot handle such cases. Schiffer credits the observation to Jeremy Landen.

¹⁸ Leading advocates of SSI are split on the issue of these principles. Stanley (2005b) explicitly endorses at least (PK) and Hawthorne (2004a, esp. p. 24-28) expresses considerable sympathy for both (PK) and (CK). However, Fantl and McGrath (2009, ch. 1) reject at least (CK).

Recall the first argument to that effect that we encountered. Suppose that Mary truly believes that Peter is in the office. Mary's belief is based on the fact that she has seen Peter's coat hanging in the office hallway and on the fact that a reliable colleague has just informed Mary that Peter is in the office. Ordinarily, we would take Mary to know that Peter is in the office. Nevertheless, it seems that there's a chance, for Mary, that Peter is not in the office. (There's a chance, for Mary, that Peter has just walked out, there's a chance, for Mary, that her colleague was lying, and so on.) Similarly, it seems true for Mary to utter 'There's a chance that Peter isn't in the office'.

As we saw, the traditional invariantist has problems accounting for such judgments. Can the advocate of SSI endorse (CK) and yet nevertheless provide an explanation of judgments like these?

In the example, we need not suppose that there is much at stake in the question of whether Peter is in the office in order for Mary's utterance of 'There's a chance that Peter isn't in the office' to seem true. It therefore seems that, given the character of Mary's evidence, a version of SSI that ties knowledge simply to the practical environment of the subject will have the result that Mary knows that Peter is in the office. Given (CK), it will then follow that there is in fact no chance, for Mary, that Peter is not in the office. This obviously seems to make trouble for the apparent truth of her claim to the effect that there is.

In response, it might be proposed that in considering Mary's utterance of 'There's a chance that Peter is not in the office', it becomes salient to us that, given her evidence, there are certain practical environments in which she would be unable to rely on the proposition that Peter is in the office in her practical reasoning. It might be proposed that (a) we realise that in such circumstances, Mary would not know that Peter is in the office, and (b) that we take it that knowledge does not come and go simply with changes in practical environment (cf. Hawthorne 2004a, p. 180 & 182-183). That might explain why we judge that Mary's utterance seems true.

However, the response is clearly problematic. Even supposing that it does indeed become salient to us that Mary would be unable, in certain practical environments, to rely on the proposition that Peter is in the office in her practical reasoning (not an

obvious claim), the response relies on substituting one troubling mistake in our judgments for another. If the relevant form of SSI is true, knowledge *does* come and go simply with changes in practical environment. Indeed, that is a core of the thesis. But then why would competent speakers take it that knowledge does *not* come and go simply with changes in practical environment? (Note that this is not something we do with other plausibly practical-environment sensitive notions, like the need for further inquiry.) The proposal here is clearly a troubling one – at the very least, pursuing this kind of response represents a notable cost.

There is another way to try and invoke the apparently stakes-sensitive aspect of SSI to explain, for instance, why an utterance of ‘There’s a chance that Peter is not in the office’ by Mary seems true. Suppose the stakes-sensitive aspect of SSI is (at least partially) reflected in some principle like (1a). It might then be proposed that which propositions it is appropriate for a subject to rely on in her practical reasoning is partially determined by the error possibilities she is considering. In the previous chapter, it was suggested that by claiming that there is chance, for her, that she’s mistaken about Peter being in the office, Mary might bring to salience possibilities in which she has exactly similar evidence that Peter is in the office but in which Peter is not in the office. It might therefore be proposed that, simply as a result of such possibilities becoming salient, it is no longer appropriate for Mary to rely on the proposition that Peter is in the office in her practical reasoning. It would then follow from (1a) that, in her current circumstances, Mary does not know that Peter is in the office. Those proponents of SSI who seek to endorse (CK) could therefore maintain that there *is* a chance, for Mary, that Peter is not in the office, and that it is therefore unsurprising that her claim that there is such a chance seems true.

However, it seems rather implausible that simply reflecting on error possibilities should limit the range of propositions it is appropriate for a subject to rely on in her practical reasoning.^{19,20} Suppose Peter and Emily drive to the eastern edge of the

¹⁹ Except in the (unusual) case of those propositions that concern what error possibilities the subject is currently entertaining.

²⁰ This is one juncture at which the differences between (1a-c) are relevant. In particular, this kind of suggestion might not get off the ground if it is (1c) that expresses the important relationship between knowledge and the subject’s practical environment. For example, on the kind of characterisation of ‘acting as if p’ offered by Fantl and McGrath (2002), whether it is

forest together and leave their car there. They then walk off in different directions for half an hour. Now consider the following practical reasoning:

(1) The car is parked at the eastern edge of the forest.

(2) If the car is parked at the eastern edge of the forest, I should go there to get home.

Therefore,

(C) I should go to the eastern edge of the forest to get home.

It does not seem particularly plausible that *simply* in virtue of the fact that Emily is currently considering the possibility that the car has been stolen but Peter is not, Peter can appropriately engage in the above practical reasoning but Emily cannot. (Of course, considering that kind of possibility might influence which propositions Emily does in fact rely on in her practical reasoning, but the question is whether it influences which propositions it is *appropriate* for her to rely on.)

This proposal will also share in the implausibilities of a final option. That option is to extend SSI to explicitly include the claim that in order to know some proposition, the subject must be able to rule out the alternatives to that proposition salient in her conversation (or something along these lines). Suppose, as suggested earlier, that by claiming that there is a chance, for her, that she's mistaken about Peter being in the office, Mary brings to salience a possibility in which she has exactly similar evidence that Peter is in the office but in which Peter is not in the office. If that kind of possibility becomes salient to Mary, and something like the salience-based condition just proposed is on the right track, then it will follow that, in the example, Mary does not know that Peter is in the office. It is then open to the advocate of (CK) to maintain that there is a chance, for Mary, that Peter is not in the office, and that is why her utterance of 'There's a chance that Peter is not in the office' seems true.

appropriate for a subject to act as if p is clearly independent of the error possibilities she is currently considering (except in such unusual cases as those in which p is about the error possibilities the subject is currently considering).

There are a variety of problems with such a manoeuvre. First – and perhaps most importantly in the present context – it does not cover all of the relevant cases. Suppose that Mary is aware that a friend, Joanna, believes on the basis of similar evidence to her that Peter is in the office, but that Joanna is not even considering the question of whether Peter is in the office (nor is Peter’s location of practical significance to her). In that case, even assuming a version of SSI that embraces a salience-condition, Joanna will know that Peter is in the office. Nevertheless, it seems true not only for Mary to utter ‘There’s a chance that Peter is not in the office’, thereby indicating something about her own epistemic state, but *also* true for her to utter ‘There’s a chance, for Joanna, that Peter is not in the office’. Even on the assumption that a salience-sensitive form of SSI is true, our ‘chance’ judgments still represent a problem for SSI.

Second, the proposal is costly in that it multiplies the problems outlined in §2. For example, it seems that the advocate of SSI will also have to endorse the truth of such statements as ‘Mary does not know that Peter is in the office. But if Mary were not considering the possibility that Peter has just walked out of the office, she would know that he’s in the office’. Such statements seem clearly false.

There are some more general problems too. Prominent advocates of SSI take it to be an advantage of their view that they can respect the kind of link between knowledge and practical reasoning expressed in (2a).²¹

- (2) a. S can appropriately rely on p in her practical reasoning only if S knows that p

Indeed, Stanley (2005a) takes our Bank Case judgments to reveal (perhaps above all things) a commitment to this kind of principle.²² The problem is that endorsing a salience-sensitive version of SSI threatens make trouble for a commitment to (2a). As argued above, it is rather implausible to suggest that simply by reflecting on error possibilities, a subject limits the range of propositions it is appropriate for her to rely

²¹ Stanley (2005a, esp. p. 11), Hawthorne and Stanley (2008), and (more tentatively) Hawthorne (2004a, ch. 4), endorse some version of (2a).

²² See Stanley (2005a, esp. p. 11).

on in her practical reasoning.²³ But this is a simple consequence of a form of SSI that combines an explicit salience-sensitivity condition and a commitment to (2a).

Another more general problem.²⁴ Dretske (2005) objects to *contextualism* on the grounds that it has the following problematic implication:

philosophers who spend time worrying about heavyweight implications (How do I know I'm not dreaming? How do we know there is a material world?) are the most ignorant people in the world. Not only don't they know these heavyweight implications ... they don't (like everyone else) know the things (that there are cookies in the jar, zebras in the pen) that imply them. This, of course, includes almost everything they thought they knew.

This result, it seems to me, is pretty bizarre – *more* bizarre, in fact, than abandoning closure. (Dretske 2005, p. 18-19).

It is often pointed out that contextualism is not much impugned by this kind of objection. Intended as an objection to contextualism, it appears to rest largely on a use-mention confusion.^{25,26} However, this species of objection does have force against the form of SSI under consideration. If knowledge comes and goes simply depending on the error possibilities the subject is considering – as this form of SSI would have it – then simply by engaging in reflection on one's knowledge, it can go away. In the same vein, the layman may know what the concerned philosopher does not *simply* in virtue of the latter, but not the former, reflecting on how he might have been led into error. Both results can seem bizarre.

Indeed, a similar problem also seems to afflict those forms of SSI which tie knowledge to the practical environment of the subject.²⁷ It can seem pretty bizarre that those who care more about whether things are true can know a great deal less

²³ Except in the unusual case of those propositions that concern which error possibilities the subject is currently considering.

²⁴ See Hawthorne (2004a, p. 167). See also Hawthorne (2004a, p. 166, n. 18) for another, related, problem (one that Hawthorne credits to Stewart Cohen).

²⁵ See *inter alia* Hawthorne (2004a, p. 167; 2005) and DeRose (2009, p. 222-225).

²⁶ Note that Dretske's formulation of the objection is a little misleading (even setting aside use-mention worries). The relevant theoretical implication is at most that, *whilst they are actively engaged in a certain kind of epistemological reflection*, philosophers are extremely ignorant (where ignorance is understood as a lack of knowledge).

²⁷ See, for example, Russell & Doris (2008).

than those who care little on exactly the same evidence (or in virtue of occupying the same epistemic position). The carefree marijuana smoker may know what the concerned politician does not *simply* in virtue of being carefree? This implausibility in even the most widely defended forms of SSI is worth reflecting on.

These problems illustrate the serious difficulties faced by the proponent of SSI in addressing even the most straightforward of the problems, highlighted in the previous chapter, with a traditional invariantist endorsement of (CK) and (PK). Unsurprisingly, the other problems highlighted with a traditional invariantist endorsement of (CK) and (PK) – see §2 and §3 of Chapter 3 – also afflict SSI. The issues are sufficiently similar that I leave it to the reader to determine that SSI fares at least as badly there.

In fact, there are further problems besides those reviewed in the previous chapter. In particular, it seems that if the proponent of SSI ties knowledge to *either* the practical environment *or* the error possibilities salient to the subject, but also endorses (CK), she will be left with the result that whether or not there's a chance, for the subject, that some proposition *p* is true can vary simply with changes in her stakes (or what error possibilities she is considering).²⁸ Claims such as the following will then come out true: 'There's a chance, for her, that she's mistaken. But if there were merely less at stake for her (or she were merely considering fewer error possibilities), there would be no chance, for her, that she's mistaken'. But such claims seem *clearly* infelicitous. Similar problems arise for the proponent of SSI's endorsement of (PK).

²⁸ The combination of SSI and (CK) doesn't *quite* entail that whether there's a chance, for the subject, that some proposition is true will come and go simply with changes in the practical environment of the subject (or the error possibilities salient to her). In particular, the advocate of SSI can endorse (CK) and avoid the problems about to be highlighted in the main text by also endorsing the following principle (a parallel principle can be envisaged to deal with the error possibilities salient to the subject):

(CK)*** If a subject *S* would know that *p* if she had nothing at stake in whether *p* is true (and her epistemic circumstances were otherwise unchanged), then there's no epistemic chance, for *S*, that *p*.

However, when combined with SSI, this principle leads to many oddities of its own. For instance, if Peter is in high stakes circumstances, the following may come out true: 'There's no chance, for Peter, that he's wrong about *p*, but he doesn't know that *p*' and 'There's no chance, for Peter, that he's wrong about *p*, but he needs to check further' (and also in combination: 'There's no chance, for Peter, that he's wrong about *p*, but he doesn't know that *p* and he needs to check further').

The upshot of this section is two-fold. First, there seem to be clear costs associated with an SSI endorsement of (CK) and (PK). (We shall uncover some more in the discussion of (MPC) in §5.) Second, it was shown to be very problematic to tie knowledge to the error possibilities salient to the subject (either through introducing an explicit ‘salience condition’ on knowledge, or by proposing that such a condition falls out of an endorsement of (1a)). This leads to a further conclusion. If knowledge is not tied to the salience of error possibilities in this way, the advocate of SSI is *additionally* left (or so it would seem) with no explanation of the HIGH-LOW contrasts apparently elicited by a change in the salience of error possibilities alone (even in those examples in which subject and attributor are one and the same). In regard to that particular species of HIGH-LOW contrast, it therefore seems that the advocate of SSI is in exactly the same troubled boat as the traditional invariantist.

§6. Rejecting (CK) and (PK).

Can the advocate of SSI explain the problematic chance-knowledge conjunctions without endorsing (CK) and (PK)? The most noted alternative to such an explanation – an alternative considered with some sympathy in Fantl and McGrath (2009) – is a pragmatic explanation of the infelicity of utterances such as those of those form ‘I know that p, but there’s a chance that not-p’.²⁹

If (CK) is false, knowledge that p will be compatible with a small epistemic chance, for the subject, that not-p. The natural way to cash out a version of SSI that rejects (CK) is therefore to suggest that the practical environment of the subject determines just how small the epistemic chance that not-p needs to be in order for her to know that p. For example, if it matters a great deal to Mary whether Peter is in the office, then the epistemic chance, for Mary, that Peter is not in the office might have to be *very* small in order for Mary to know that Peter is in the office. But if it matters little to Mary whether Peter is in the office, Mary can know that Peter is in the office compatibly with there being a greater epistemic chance, for Mary, that he isn’t.

²⁹ See Fantl and McGrath (2009, p. 22-23).

Recall the proposal, introduced in the previous chapter (§4), that the reason utterances of the form ‘I know that p, but there’s a chance that not-p’ seem infelicitous is, firstly, that in order for the utterance of the second conjunct – ‘there’s a chance that not-p’ – to be relevant, the chance, for the speaker, that not-p must be large enough to be relevant to the conversation. And, secondly, that if that chance were large enough to be relevant to the conversation, then the utterance of the first conjunct – ‘I know that p’ – would be false.³⁰

One set of problems with that pragmatic proposal, insofar as it was offered in defence of traditional invariantism, centred on the fact that, in high stakes circumstances, a chance could be large enough to be relevant to the conversation without being large enough to preclude knowledge. That would seem to be less of a problem for a view on which, in high stakes circumstances, potentially even a *very* small chance, for the subject, that she is mistaken can translate into a lack of knowledge. Thus, proponents of SSI may well be in a position to avoid such difficulties.³¹

However, the proposal is susceptible to other difficulties, including some based on the objections outlined in the previous chapter. For example, we are told that the infelicity of an utterance by Mary of ‘I know that Peter is in the office, but there’s a chance that he isn’t’ lies in some relevance-based pragmatic effect. If that’s so, it seems that the effect should be cancellable. However, it does not seem felicitous for Mary to utter something along the lines of ‘It’s not really relevant right now, but bear in mind that even though I know that Peter is in the office, there’s a chance I’m mistaken’.

Furthermore, the infelicity of various closely related utterances cannot be explained away in a similar fashion. Suppose that it matters a great deal to John and Mary whether or not Peter is in the office. Unbeknownst to John, Mary has secured some very good evidence that Peter is in the office – she saw him there five minutes ago.

³⁰ This proposal (or something very similar) is found in Dougherty & Rysiew (2009). (Note that Dougherty and Rysiew present the account in an attempt to defend traditional invariantism.)

³¹ See Fantl and McGrath (2009, p. 22-23)

Suppose that John suggests to Mary that Emily might be able to confirm for them whether Peter is in the office, and asks Mary whether Emily might know whether Peter is in the office. Mary replies, ‘Emily knows that Peter is in the office, but there’s a chance, for her, that he isn’t’. If (CK) is false, Mary’s response can easily be true. But it seems infelicitous.

A parallel pragmatic explanation of that infelicity would have it that the reason Mary’s utterance sounds infelicitous is that the relevance of the second conjunct entails the falsity of the first. However, John and Mary are in a high stakes practical environment. It follows that in order for the chance, for Emily, that John isn’t in the office to be relevant to John and Mary’s conversation, it need only be very small. But *Emily* need not be in a high stakes practical environment. Even given SSI then, a very small chance, for Emily, that Peter is not in the office can be perfectly compatible with her knowing that he is in the office. It follows that Mary’s utterance of the second conjunct – ‘there’s a chance, for Emily, that Peter isn’t in the office’ – can easily be relevant compatibly with her utterance of the first conjunct – ‘Emily knows that Peter is in the office’ – being true. The proposed pragmatic explanation does not seem sufficiently general.

Finally, this kind of pragmatic explanation will also fall afoul of the ‘Suppose Test for Contradiction’ (Chapter 3, §4). Imagine that Mary and John are in a high stakes situation and that John utters ‘Suppose Emily knows that Peter is in the office, but there’s a chance, for her, that he isn’t’. Such an utterance seems infelicitous. This strongly suggests that the embedded conjunction expresses a contradiction and that (CK) is true.

The natural response to this objection is to attempt some other explanation of the infelicity of John’s ‘Suppose...’ utterance – one which doesn’t involve claiming that the embedded conjunction expresses a contradiction. Perhaps the most promising candidate appeals to the suggestion that John’s ‘Suppose...’ utterance could be relevant only if the chance, for Emily, that Peter is not in the office is sufficiently large to be relevant to John’s conversation. If that were the case, it might be suggested that the only way John’s utterance could be understood as relevant is if he were

introducing an impossible supposition. That might account for the infelicity of that utterance.

However, for reasons given in the discussion of the preceding objection, even supposing that the chance, for Emily, that Peter is not in the office is sufficiently large to be relevant to John's conversation (which is taking place in a high stakes practical environment), Emily can nevertheless know that Peter is in the office. Emily need simply be in a low stakes practical environment. This proposal for accounting for the infelicity of John's utterance of 'Suppose Emily knows that Peter is in the office, but there's a chance, for her, that he isn't' therefore fails. It is very unclear that there is any other.

The pragmatic account of the infelicity of utterances such as those of the form 'I know that p, but there's a chance that not-p' seems unsuccessful. Is there an alternative? One alternative that I shall briefly consider here is that a salience-sensitive form of SSI is true. On such a view, a subject must rule out those alternatives to p that are salient to her in order to know that p. It was suggested earlier that in claiming that there is a chance, for her, that not-p, a speaker might make salient a possibility in which she appears to have just the same evidence that she currently does for p, but in which p is false. If that does indeed occur, then on a salience-sensitive form of SSI, it will follow that the speaker does not know that p. It might therefore be proposed that the reason utterances of the form 'I know that p, but there's a chance that not-p' typically seem infelicitous is that asserting the second conjunct renders the first conjunct false.

Of course, we have already seen that there are a number of serious difficulties with a salience-sensitive form of SSI. But even setting those aside, this form of SSI will not be of much assistance in the present context, since the proposal does not enable us to account for all of the relevant judgments. The problem is that, as we have already seen, such conjunctive claims seem similarly problematic when couched in the third-person. For instance, suppose that Mary utters, 'Emily knows that Peter is in the office, but there's a chance, for her, that he isn't'. Mary's utterance seems infelicitous, but if (CK) is false, it seems that Mary's response can easily be true.

A parallel explanation of that infelicity would again invoke the idea that asserting the second conjunct renders the first conjunct false, on the grounds that by asserting that there is a chance, for Emily, that Peter is not in the office (i.e. the second conjunct), Mary makes salient some possibility in which Emily has just the same evidence, but in which Peter is not in the office. However, by Mary claiming that there is a chance, for Emily, that Peter is not in the office, she need not make salient *to Emily* (i.e. the subject) some possibility in which Emily has just the same evidence, but in which Peter is not in the office. Emily could be in an entirely different conversation, reflecting on entirely different questions. And if such a possibility is not salient to Emily, there seems no obstacle, even granting a salience-sensitive form of SSI, to Emily being able to know that Peter is in the office. The proposal is therefore insufficiently general.

In the absence of a more promising explanation, the infelicity of utterances such as ‘I know that p, but there’s a chance that not-p’ and ‘I know that p, but it might be that not-p’ represent a serious objection to those advocates of SSI who reject (CK) and (PK). Furthermore, the additional judgments introduced in the course of the preceding objections represent further problems for those proponents of SSI who reject (CK) and (PK). For example, given a rejection of (CK), it seems that the following utterance expresses a consistent supposition: ‘Suppose that Emily knows that Peter is in the office, but that there is a chance, for her, that he isn’t’. That the utterance seems infelicitous therefore represents a general problem for those sympathetic to SSI but not (CK).

The upshot of the discussion of the preceding two sections is that the advocate of SSI is faced with a dilemma which parallels that of the traditional invariantist. There are notable costs associated with endorsing principles (CK) and (PK), but there also seem to be considerable troubles associated with rejecting those principles.

§7. Multi-Premise Closure.

Issues pertaining to (MPC) further intensify that dilemma.

(MPC) If a subject S knows propositions p_1 to p_n and competently deduces q from p_1 to p_n – retaining throughout her knowledge of p_1 to p_n – then S knows q .

Consider first the possibility that the proponent of SSI rejects (MPC). It was argued in the previous chapter that, at least for the traditional invariantist, a rejection of (MPC) entails a rejection of (CK) and (PK). One prominent argument to that effect was the following. Suppose Peter knows a series of propositions p_1 to p_n . Given (CK), it follows that, for each p_i , there is no chance, for Peter, that p_i is false. It seems bizarre, however, that Peter could truly claim, for each of p_1 to p_n , that there's no chance, for him, that he's mistaken *and* truly claim that he believes the conjunction of those propositions on the basis of competent deduction from the conjuncts, but then truly claim that, *nevertheless*, he does not know the conjunction. *How* could he not know the conjunction?³²

If SSI licenses the truth of all these three claims together, that strongly suggests that if the proponent of SSI rejects (MPC) she must reject (CK) as well. Does SSI license the truth of these three claims? If knowledge is merely tied to the practical environment of the subject – if some principle like (1a) is true, for instance – SSI seems to provide little in the way of tools to help resist that conclusion. However, a salience-sensitive form of SSI *might* seem to lessen the difficulties associated with endorsing (CK) but not (MPC). It could be proposed that, in circumstances in which (MPC) fails – and so the relevant conjunction is not in fact known (perhaps because it is very long) – Peter makes salient certain error possibilities by claiming not to know the conjunction. In particular, perhaps he makes salient the possibility in which he appears to have just the same evidence for the conjunction, but in which one of the conjuncts is false. (See Chapter 3, §3 for relevant discussion.) Given a salience-sensitive version of SSI, Peter might then also fail to know one of the conjuncts when he claims not to know the conjunction. If that's right, then Peter might simply be *unable* – even if (CK) holds

³² But couldn't there be a chance, for Peter, that the deduction was not fully competent – and in consequence, perhaps, a chance, for Peter, that the conjunction is false? If that were so, it might explain how Peter could still fail to know the conjunction. In the present context, this source of concern can be eliminated simply by assuming that Peter knows that the deduction was fully competent – since if Peter knows that, then given (CK), it will follow that there is no chance, for Peter, that the deduction was not fully competent. See Chapter 3, §2 for further discussion.

and (MPC) fails – to truly claim, for each conjunct, that there’s no chance, for him, that it’s false and at the same time truly claim that he does not know the conjunction.

However, this manoeuvre involves endorsing the kind of salience sensitive form of SSI that was shown to be so problematic in §5, above. Furthermore, it will do nothing to assuage a similar argument to the effect that a rejection of (MPC) entails a rejection of (CK) – one couched in the third person. Suppose once again that Peter knows a series of propositions p_1 to p_n . Given (CK), it follows that for each p_i there is no chance, for Peter, that p_i is false. It seems equally bizarre, however, that *Mary* could truly claim, for each of p_1 to p_n , that there’s no chance, for Peter, that he’s mistaken *and* truly claim that Peter believes the conjunction of those propositions on the basis of competent deduction from the conjuncts, but then truly claim that, *nevertheless*, Peter does not know the conjunction. Once again, *how* could he not know the conjunction?

In response to this argument, it cannot be suggested that in virtue of Mary claiming that Peter does not know the relevant conjunction, certain possibilities of error will become salient to Peter (i.e. the subject). (And as a result of which Peter loses knowledge of some conjunct when Mary claims not to know the conjunction.) For all that has been said, Peter could be sound asleep! The proponent of SSI therefore seems to be able to do little to resist the force of this argument, which strongly suggests that a rejection of (MPC) entails a rejection of (CK). The situation is similar for the other, closely related arguments, to the effect that a rejection of (MPC) entails a rejection of (CK) that we considered in the previous chapter. In short, it seems that a rejection of (MPC) entails a rejection of (CK), even for the advocate of SSI.

The advocate of SSI who rejects (MPC) will therefore have to find an explanation of the infelicity of utterances of the form ‘I know that p , but there’s a chance that not- p ’ that does not involve appeal to (CK). However, as argued in the previous section, no such explanation appeared to be forthcoming. Furthermore, there are a number of independent considerations that speak against a rejection of (MPC). For example, it seems that were a speaker to claim to know each of a series of propositions p_1 to p_n , but then claim not to know their conjunction, she would be going back on at least one

of her earlier claims to know. This can be readily understood on the assumption that (MPC) holds, but is much harder to understand on the assumption that it doesn't.

If (MPC) is rejected, then the SSI theorist should presumably seek to explain judgments of this kind away. However, it does not seem that SSI supplies much in the way of materials, over and above what is available to the traditional invariantist, that are likely to prove a great deal of assistance in that explanatory project. The reason is that it seems that we can simply suppose that the relevant subject remains in a practical environment in which neither the propositions p_1 to p_n , nor their conjunction, are of practical import. Furthermore, we can also suppose that the subject is not considering any potentially pertinent error possibilities. Perhaps this is most readily achieved by placing the relevant dialogue in the third-person and stipulating that the subject has nothing resting on the truth or falsity of those propositions (or their conjunction) and is not even considering whether or not they are true. The pertinent dialogue would then take something like the following form:

Emily: Does Peter know that he will go to the theatre tomorrow?

Mary: Yes

Emily: Does Peter know that his daughter will be in California next month?

Mary: Yes

Emily: Does Peter know that the temperature in London will not exceed forty degrees centigrade tomorrow?

Mary: Yes

...

Emily: Has Peter competently deduced that he will go to the theatre tomorrow and that your daughter will be in California, and ...?

Mary: Yes

Emily: Does Peter know that he will go to the theatre tomorrow and that his daughter will be in California, and ...?

Mary: No, of course he doesn't know that.

As with the first-person correlate of this dialogue, it seems that in her final admission, Mary is going back on one her earlier knowledge claims. This is obviously difficult to

understand on the assumption that, in such cases, (MPC) fails. It is unclear how the SSI theorist could successfully account for this.

The alternative is for the proponent of SSI is to *endorse* (MPC). It was argued in the previous chapter that, for the traditional invariantist, this entails a commitment to (CK) and (PK). Similar consideration apply in the case of SSI. If (MPC) holds but (CK) fails, then a subject can know some very long conjunction, $p_1 \& p_2 \dots \& p_n$, on the basis of competent deduction from its conjuncts, even though, for each p_i , there is a small epistemic chance, for her, that p_i is false. On any plausible notion of epistemic chance, this will translate into a large epistemic chance, for the subject, that the conjunction is false. But it seems very implausible that a subject can know a proposition even though there is a large epistemic chance, for her, that the proposition is false (Hawthorne 2004a, p. 182-183).³³

The proponent of SSI who endorses (MPC) would therefore appear to be left with those problems that afflict her endorsement of (CK) and (PK). The kind of problems she will face were reviewed in §5, above. But even setting aside the relationship between an endorsement of (MPC) and an endorsement of (CK) and (PK), endorsing (MPC) leaves the SSI theorist with a number of exceedingly troubling results. For example, suppose Peter knows a large series of contingent propositions, r_1 to r_n , concerning the future. The conjunction of those propositions, we may suppose, has a high objective chance of being false. Suppose Peter knows this, but that Peter has also competently deduced the conjunction from the conjuncts. In that case, given (MPC), it follows that Peter knows the conjunction, but also that he knows that the conjunction has a high objective chance of being false (Hawthorne, 2004a, p. 185).

The SSI-machinery may offer *some* comfort here, but not much. In particular, it may be proposed that we recognise that, given the high objective chance that the conjunction is false, there will be some circumstances in which it is not appropriate for Peter to rely on the conjunction in his practical reasoning. It might be proposed that this fact becomes salient when we consider whether Peter knows the conjunction and therefore, given that we assume that knowledge does not come and go simply

³³ See Chapter 3, §5.1, for further discussion.

with changes in practical environment, we judge that Peter does not know the conjunction (cf. Hawthorne 2004a, p. 180 & 182-183).

However, as noted earlier, it is clearly problematic for the advocate of SSI to appeal to the claim that we mistakenly assume that knowledge does not come and go with changes in practical environment. In effect, the theorist is postulating an error to explain an error. This might not be so troubling if the error postulated were one that could be readily explained away (that is, if the error were not *significant*) or if the error postulated seemed to be present in our judgments concerning other apparently practical-environment-sensitive notions. But it is not.

Furthermore, this kind of manoeuvre does little to explain away the force of the arguments which can be used to bolster the initial implausibility of the claim that Peter knows the conjunction but also knows that the conjunction is highly objectively likely to be false. For example, if the subject knows that there's a very high objective likelihood that some proposition *p* is false, it seems that the subject should place only a low degree of confidence in *p*. But it seems absurd to suggest that even though a subject should place only a low degree of confidence in *p*, she nevertheless knows that *p*. This supplies us with a further argument to the effect that Peter does not know the objectively unlikely conjunction. It is far from clear that a similar explanatory strategy could be plausibly deployed to assuage the force of such additional arguments.³⁴

In this section, we have seen that endorsing (MPC) seems to require endorsing (CK) and (PK), and that rejecting (MPC) seems to require rejecting (CK) and (PK), even for the proponent of SSI. More importantly, we have seen that issues surrounding (MPC) compound the dilemma highlighted at the end of the previous section. On the one hand, rejecting (MPC) leads to difficulties over and above those previously seen with a rejection of (CK) and (PK). In particular, a rejection of (MPC) leaves the

³⁴ Another such argument (see Chapter 3, §5.2) rests on the observation that if the subject knows that there's a very high objective likelihood that some proposition *p* is false, it seems that the subject's belief that *p* is not (at all) safe from error. If knowledge requires that a belief be safe from error (or even reasonably safe from error), this supplies us with a further argument that Peter does not know the conjunction. It is not clear that a similar explanatory strategy could be plausibly deployed to assuage the force of this kind of argument either.

proponent of SSI seemingly unable to account for the fact that a speaker who claims to know a series of propositions, but then claims not to know their conjunction, seems to be going back on one her earlier claims to know. On the other hand, endorsing (MPC) leads, as we have just seen, to some extremely counterintuitive results.

§8. Summary.

In this chapter, we have reviewed the prospects for Subject Sensitive Invariantism – a conservative form of anti-intellectualism that retains a commitment to invariantism and to the thesis that we know much of what we take ordinarily ourselves to know. It has been shown that this position is afflicted by a number of serious problems.

In the first instance, it was shown that those forms of SSI that tie knowledge to the subject's practical environment face problems with respect to our judgments concerning third-person cases, counterfactuals linking stakes and knowledge, and knowledge claims concerning pairs of subjects occupying the same epistemic position but different practical environments (§2). These judgments also reveal how at odds knowledge is with other apparently stakes-sensitive notions, such as the need for further inquiry. The most promising attempts to explain away these judgments through some pragmatic, psychological, or other mechanism, were shown to be unsuccessful (§3).

The plausibility of a form of SSI that ties knowledge to the error possibilities salient to the subject was also considered. It was shown that such a position is very problematic (§4). But if such a position is untenable, it seems that the proponent of SSI is entirely unable to account for the HIGH-LOW contrast, as that contrast seems to manifest itself in examples that feature a shift merely in those error possibilities salient to the attributor. In this respect, it seems that SSI is no better off than traditional invariantism.

Finally, it was shown that the proponent of SSI, whether she claims that knowledge is tied to the subject's practical environment *or* the error possibilities salient to the subject (or both), is unable to account for our judgments concerning the relationships between knowledge and chance. In that arena, the advocate of SSI faces a similar

dilemma to the traditional invariantist. On the one hand, endorsing the connection between knowledge and epistemic chance expressed in (CK) leads to serious problems with respect to those judgments which purport to show that, for much of what we ordinarily take ourselves to know, there is a chance, for us, that we are mistaken. The endorsement of (CK) also entrains a commitment to (MPC) which brings even more severe problems – for instance, that a subject can know a proposition, but also know that there is a high objective chance that that proposition is false. On the other hand, rejecting (CK) and (MPC) leaves the advocate of SSI unable to account for those judgments that seem to support (CK) – such as the infelicity of utterances of the form ‘I know that p, but there’s a chance that not-p’ – and unable to account for other judgments that lend support to (MPC).

These difficulties are unquestionably severe. Of course, if all other possible accounts are in an equally dire situation, these concerns might seem less troubling. In the next and final chapter, I argue that all views are *not* in such an equally dire situation.

Chapter 5: Scepticism and Contextualism.

The severity of the problems encountered by non-sceptical forms of invariantism suggests that it is worth exploring the possibility that some form of *sceptical* invariantism is correct. Sceptical invariantists endorse an invariantist semantics for ‘knows’, but – for various possible reasons – reject the assumption that we know much of what we ordinarily claim to know. The form of sceptical invariantism that I wish to explore in the present work is one according to which in order to know some proposition *p*, a subject must occupy a *very* strong epistemic position with respect to *p*, but that this is rarely, if ever, the case.

Philosophers often balk at endorsing forms of scepticism. It will therefore be useful to begin by stressing that there is a strong sense in which this *particular* form of sceptical invariantism isn’t especially radical. First, notice that it’s not as though it takes much prodding to get a good number of ordinary speakers to admit that ‘you never really know anything’. Furthermore, the sceptical invariantist (at least of this stripe) is not, at least *prima facie*, committed to such claims as that we do not occupy a strong epistemic position with respect to many of our beliefs, that those beliefs are not justified, or that those beliefs are not appropriate. In other words, it’s not as though the position is entirely anathema to ordinary thinking and not as though, at least *prima facie*, it spells trouble for the truth of many important aspects of our epistemic discourse (our use of ‘knows’ aside), or the appropriacy of our belief forming practices.

Of course, the sceptical invariantist faces an obvious – and pressing – challenge. If we know very little of what we claim (or at least, that we *prima facie* appear to claim) to know, it can seem like we are engaged in large scale error in our use of ‘knows’. There is therefore an especially pressing burden upon the sceptical invariantist to make sense of our everyday use of ‘knows’. If this challenge cannot be met, then whatever other benefits sceptical invariantism might offer, it seems that the theory should be rejected.

The primary goal of this chapter is to argue – at least tentatively – for the conclusion that sceptical invariantism represents ‘The Best Invariantism’, in the sense that it is

superior to both traditional and subject sensitive invariantism. But this goal is not sought without the suspicion that sceptical invariantism might be the correct theory. Whilst a full comparison of the merits of sceptical invariantism to the various forms of contextualism is beyond the scope of the present work, a secondary goal of this chapter is to suggest – in at least a preliminary way – how the sceptical invariantist might best attempt to show the superiority of her view to that of the standard contextualist.

In the first section of this chapter, I outline some of the central benefits of sceptical invariantism, drawing on the phenomena we have encountered so far. I then consider some proposals for how to make sense of our everyday use of ‘knows’, and some further objections. In the final section, I consider the relationship between sceptical invariantism and contextualism.

§1. Initial Advantages.

In this section, I shall sketch some initial advantages of sceptical invariantism, and the contours of the kind of sceptical invariantism that I take to be the most plausible. First, consider some pertinent judgments relating knowledge and inquiry, and knowledge and appropriate practical reasoning, that we encountered in previous chapters. In particular, recall that utterances of the following typically seem infelicitous, ‘I know that p, but I must inquire further’ and ‘I know that p, but I can’t rely on p in my practical reasoning’. The infelicitousness of such utterances suggests that something like the following principle is true:

- (1) a. If S knows that p, then S can appropriately rely on p in her practical reasoning.

This principle presented a problem for the traditional invariantist because, in certain practical environments, a subject seems to have to occupy a *very* strong epistemic position with respect to a proposition p in order to appropriately rely on p in her practical reasoning. In conjunction with (1a), and the traditional invariantist’s commitment to intellectualism, this led to the conclusion that we know very little of what we ordinarily take ourselves to know. In contrast, the subject sensitive

invariantist was able to embrace (1a) whilst avoiding this sceptical conclusion – but only, as we saw in the previous chapter, at the (expensive) cost of intellectualism. However, the sceptical conclusion represents no problem for the sceptical invariantist; she can therefore embrace the attractive thesis of intellectualism *and* (1a). That is surely an advantage.

There are further advantages in the arena of knowledge and chance. The natural explanation of the infelicity of utterances of the form ‘I know that p, but there’s a chance that not-p’ and ‘I know that p, but it might be that not-p’ appeals to the truth of (CK) and (PK).

(CK) If S knows that p, then there’s no epistemic chance, for S, that not-p.

(PK) If S knows that p, then it’s epistemically impossible, for S, that not-p.

The chief objections to traditional invariantist and subject sensitive invariantist endorsement of (CK) and (PK) stemmed from a number of arguments and judgments suggesting that, for much (or perhaps even all) of what we ordinarily take ourselves to know, there’s an epistemic chance, for each of us, that we are wrong. In conjunction with (CK) and (PK), these arguments and judgments seemed to force the conclusion that we know very little. But once again, such a sceptical conclusion represents no problem for the sceptical invariantist. It will be instructive to briefly review two of the relevant arguments.

One simple argument rested on our judgments regarding epistemic uses of expressions like ‘chance’ and ‘might’. For example, even supposing that Mary occupies a strong epistemic position with respect to the proposition that Peter is in the office – a reliable colleague told her that Peter is in the office and she saw Peter’s hat in the office hallway – it still seems true for her to utter ‘There’s a chance that Peter is not in the office’. Indeed, it can seem true to utter, for virtually any belief that p, ‘There’s a chance that not-p’. Such judgments suggest that there’s an epistemic chance, for each of us, that we are wrong, for many of those beliefs we ordinarily take to amount to knowledge.

A second argument – based on considerations pertaining to objective chance – stemmed from two very plausible claims. The first was that if a subject S knows that the objective chance that some event e occurs at t is greater than 0, then S's credence that e will occur at t should be greater than 0.¹ The second claim was that principle (CC), expressing a very plausible connection between rational credence/confidence and epistemic chance, is true.

(CC) If it is not rational for S to assign a credence of 0 to p , then the epistemic chance, for S, that p is greater than 0.²

Contemporary physics suggests that for all (or nearly all) contingent propositions about the future, there's an objective chance that they are false. If we can know what our best physics seems to teach us, then these two highly plausible claims suggest that there will be an epistemic chance, for each of us (or at least those of us acquainted with contemporary physical theory), that we are wrong about all contingent propositions concerning the future.

It seems rather implausible for proponents of traditional invariantism and SSI to reject the idea that our best physics can furnish us with knowledge about the world. Proponents of those theories are therefore pressured to embrace the conclusion that there will be an epistemic chance, for each of us, that we are wrong about all (or nearly all) contingent propositions concerning the future.

In conjunction with (CK), these arguments appear to lead to the sceptical conclusion that we know very little. This consequence obviously represents no objection to the sceptical invariantist. In consequence, it might seem that the sceptical invariantist should straightforwardly embrace these arguments alongside principle (CK). But is that the best reaction for her?

¹ As we saw in Chapter 3, §2.1, this principle should come with certain qualifications – for instance, a clause barring the presence of God-like access to the future. These qualifications are irrelevant to the present discussion.

² The sense of rationality intended in (CC) is epistemic, rather than practical, in character. Suppose that Peter is threatened with a horrible death if he places full confidence in the truth of some proposition p . In that case, even supposing that there's no chance, for Peter, that not- p , there is an obvious sense in which he nonetheless ought not to place full confidence in p . This is not the kind of rationality intended in (CC).

It seems that the appropriate reaction for the sceptical invariantist to the first argument is indeed to straightforwardly embrace it and its conclusion (alongside principle (CK)). However, the situation is slightly more complicated with respect to the second argument. The reason is that the sceptical invariantist – at least of the stripe we are interested in – is unlikely to allow that our best physics can furnish us with *knowledge* about the world. She may well therefore reject the idea that people ever *do* know that the objective chance of some event is greater than 0, because even acquaintance with contemporary physics cannot put us in the kind of very strong epistemic position required to know such claims. For the sceptical invariantist, the second argument will not get off the ground.

However, the argument can be readily resuscitated, since the following claim also seems plausible: if S occupies a strong epistemic position – that typified by very good, all-things-considered, evidence – with respect to the proposition that the objective chance that event *e* occurs at *t* is greater than 0, then S's credence that *e* will occur at *t* should be greater than 0.³ The sceptical invariantist – at least of the stripe we are interested in – sensibly accepts that our best physics can provide us with very good, all-things-considered evidence concerning matters about the world. (In this respect, she is just like the traditional invariantist and the subject sensitive invariantist.) Given principle (CC), it will therefore follow that for all (or nearly all) those contingent propositions we take ourselves to know about the future, there is an epistemic chance, for us (or at least for those of us acquainted with contemporary physics), that we are wrong. It seems that even granting that our best physics cannot deliver knowledge, considerations pertaining to objective chance nevertheless suggest that there is an epistemic chance, for us (or at least for those of us acquainted with contemporary physics), that we are wrong about nearly all contingent propositions concerning the future.

As noted above, the sceptical invariantist, unlike the proponent of traditional invariantism or SSI, can happily grant this conclusion whilst maintaining (CK). The sceptical invariantist is therefore in the enviable position of being able to explain the

³ Once again, this principle should come with certain qualifications – e.g. a clause barring the presence of God-like access to the future. Such qualifications are irrelevant here.

infelicity of conjunctions such as those of the form ‘I know that p , but there’s a chance that not- p ’ by endorsing (CK) *and* respect those judgments and arguments that purport to show that, for much of what we ordinarily take ourselves to know, there’s a chance, for us, that we’re wrong. (Parallel remarks apply to conjunctions of the form ‘I know that p , but it might be that not- p ’ and principle (PK).) The dilemma that plagued traditional invariantism in Chapter 3 (and SSI in Chapter 4) therefore evaporates. This is surely also an advantage.

The situation is similar with respect to (MPC):

(MPC) If a subject S knows propositions p_1 to p_n and competently deduces q from p_1 to p_n – retaining throughout her knowledge of p_1 to p_n – then S knows q .

It seems that, as with (CK) and (PK), the sceptical invariantist can embrace (MPC) without facing the kind of problems that upset traditional invariantist and subject sensitive invariantist endorsement of that principle.⁴ For example, recall that one apparent consequence of a traditional invariantist or subject sensitive invariantist endorsement of (MPC) was that a subject could know some long conjunction, but also know that the conjunction is exceedingly objectively unlikely to be true. This result came about because, on those views, it seems that a subject will be able to know each of a large series of propositions p_1 to p_n , where each p_i has a very small objective chance of being false. But the sceptical invariantist can block the troublesome result simply by maintaining that we never know each proposition p_i . The reason is simple: it is open to her to claim that if a proposition has an objective chance of falsehood

⁴ Of course, considerations relating to a chance of error in the performance of even a competent deduction may well show that, as currently formulated, (MPC) is false for *everyone*, including the sceptical invariantist. (See Lasonen-Aarnio (2008) for relevant discussion.) These issues are not central to the discussion in hand. One way to see this is to notice that concerns based around this issue can be eliminated by replacing (MPC) with something like (MPC)*.

(MPC)* If a subject S knows propositions p_1 to p_n and competently deduces q from p_1 to p_n – in such a way that there is no epistemic chance, for S , that not- q traceable to the performance of the deduction, and retaining throughout her knowledge of p_1 to p_n – then S knows q .

(MPC)* can still be exploited to create serious trouble for traditional invariantism and SSI (it will allow those that differ from us *only* in the respect that they are ‘infallible reasoners’ to know some proposition p but also know that p is highly objectively unlikely to be true) but not sceptical invariantism.

greater than 0, then – barring some God-like access to the future – that proposition cannot be known.

Let's sum up these advantages. The sceptical invariantist is able to respect pertinent judgments concerning the relationship between knowledge and practical reasoning – such as that utterances of 'I know that p, but we must inquire further (into whether p)' seem infelicitous – by endorsing principle (1a). She is also able to do this whilst endorsing the attractive thesis of intellectualism, and thus avoiding the problems – highlighted in Chapter 4 – that attend invariantist rejection of that (attractive) thesis. In addition, the sceptical invariantist is able to endorse apparently well-supported and attractive principles (CK), (PK) and (MPC), *and* respect those judgments and arguments that seem to favour the conclusion, that, for much of what we ordinarily take ourselves to know, there's a chance, for each of us, that we are wrong. She is therefore able to avoid the pressing dilemmas that beset traditional and subjective sensitive invariantists in regard to those principles and those judgments/arguments.

These represent clear and undeniable advantages of sceptical invariantism. Of course, they come at a notable *prima facie* cost. In committing to the kinds of principles and arguments outlined above, we have a clearer conception of just what requirements the sceptical invariantist places on knowledge, and it's clear that they are *tough*. For example, in accepting the conclusions of the argument centred around objective chance and principle (CC), it will follow, on the assumption that our world is 'chancy', that no matter how much evidence we gather – unless perhaps via some clairvoyance-like method – we will be unable to know contingent propositions about the future. Similarly, given (1a) and intellectualism, it will follow that a subject can know some proposition p only if, *no matter what her practical environment*, she can rely on p in her practical reasoning.⁵ It seems that in order to win the advantages outlined above, the sceptical invariantist is going to have to be a fairly hard-nosed sceptic – one who claims that all (or nearly all) of those beliefs we hold do not amount to knowledge.

⁵ The characterisation is of course a little loose, since the proposition could itself be about features of the subject's practical environment.

The sceptical invariantist is therefore under strong pressure to make sense of our everyday ‘knowledge’-ascribing (and denying) practices. It is to that task that I now turn.

§2. The HIGH-LOW Contrast.

As noted in the introduction, the central challenge for the sceptical invariantist is to make sense of our ordinary use of ‘knows’. In what follows, I shall explore the proposal, familiar from Unger (1979, 1984) and, more recently, Schaffer (2004) and Davis (2007), that this is best achieved by appeal to some broadly pragmatic mechanism. The proposal can be brought out by considering the Bank Case once more. Recall that in both LOW and HIGH, DeRose occupies a modest epistemic position with respect to the proposition that the bank will be open on Saturday. It seems (at least initially) true and appropriate for DeRose to self-ascribe knowledge of that proposition in LOW, but true and appropriate for him to deny himself knowledge of that proposition in HIGH.

The sceptical invariantist claims that DeRose does not know that the bank will be open in HIGH or LOW. The most pressing challenge is therefore to explain away our judgment to the effect that it seems true and appropriate for DeRose to self-ascribe knowledge in LOW. On a pragmatic treatment, this will be because DeRose’s false utterance in LOW pragmatically conveys a truth – perhaps that he is sufficiently close to knowing that the bank will be open for present purposes, or that he is in a strong enough epistemic position to rely on the proposition that the bank will be open in his practical reasoning.

In Gricean terms, it is quite straightforward to offer an explanation of the required kind. In uttering ‘I know that the bank will be open’ in LOW, DeRose is uttering something that cannot plausibly be true. He is therefore in apparent violation of the maxim of Quality.⁶ But on the assumption that DeRose is being co-operative, conversational participants will assume that he is trying to convey something other

⁶ Grice (1989, p. 27) elucidates the maxim of Quality as consisting of a ‘supermaxim’ – ‘Try to make your contribution one that is true’ – and two submaxims: (i) Don’t say that which you believe to be false, and (ii) Don’t say that for which you lack adequate evidence.

than what his utterance semantically expresses – perhaps that he is close enough to knowing that the bank will be open for present purposes. DeRose’s utterance of ‘I know that the bank will be open’ in LOW seems true and appropriate then, because (although false) it pragmatically conveys a truth.

(The reader will also notice that on the sceptical invariantist account, DeRose’s utterance of ‘I don’t know that the bank will be open’ in HIGH is unlikely to convey simply what it semantically expresses. On the assumption that sceptical invariantism is true, it is presumably trivial that DeRose does not know that the bank will be open. Thus, although his knowledge denial in HIGH is true, it is plausible that it will pragmatically convey the truth that DeRose is not close enough to knowing that the bank will be open for present (contextually indicated) purposes (or something along these lines).

The relevant Gricean explanation of this fact might be centered around the maxim of Quantity – in particular, the injunction to one make one’s conversational contribution as informative as is required.⁷ DeRose’s utterance of ‘I don’t know that the bank will be open’ in HIGH is trivially true and therefore (superficially at least) altogether uninformative. He is therefore in (superficial) violation of the maxim of Quantity. But conversational participants will assume that DeRose is striving to be co-operative (and relevant). In consequence, his utterance of ‘I don’t know that the bank will be open’ in HIGH will pragmatically convey (the truth) that he is not close enough to knowing that the bank will be open for present purposes (or something along those lines). Of course, that the theory has the consequence that DeRose’s knowledge denial in HIGH is not only true but *also* conveys a truth does little to problematise a sceptical invariantist explanation of why that utterance might seem true and appropriate. Indeed, the issues surrounding DeRose’s knowledge denial in HIGH are not central to the discussion to follow. I shall therefore keep the focus (at least in the main) on DeRose’s positive knowledge ascription in LOW.)

Often, sceptical invariantists will not couch their proposal in explicitly Gricean terms. Instead, the pragmatic effect will be described as being of a certain *kind* – for

⁷ See Grice (1989, esp. p. 26).

instance, a species of exaggeration, metaphor, or loose use – and we will then be provided with some examples of expressions of the same kind. This strategy has some advantages. On the one hand, it doesn't tie the sceptical invariantist to the Gricean framework, and the attendant Gricean explanation. But on the other hand, it still enables us to evaluate the proposal fairly effectively. For example, we can compare the behaviour of 'knows' with the behaviour of those expressions that it is suggested are of the same kind. We can even raise questions about whether a pragmatic treatment is appropriate for *any* of the relevant expressions. For instance, if a 'loose use' account of various expressions, including 'knows', is offered, it might be objected that loose-use is a semantic, rather than a pragmatic, phenomenon.

Of course, pursuing the Grice-independent strategy may leave us without a complete picture of the pragmatic mechanism in play, and its relationship to other pragmatic mechanisms. But for present purposes, this is probably not so important. To compare: it is not fair to criticise the contextualist about 'knows' because she hasn't provided a full account of context-sensitivity in natural language.

Let's consider one prominent implementation of the Grice-independent strategy. It is to my mind one of the more plausible sceptical invariantist accounts. It will also introduce some expressions that will be useful in the discussion to follow. The account is found in Davis (2007). Davis favours a 'loose use' proposal, and he likens 'knows' to expressions such as 'all gone', 'completely flat', 'two hours', and 'two meters'.

First, notice that a similar contrast to the HIGH-LOW contrast can be found with 'all gone'.

Case I. Peter picks up the vodka bottle and sees that someone has poured the last glass; there are only the remnants left inside. Suppose Mary is about to go to shopping. She asks Peter 'Is the vodka all gone?'. Peter replies 'Yes, the vodka's all gone. We need to buy some more'. His utterance seems true and appropriate.⁸

⁸ Davis (2007) suggests our judgments to the effect that the relevant utterances in examples like *Case I and Case II* are true and appropriate are just as strong as our judgments to the effect that DeRose's utterances of 'I know that the bank will be open' in LOW, and 'I don't

Case II. As in *Case I*, Peter has picked up the vodka bottle and seen that there are only the remnants left inside. Peter's son approaches him and explains that he needs some traces of alcohol for his science experiment. He asks, 'Is the vodka all gone?'. Peter replies, 'No, not quite. There's probably enough for you.' His utterance seems true and appropriate.

Davis proposes a pragmatic account of this contrast. He suggests that an utterance of a sentence like 'The X is all gone' is true just in case there is *absolutely none* of the X remaining (or something along those lines). If that's right, Peter's utterance of 'The vodka's all gone' in *Case I* is false. Davis proposes that the reason Peter's false utterance is *Case I* seems true and appropriate is that the speaker is using 'all gone' *loosely* to convey that it is sufficiently close to all gone for present conversational purposes.

Davis offers a parallel explanation for why DeRose's utterance of 'I know that the bank will be open' in LOW seems true and appropriate. He proposes that DeRose is able to use 'I know that the bank will be open' loosely to convey that he is sufficiently close to knowing that the bank will be open for present, contextually indicated, purposes. It comprises a central aspect of his defence of the 'loose use' account of our Bank Case judgments that 'knows' behaves in pertinent respects like expressions such as 'all gone'. (Davis suggests that there exists a broad range of language items that we use loosely. For instance, he suggests that loose use leads to similar 'HIGH-LOW' contrasts with expressions of measurement like 'two hours' and 'two meters', expressions of epistemic possibility like 'impossible', and expressions featuring certain adjectives modified by 'completely', such as 'completely flat' and 'completely empty'.)

It is of course possible to pursue a Gricean explanation of the pragmatic effect putatively present in examples like *Case I*. One possible such explanation would appeal once again to the maxim of Quality. In *Case I*, Peter has seen that there are remnants left in the bottle, and is thus aware that the vodka is not (strictly speaking)

know that the bank will be open' in HIGH, are true and appropriate. If you disagree in this particular instance, see Davis (2007) for a raft of further examples.

all gone. Further, Mary recognises that Peter is very unlikely to have confirmed (or even to have attempted to confirm) that it is (strictly speaking) all gone.⁹ It might therefore seem to us (and to Mary) that Peter is in superficial violation of the maxim of Quality.¹⁰ However, on the assumption that Peter is being co-operative (and relevant), it is natural for us (and for Mary) to take him to be conveying that the vodka is sufficiently close to being all gone for present (shopping-related) purposes. This could account for why his utterance seems true and appropriate.

Davis is not overly sympathetic to the Gricean picture.¹¹ But on the assumption that the relevant phenomena associated with uses of ‘all gone’ is broadly pragmatic rather than semantic in character, these considerations suggest that those sympathetic to a Gricean picture concerning ‘knows’ might also be *prima facie* reasonable in defending their proposal by appeal to our judgments concerning expressions like ‘all gone’.

In what follows, I shall largely remain neutral on the issues of which particular pragmatic proposal might be correct – gesturing at the relevant pragmatic mechanisms, exploring the similarities of ‘knows’ to other expressions, and responding to some objections.

§3. Objections I: Cancellability, Retraction, Disagreement.

How does the sceptical invariantist’s pragmatic proposal fare with respect to the objections raised in Chapter 2 against a *traditional* invariantist pragmatic proposal? Perhaps the most central problems encountered there concerned disagreement, retraction, and cancellability. It is immediately noticeable that the pragmatic effect the sceptical invariantist alleges is associated with knowledge ascriptions does not seem to be cancellable either: it is not felicitous for DeRose, in LOW, to utter such sentences as ‘I know that the bank will be open, but I’m not close enough to knowing that it will

⁹ Even if Peter has thrown out the bottle, has he checked that there are no small splashes in the drinks cabinet?

¹⁰ As remarked in fn. 6, Grice (1989, p. 27) elucidates the maxim of Quality as consisting of a ‘supermaxim’ – ‘Try to make your contribution one that is true’ – and two submaxims: (i) Don’t say that which you believe to be false, and (ii) Don’t say that for which you lack adequate evidence. Either or both submaxims might be in play in this explanation.

¹¹ See Davis (2007, p. 412-3, and 1998).

be open for present purposes'. However, this presents no objection to the sceptical invariantist, since she has a ready explanation of *why* there is such a failure of cancellability – viz. the cancellation attempt expresses a contradiction. The sceptical invariantist is therefore able to avoid a central problem that plagued traditional invariantist (and subject sensitive invariantist) pragmatic manouvering.¹²

The situation with respect to our judgments concerning retraction and disagreement is somewhat more complicated. Let's take the issue of retraction first. Recall the retraction judgments that seemed to create trouble for a traditional invariantist pragmatic proposal.

Suppose that, earlier today (Friday), DeRose was in the situation described in LOW, but that he now finds himself in a situation resembling HIGH – perhaps a friend who has a very important cheque coming due approaches DeRose, explains his situation, raises the possibility that the bank has recently changed its hours, and asks DeRose 'Do you know that the bank will be open?'. Suppose DeRose responds by uttering 'I don't know that the bank will be open'. His wife then questions him 'Didn't you say earlier "I know that the bank will be open"?'.

In that case, it was suggested that it seems true and appropriate for DeRose to utter 'I spoke falsely/what I said earlier was false' and false and inappropriate to utter 'I spoke truly/what I said earlier was true'. It was also suggested that it seems appropriate for DeRose to utter 'I take that back'.

Such judgments are rather mysterious on the assumption that *traditional* invariantism is true because DeRose's utterance of 'I know that the bank will be open' in LOW has all of the following features: it (a) is true, (b) straightforwardly conveys the truth that it semantically expresses, and (c) does not convey anything false. It was also observed

¹² Notice that the pragmatic effect allegedly associated with DeRose's knowledge denial in HIGH – viz. that that utterance conveys that DeRose is not close enough to knowing that the bank will be open for present purposes – *is* cancellable. For instance, an utterance by DeRose in HIGH of 'I don't know that the bank will be open, but I am close enough to knowing that it will be for present purposes' seems felicitous (although false). See Davis (2007, p. 411) for relevant discussion.

that we do not seem to see the same sorts of judgments in parallel cases featuring other pragmatic phenomena (both ‘controversial’ and ‘uncontroversial’ in character).

According to the sceptical invariantist, DeRose’s utterance of ‘I know that the bank will be open’ in LOW is *false*. It therefore lacks property (a). Furthermore, despite being false, DeRose’s utterance seems true and appropriate. This suggests that DeRose’s knowledge ascription in LOW does *not* convey what it semantically expresses, but instead conveys something else (something true). It therefore also lacks property (b). It follows that a retraction-based argument cannot be pushed against the sceptical invariantist – at least not by appeal to the same judgments and on the same grounds. Perhaps more importantly, given the fact that what DeRose semantically expressed by his utterance of ‘I know that the bank will be open’ in LOW is *false*, it might seem wholly unsurprising that it seems true and appropriate for DeRose to respond to his wife’s query, in the context resembling HIGH, with ‘What I said earlier is false’. This suggests that our retraction judgments represent no serious objection to the sceptical invariantist. (I explore more fully the relationship between our retraction judgments and sceptical invariantism in §3.2.)

Our third problem concerned disagreement. In this instance, it *does* seem that if the judgments reported in Chapter 2 are the ones we have, there exists a difficulty with the sceptical invariantist position (or at least a pragmatic defence of that position).

Suppose Emily is in a context similar to LOW in which the stakes are low and relatively distant error possibilities are not being considered. Emily overhears the conversation between DeRose and his wife in LOW, and utters ‘Keith knows that the bank will be open’.

Now suppose that Tom is also apprised of the nature of DeRose’s epistemic position with respect to the proposition that the bank will be open tomorrow, but that Tom occupies a different context that is similar to HIGH. There is a lot at stake for Tom in the question of whether the bank will be open tomorrow, and he is considering the possibility that the bank has changed its hours in the past two weeks. Tom utters ‘Keith doesn’t know that the bank will be open’.

It might seem that Emily and Tom disagree. In the same vein, suppose that Kurt is acquainted with Emily and Tom's respective situations and overhears both of their utterances. It might seem true and appropriate for Kurt to report the situation with the following: 'Emily and Tom disagree about whether Keith knows that the bank will be open'.

In parallel examples featuring expressions like 'all gone', there is *no* strong sense of disagreement. Recall *Case I* once again.

Case I. Peter picks up the vodka bottle and sees that someone has poured the last glass; there are only the remnants left inside. Suppose Mary is about to go shopping. She asks Peter 'Is the vodka all gone?'. Peter replies 'Yes, the vodka's all gone. We need to buy some more'.

Suppose that Sarah, Peter's daughter, has also looked at the vodka bottle, and seen that there are only some remnants inside. Peter's son approaches Sarah and explains that he needs some traces of alcohol for his science experiment. He asks, 'Is the vodka all gone?'. Sarah replies, 'No, not quite. There's probably enough for you.'

It is really not clear that Peter and Sarah *disagree*. Furthermore, it is not clear that it would be appropriate for some third-party acquainted with Peter's and Sarah's respective situations to report them as disagreeing; an utterance of 'Peter and Sarah disagree about whether the vodka is all gone' is not clearly true or appropriate.

Indeed, if a pragmatic account of the phenomena associated with 'all gone' is on the right lines, it should probably be unsurprising that there is not much sense of disagreement here. The reason is that on such an account, Peter and Sarah are using their utterances to convey different and perfectly compatible propositions. (Peter's utterance conveys that the vodka *is* sufficiently close to all gone for *his conversational purposes*. Sarah's utterance conveys that that the vodka is *not strictly speaking* all gone, or perhaps that it is not sufficiently close to all gone for her conversational purposes.) Similar remarks would seem to apply to a pragmatic account of the analogous Bank Case judgments. Our disagreement judgments might therefore seem to constitute a strong objection to such an account.

The sceptical invariantist appears able to avoid the problems associated with cancellability and retraction, but to face difficulty with respect to our disagreement judgments. In the next section (§3.1), I wish to cast doubt on this last claim. In §3.2, I revisit our retraction judgments. I then move on (§5-8) to consider further objections to sceptical invariantism.

§3.1 Disagreement.

Is there really a contrast between our disagreement judgments concerning cases featuring expressions like ‘all gone’ and cases featuring ‘knows’? It might seem as though there is in fact *not* much sense of disagreement between Tom and Emily in the relevant extension of the Bank Case, and that it is rather unclear that Kurt’s disagreement report concerning them would be true and appropriate. I am certainly inclined to such a reaction. It is also supported by the comments of others in the literature.

For instance, DeRose (2009, p. 157-159) suggests that once we properly set the cases up, there is little sense of *contradiction* between speakers like Emily and Tom.¹³ This supports the conclusion that we would judge that there’s little sense of disagreement between them. Indeed, discussion lending support to the idea that there we would not judge that there is disagreement between speakers like Emily and Tom can be found ever before judgments about disagreement gained prominence in the contextualism debate. For instance, Stephen Schiffer, an early *critic* of contextualism (the view that the presence of disagreement has traditionally been thought to present problems for), writes:

If you claim to know where your car is, and someone challenges you, 'But how can you be sure it wasn't towed or stolen?', you'll merely get impatient at the questioner's obtuseness: it ought to have been mutual knowledge between you that you were speaking casually. (Schiffer 1996, p. 328).

¹³ Though see fn. 44 below.

If it really is mutual knowledge that the conversational participants are speaking ‘casually’ in ‘LOW’ contexts – those contexts marked by low-stakes conversational purposes, in which more distant error possibilities are not being taken seriously –, that lends support to the claim that we would *not* judge that those participants disagreed with conversational participants in the stricter ‘HIGH’ contexts – those marked by high-stakes conversational purposes and/or in which more distant error possibilities are being taken seriously. (Even on the supposition that those in the LOW context have uttered, regarding a subject S and proposition p, ‘S knows p’ and that those in the HIGH context have uttered, regarding that same subject and proposition (and at the same time), ‘S does not know p’.) It might seem rather surprising if there was much sense of disagreement in cases in which one speaker utters a sentence ‘A’ in one context and another speaker utters ‘Not-A’ in a different context, but one speaker is speaking casually and the other is not.

These considerations suggest that we don’t judge in the way that it has thus far been portrayed that we do. This is surely good news for the sceptical invariantist. But it does raise the question *why* some individuals might have been tempted to claim we judge that there *is* disagreement in the relevant cases. There are a number of possible explanations; I do not intend the proposals below to be exhaustive.¹⁴

(i) One proposal is that we have been focusing on poorly described cases.¹⁵ On the sceptical invariantist’s (pragmatic) account, Tom’s knowledge denial will convey that, for purposes like his own, he is *not* close enough to knowing that the bank will be open, and Emily’s knowledge ascription will convey that, for purposes like hers, DeRose *is* close enough to knowing that the bank will be open. If not enough has been done to focus attention on the substantial difference between the relevant purposes, that might lead us to (perhaps implicitly) take it that the speakers are conveying

¹⁴ There may not be much of an explanatory burden for the sceptical invariantist here. I know of no explicit suggestion (in the literature) that we would judge in this way *in cases like that of Emily and Tom*, though it seems to be implied that we would judge this way in MacFarlane (2007, p. 19-21). The suggestion that we judge that the relevant parties do *not* disagree in these cases is strongly suggested by what is said in DeRose (2009, p. 157-159), and is implied by what is said in Schiffer (1996, p. 328). (Notice that the situation contrasts with our retraction judgments: these have been *much* more widely reported in the relevant cases as being as they have been thus far been portrayed to be.)

¹⁵ See DeRose (2009, esp. p. 158-159).

incompatible propositions.^{16,17} If we are (implicitly or otherwise) assuming that the two speakers are conveying incompatible propositions, it would then be unsurprising if we were tempted to judge that they disagree.

(ii) A rather different alternative. This option draws on the plausible idea that there is more than one sense of ‘disagreement’. It seems acceptable for the sceptical invariantist to maintain that the cases make salient the fact that Tom disagrees *with what Emily said*, where what is ‘said’ is understood as referring to what is semantically expressed by her utterance. Tom has uttered ‘DeRose does not know that the bank will be open’ in his context. His utterance semantically expresses the proposition that DeRose does not know that the bank will be open, and conveys, for reasons similar to those that we have been exploring, that DeRose is not close enough to knowing that the bank will be open for (Tom’s) present purposes. Tom has therefore uttered a sentence, in his context, that semantically expresses, and also that conveys, something that is incompatible with the proposition that DeRose (strictly speaking) knows that the bank will be open. He has therefore uttered a sentence that semantically expresses, and that conveys, something that is incompatible with what Emily said. (This is all perfectly acceptable from the sceptical invariantist’s viewpoint – the sceptical invariantist will just maintain that Emily conveyed something other than what she ‘said’.)

It might therefore be proposed that whilst Tom and Emily do not disagree – there is no genuine interpersonal disagreement between them – some individuals might have been tempted to judge that there *is* disagreement between them in virtue of a failure to properly distinguish the (true) claim that Tom disagrees with what Emily said, and the

¹⁶ Similarly, to the extent that we neglect to take into account the speakers’ purposes, it may be that we just (presumably implicitly) take them to be conveying something similar purpose-wise – perhaps that DeRose is and is not (respectively) close enough to knowing that the bank is open for common practical purposes.

¹⁷ In this vein, it might be suggested that there is a risk associated with case descriptions that are rather *abstract* in nature. A description (like the one given earlier) featuring such language as ‘Tom occupies a context in which distant error possibilities are being considered and the stakes are high, and Emily occupies a context in which such possibilities are not being considered and the stakes are low’ is much more abstract in nature than the Bank Case descriptions quoted in the Introduction. It could be that this leads to a failure to fully envisage the relevant contexts, and thus to a failure to recognise the fact that, given their different conversational purposes, speakers like Tom and Emily really do convey different (and compatible) things by their respective knowledge ascriptions and denials.

(false) claim that Tom and Emily disagree.¹⁸ (On this picture, the absence of genuine interpersonal disagreement will be tied up with the fact that Tom and Emily *conveyed* perfectly compatible propositions.) This rather different proposal is thus that the differing judgments in regard to disagreement are (relatively harmless) reflections of the fact that there is more than one sense of ‘disagreement’.

Let’s take stock. It seems that our disagreement judgments might not have the quality suggested in Chapter 2. This suggests that the sceptical invariantist is not so poorly placed with respect to those judgments as may have first appeared. Furthermore, to the extent that there is dispute about the relevant judgments, it seems that the sceptical invariantist might well be in a position to explain its presence.¹⁹ This reinforces the idea that the sceptical invariantist stands in no bad position with respect to our disagreement judgments.

However, the above discussion raises important questions about the critique presented in Chapter 2. In particular: does this more complicated outlook in re our disagreement judgments undermine the disagreement-based objection to (pragmatic defences of) traditional invariantism presented in Chapter 2? I address that issue in §4. First, it will be prudent to look once more to our *retraction* judgments, for they have also been the subject of a somewhat similar dispute.

§3.2 Retraction.

It has been repeatedly suggested that our retraction judgments are as they have been portrayed to be thus far – or at least, that it seems in some important sense correct (be it true, appropriate or both true and appropriate), for DeRose, in a context like HIGH, to utter ‘What I said was false/I spoke falsely’ upon having his earlier knowledge ascription (made in LOW) questioned (see *inter alia* Hawthorne 2004a, Cohen 2004b, Williamson 2005a, Stanley 2005a, MacFarlane 2005, and Davis 2007). In his recent

¹⁸ An alternative: the fact that Tom disagrees with what Emily ‘said’ might ground a second real (but perhaps less important) sense in which there is interpersonal disagreement between Tom and Emily, and in virtue of which a report like ‘Tom and Emily disagree’ can be true. John MacFarlane has recently explored some ideas in (something like) this spirit – see MacFarlane (ms.).

¹⁹ As noted above (fn. 14) it is not clear that this represents a genuine explanatory burden on the sceptical invariantist.

(2009) book, DeRose does not talk about these particular judgments, but he does talk more generally about a typical ‘HIGH’ subject claiming ‘What s/he said is false’ in reference to the knowledge ascription of some typical ‘LOW’ subject. He suggests that in such cases it is far from clear that the oft-reported judgment – that the relevant utterance of ‘What I said is false’ is in some important respect correct – is the judgment that we do in fact have. He personally reports judging that such ascriptions of falsity are themselves *false* and asserts that “you will find a lot of speakers siding” with him on that point (DeRose, 2009, p. 161).

What significance does this apparent dispute carry? Given our present focus on pragmatic defences of sceptical invariantism, it will be instructive to first look at our retraction-like judgments with respect to expressions like ‘all gone’, ‘completely flat’, and so on. Recall *Case I*.

Case I. Peter picks up the vodka bottle and sees that someone has poured the last glass; there are only the remnants left inside. Suppose Mary is about to go to shopping. She asks Peter ‘Is the vodka all gone?’. Peter replies ‘Yes, the vodka’s all gone. We need to buy some more’. Peter’s utterance seems true and appropriate.

Suppose that later that day, Peter’s son approaches Peter and Mary and explains that he needs some traces of alcohol for his science experiment. He asks, ‘Is the vodka all gone?’. Peter replies, ‘No, not quite. There’s probably enough for you.’ Peter’s utterance seems true and appropriate.

Suppose that Mary then questions Peter: ‘But didn’t you say earlier “the vodka is all gone”?’

Suppose Peter responds to Mary’s inquiry by uttering ‘What I earlier said was false’ or ‘I spoke falsely’. In that case, it seems – at least to me and those that I have questioned – that there is a strong temptation to judge that his utterance would be true and appropriate (although also to judge that it would be more clearly appropriate were he to follow up with ‘but obviously I was speaking loosely’, or something along those lines). Notice also that Mary’s question tends to smack of either pedantry or mild idiocy.

Are these judgments good news for the sceptical invariantist? First, notice that our judgments regarding the truth and appropriateness of the relevant response – viz. ‘What I said is false’ – are similar in the extension of *Case I* presented above (featuring ‘all gone’) and the relevant extension of the Bank Case (featuring ‘knows’). Second, notice that the reaction to Mary’s question in this case – pedantry or idiocy – mimics the reaction that Schiffer suggests is natural in the ‘knows’ case. That is, that the challenge to DeRose’s earlier knowledge ascription (in LOW) is one that is likely to annoy a (typical) speaker, because it should have been mutual knowledge that he was speaking ‘casually’. This is surely good news.

But why might DeRose (and his ‘many other speakers’) judge that an utterance of ‘What I said earlier is false’ is *false* in the ‘knows’ case? The natural proposal – on the assumption that sceptical invariantism is correct, and that a pragmatic treatment similar to the proposed pragmatic treatment of ‘all gone’ is on the right lines – is that individuals that judge in this way are focusing on what was *conveyed* by the DeRose’s earlier knowledge ascription in LOW, rather than on what it semantically expressed. On the sceptical invariantist account, that earlier knowledge ascription conveyed the *truth* that DeRose was close enough to knowing that the bank will be open for present purposes (and did not convey what it semantically expressed – viz. that he knew that the bank will be open). This could account for why some individuals might judge that an utterance by DeRose of ‘What I said is false’, in reference to his earlier knowledge ascription in LOW, would seem false. Returning to the ‘all gone’ case, it can indeed seem that there is a reading of ‘What I said is false’ on which it seems false – a reading that is naturally brought out by focusing on the fact that Peter’s earlier utterance of “The vodka is all gone” plausibly conveyed that it was close enough to being all gone for their (shopping-related) purposes.²⁰

How might we accommodate these uses of ‘What is said’ from a more theoretical vantage point? One possible proposal is that there is a kind of context sensitivity in expressions like ‘What I said’. In particular, it might be suggested that they can pick

²⁰ It may also be that DeRose reports these particular judgments, rather than those expressed by the majority, because his focus is on utterances of the form ‘What s/he *claimed*’ rather than ‘What s/he *said*’. See DeRose (2009, p. 160-161)

out *either* what was semantically expressed *or* what was conveyed by the utterance in question (and perhaps they can pick out other things besides).²¹ In the Bank Case then, an utterance by DeRose of ‘What I said is false’ is true on the *what was semantically expressed* interpretation, but false on the *what was conveyed* interpretation. This contrast might account for the dispute over our retraction judgments.²²

There is a more speculative explanation for the dispute over our retraction judgments, one that I shall mention only to set aside. Perhaps the correct account of ‘knows’ for some language users is a sceptical invariantist one, whilst for other language users it’s the contextualist one. On the contextualist account, DeRose’s knowledge ascription in LOW semantically expresses (and conveys) a truth. In consequence, an utterance by DeRose in the context resembling HIGH of ‘What I said is false’ should presumably seem false on such a view. On the sceptical invariantist account, in contrast, DeRose’s knowledge ascription in LOW semantically expresses a falsehood, and so the utterance of ‘What I said is false’ might be expected to seem true. Perhaps the divergence in judgments between DeRose (and his ‘many other speakers’) and the majority of other philosophers reflects a difference in the correct semantic theory for the language of those speakers.

Notice that this proposal presupposes something like an *idiolect*-based account of meaning. Such an account doesn’t preclude the possibility that there might still be a correct semantic theory for a group of language users (English speakers, for example). But the correct semantic theory for the group will be derived from the correct

²¹ For somewhat related discussion on flexibility in our use of ‘What I/he/she said’, see Lewis 1980.

²² This no doubt oversimplifies the situation somewhat. An utterance of ‘What I said is false’, made in reference to an earlier knowledge ascription, will be trivially true on the *what was semantically expressed* interpretation (the sceptical invariantist claims that *no* knowledge ascriptions are true – or that almost none are). Thus, when used on that interpretation, ‘What I said is false’ might often convey something along these lines: what my utterance semantically expressed (e.g. that DeRose knows that the bank will be open) is not close enough to being true for our present purposes. This does not undercut the proposed explanation. It just suggests that, on the *what was semantically expressed* interpretation, DeRose’s utterance of ‘What I said is false’, in the context resembling HIGH, will not only be true, but *also* convey the truth that, for present purposes, *that DeRose knows that the bank will be open* is not close enough to being true.

semantic theory for the individuals that make up that group. If there is sufficient divergence among individuals' idiolects, this might in turn lead to semantic indeterminacy at the group level. In the case at hand, this raises the possibility that *neither* contextualism *nor* sceptical invariantism (+pragmatics) is correct at the level of English speakers, for example. (Given that what is *conveyed* by knowledge ascriptions is for the most part identical across the two accounts, language users whose idiolects differ in the respect at hand will nevertheless communicate successfully.) This is a slightly more radical way to accommodate our disputed retraction judgments. (Of course, it also means giving up on sceptical invariantism – to a certain extent.)

So it seems that the sceptical invariantist can account for our retraction judgments, and for the dispute that has recently arisen regarding them. But are there more pressing objections to the pragmatic defence of sceptical invariantism than retraction, disagreement and cancellability? I consider further objections in sections §5-8. But first I shall address the issue of how this more complicated outlook in re our retraction and disagreement judgments impacts the critique of Chapter 2.

§4. An Aside: The Critique of Chapter 2 Revisited.

Let's consider our disagreement judgments first. Insofar as it is granted that our disagreement judgments run counter to what was proposed in Chapter 2, it seems that the criticism of a pragmatic defence of traditional invariantism presented there will founder. Moreover, it seems that the traditional invariantist might accept the proposals (or simple variations on them) presented in §3.1 to account for why some individuals might have been tempted to judge that there is disagreement in the relevant cases. *Nevertheless*, it is important to recognise that the pattern of judgments consonant with a pragmatic defence of traditional invariantism may well *not* be the one that seems to have been presupposed, nor is it the same pattern that might be expected given other defences of traditional invariantism.

What about our retraction judgments? Unlike the case of disagreement, the situation here remains troubled for the traditional invariantist. First, notice that in the examples featuring those expressions that the traditional invariantist might liken 'knows' to,

there does *not* seem to be a similar plurality of readings of utterances like ‘What I said is false’. (I illustrate here with example of ‘I have eaten’, but that is inessential. The point could just as easily be made using the other (less controversial) examples of pragmatic phenomena featured in Chapter 2 – for example, with the appropriate variation on Grice’s handwriting case.)

*Case I**. Mary and John are discussing whether Peter has ever eaten. Mary is aware that Peter must have eaten at some point in order to still be alive. Mary utters ‘Peter has eaten’.

In that case, Mary’s utterance seems true and appropriate.

Now suppose that a short time later Mary and John are approached by a friend, who asks whether they should get Peter something to eat. Mary is aware that Peter has not eaten all day. In that new context, Mary utters ‘Peter has not eaten’. Once again, Mary’s utterance seems true and appropriate.

Suppose John then questions Mary: ‘Didn’t you just say “Peter has eaten”?’

In Chapter 2, it was reported that were Mary to reply by uttering ‘What I said earlier is/was false’ her utterance would seem false.

In the corresponding formulation of the Bank Case, this reaction corresponds to that suggested by DeRose (2009, p. 160-161) – viz. that from the context resembling HIGH, it seems false for DeRose to say ‘What I said earlier is false’ in reference to his earlier knowledge ascription in LOW. So *that* judgment represents no problem for the (pragmatically minded) traditional invariantist.

However, there does *not* seem to be a remotely natural reading of ‘What I said is false’ on which it expresses something true in the example featuring ‘I have eaten’. For reasons of expediency, Mary might say ‘What I said is false’ in response to John’s question (perhaps if she wanted to avoid returning to a rather pointless discussion), but even then her utterance seems clearly false. (The situation therefore contrasts with the situation we found with expressions like ‘all gone’, for which there

did seem to be the relevant pair of readings.) The absence of such a reading clearly casts doubt on the traditional invariantist's ability to plausibly account for the judgment of the majority of philosophers in the corresponding version of the Bank Case – viz. that DeRose's utterance of 'What I said is false' is true.

Furthermore, notice that the judgment that it is true and appropriate for DeRose to answer his wife's challenge to his earlier knowledge ascription in LOW with 'What I said is false, but I was speaking loosely', seems just as strong as the judgment that it would be true and appropriate for DeRose to respond with 'What I said is false'. (Indeed, the former may well seem considerably *more* appropriate than the latter.²³) But on the traditional invariantist pragmatic account, DeRose is not speaking loosely in LOW. His claim straightforwardly conveys the truth that it semantically expresses. So he is presumably speaking *strictly*. It's therefore especially mysterious why it would seem true and appropriate for him to respond in that way.

These considerations suggest that whilst the preceding discussion might undermine an objection to a pragmatic defence of traditional invariantism centered around our disagreement judgments, the objection centred around our retraction judgments retains its force. Of course, even setting aside the issues surrounding retraction and disagreement, the (to my mind) powerful objection from cancellability to such a view remains, as does the particular challenge for the popular relevance-based pragmatic accounts (see Chapter 2, §5 and §2). There also remain the more general difficulties with traditional invariantism, tied to the various relationships between chance and knowledge, which were discussed in Chapter 3. To conclude this aside: it does not seem that these disputes regarding our retraction and disagreement judgments significantly weaken the case against traditional invariantism presented earlier.

§5. Objections II: Other Judgments Concerning Chance and Knowledge.

It might seem that there are aspects of our talk of epistemic chance and possibility – some that we have not focused on in previous chapters – that are not so readily accounted for by the sceptical invariantist. Suppose the sceptical invariantist accepts –

²³ Cf. The passage from Schiffer quoted earlier; Schiffer (1996, p. 328).

on the basis of the kind of arguments and judgments reviewed in §1 – that for all or nearly all of those propositions that we take ourselves to know, there's a chance, for each of us, that we are wrong. It might then seem troubling that we often appear to appropriately (and perhaps truly) claim that various propositions are epistemically *impossible* and that there's *no* (epistemic) chance, for us, that they are true (and despite the fact that we occupy no more than a fairly strong epistemic position with respect to the relevant propositions).

For example, suppose Mary is considering who to invite to a dinner party this evening at her London home. She asks Peter 'Is there a chance that Emily is still in London?'. If Peter has just got off the phone with Emily, and Emily assured him that she was in New York, then it can seem appropriate for Peter to respond by saying 'No, there's no chance. I just spoke to her, and she's in New York'.

The best response to this kind of phenomena (at least for the sceptical invariantist) is probably to attempt a pragmatic explanation similar to the one employed to account for the HIGH-LOW contrast concerning 'knows'. The core of the account will be the claim that whilst there is indeed an epistemic chance, for Peter, that Emily is in London (since there is, for example, an epistemic chance, for Peter, that Emily was lying about her location), it nonetheless seems appropriate (and perhaps true) for Peter to claim that 'There's no chance that Emily is still in London', since he thereby conveys the truth that there's no *relevant* chance, for him, that Emily is still in London, or perhaps that it's close enough, for present purposes, to there being no chance, for him, that Emily is still in London.

On a Gricean picture, this might be explained by appeal to the maxim of Quality. It would be straightforwardly false to assert that there's no chance, for Peter, that Emily is still in London – there's a chance for nearly all such propositions –, so on the assumption that Peter is trying to be co-operative (and, in particular, relevant), he will convey that there's no *relevant* chance, for him, that Emily is in still in London, or

perhaps that it's close enough, for present purposes, to there being no chance, for him, that Emily is still in London.²⁴

So far, so good. But there are also some judgments that might seem to make trouble for the sceptical invariantist's endorsement of principles like (CK) and (PK). As pointed out at the beginning of Chapter 3, §4, it seems that we (sometimes) appropriately make certain claims that represent *prima facie* counterexamples to principles (CK) and (PK). In particular, utterances like the following can be made felicitously:

- (2) a. There's a chance that all the golfers get a hole-in-one on the thirteenth, but we know that won't happen.

The sceptical invariantist should account for the apparent felicitousness of such (in her view, contradictory) claims if she is to show a clear advantage over SSI and traditional invariantism.

It is noteworthy that utterances of sentences like (2a) tend to sound acceptable only when a distinctive tone of voice is employed (or stress is placed on 'chance' or 'knows').²⁵ If that's right, the advocate of (CK) might suggest that this indicates that the speaker is signalling that they are aware that their claim cannot (strictly speaking) be true, and is indicating to the listener that they mean to convey something other than what their utterance literally semantically expresses – perhaps that it's *so* unlikely that all the golfers get a hole in one, that it's tantamount to their knowing it won't happen.

Furthermore, the idea that utterances of sentences like (2a) do not represent genuine counterexamples to (CK) is supported by the observation that embedded under 'strictly speaking' (and when 'strictly speaking' is read as taking wide scope) (2a) cannot be uttered felicitously.

²⁴ One could of course adopt a (potentially non-Gricean) 'loose use' account of such talk instead. See Davis (2007) for relevant discussion.

²⁵ Cf. Hawthorne (2004a, p. 24, n. 60).

- (2)* a. Strictly speaking, there's a chance that all the golfers get a hole-in-one on the thirteenth, but we know that won't happen.

It seems to be a central function of 'strictly speaking' to encourage us to focus on what is literally semantically expressed by our utterances. If (CK) really failed in this case, it's therefore unclear why an utterance of (2a)* should seem infelicitous (with the 'strictly speaking' read as taking wide scope). That it does might seem to suggest that (2a) expresses a genuine contradiction, but can be felicitously uttered for some other – perhaps pragmatic – reason. The sceptical invariantist should not be overly concerned by such cases.²⁶ I conclude that there does not seem to be a strong objection to sceptical invariantism centred around our judgments concerning epistemic chance and possibility – indeed, as we saw in §1 (and the preceding discussion), the sceptical invariantist seems well-placed to capture those judgments.

§6. Further Objections to Sceptical Invariantism I: Assertion.

An issue that we have not explored up to this point concerns the relationship between knowledge and appropriate assertion. It often seems appropriate to challenge assertions by saying 'Do you know that?' and, in the same vein, to criticise a speaker's assertion by saying 'You don't know that'. It also (typically) seems inappropriate to utter sentences like the following:²⁷

- (3) Dogs bark, but I don't know that they do.

It might be suggested that these judgments are problematic for the sceptical invariantist because the best explanation of such claims is that the following principle is true.

²⁶ Note that there are even *prima facie* counterexamples to the factivity of knowledge, the principle that *S knows p* entails *p*. For instance, utterances of 'I knew that they were going to win, but they didn't' can seem felicitous in some contexts. The proper reaction here is probably not to give up on factivity, but – as in the case of (CK) and (PK) – to suggest some pragmatic explanation of why such utterances seem felicitous. One piece of evidence in favour of this strategy is that utterances of the following seem infelicitous: 'Strictly speaking, I knew they were going to win, but they didn't'.

²⁷ See, esp., Williamson (2000, ch. 11) and Hawthorne (2004a, ch. 1).

(AK) S can appropriately assert that p only if S knows that p.

If the sceptical invariantist accepts such a principle, then it follows that none (or next to none) of our assertions are appropriate. That seems a bad result indeed.²⁸

Fortunately, the sceptical invariantist can explain the relevant phenomena without recourse to principle (AK). From the perspective of the sceptical invariantist, a natural alternative to (AK) is a principle like the following:

(ACK) S can appropriately assert p only if S is close enough to knowing that p for current purposes. (See Davis, 2007, p. 422)

Obviously, some work is required to cash out what ‘current purposes’ might amount to, but the basic idea is fairly clear.²⁹ For instance, in the LOW and HIGH contexts of the Bank Case, the ‘current purpose’ is to make a *decision* whether to wait in line at the bank on Friday or leave going to the bank until Saturday.³⁰ Given that purpose, it seems natural to propose that DeRose is sufficiently close to knowing that the bank will be open for current purposes (and so able to appropriately assert that it will be open) only if he is close enough to knowing to take that proposition for granted in practical reasoning centred on that decision (or something of that ilk). Notice that this account neatly explains why it seems inappropriate for DeRose to utter the simple ‘The bank will be open tomorrow’ in HIGH.³¹

Once the sceptical invariantist has embraced a principle like (ACK), she is well placed to explain the relevant assertion phenomena. First, notice that in virtue of a pragmatic mechanism similar to those discussed in §2, the question ‘Do you know that?’ will probably *not* convey the question it semantically expresses (since it’s trivial that the addressee doesn’t, strictly speaking, know the proposition in question), but rather will convey the question ‘Are you close enough to knowing p for current

²⁸ This kind of concern is pressed in Hawthorne (2004a, p. 134).

²⁹ Note that other prominent proposals are no less vague in the details. For example, in his elaboration on the maxim of Quality, Grice speaks simply of needing ‘adequate evidence’ for one’s conversational contributions, without further elucidation of ‘adequate’.

³⁰ See Davis (2007, p. 409) for some relevant discussion.

³¹ That it seems inappropriate to utter ‘The bank will be open’ in contexts like HIGH is observed in DeRose (2002, p 177-178).

purposes?’ (or a question along those lines). But this is a fully appropriate inquiry given a principle like (ACK) – since if the addressee was not close enough to knowing *p* for her current purposes, then, given (ACK), her utterance was not appropriate.

The sceptical invariantist can also make perfect sense of the fact that one can *criticise* an assertion by saying ‘You don’t know that’. The pragmatic implication (typically) carried by utterances of ‘I don’t know that *p*’ is (something like) that the speaker is not close enough to knowing *p* for present purposes. But given (ACK), that the speaker is close enough to knowing *p* for present purposes is a necessary condition on appropriate assertion.

Finally, the following argument lends support to the conclusion that (3) should seem inappropriate given a principle like (ACK).³² On the assumption that sceptical invariantism is true, the speaker’s utterance of the latter conjunct – ‘I don’t know that they do’ – will presumably convey *either* that the speaker does not (strictly speaking) know that dogs bark, *or* that she is not close enough to knowing that they do for present purposes. (It is difficult to see how else her utterance of that conjunct could be interpreted.)

First, suppose the speaker is trying to convey that she does not (strictly speaking) know that dogs bark. In that case, she presumably occupies (or is at least trying to occupy) a context in which the ‘current purposes’ are such that to be close enough to knowing is to actually *know*. Otherwise her utterance of ‘I don’t know that they do’ would seem to be irrelevant and/or trivial. But if she *does* occupy such a context, then, given (ACK), she must (strictly speaking) know that dogs bark in order to appropriately assert the first conjunct. So on the supposition that the speaker is trying, by uttering the second conjunct, to convey that she does not (strictly speaking) know that dogs bark, her utterance of the first conjunct is inappropriate.

Let’s take the alternative, and suppose instead that the speaker is trying to convey that she is not close enough to knowing that dogs bark for present purposes. In that case,

³² Note that this account differs from the one proposed in Davis (2007, p. 422).

she is trying to convey something that stands in contradiction to what, given (ACK), is a condition on appropriate assertion of the first conjunct – viz. that she *is* close enough to knowing that dogs bark for present purposes. On either supposition then, it is unsurprising that an utterance of (3) should seem inappropriate. I conclude that there does not seem to be a strong objection to sceptical invariantism centred around principle (AK).

§7. Further Objections II: The Moorean Fact and the Mysterious Semantic Value.

The sceptical invariantist claims that we know very little (if anything). But isn't it a 'Moorean fact' that we know a lot? If so, isn't sceptical invariantism just plain false?³³

Certain responses to this challenge seem, at least to me, unpromising. Perhaps the most popular (contemporary) response is to moderate the kind of sceptical invariantism on offer. For instance, Davis (2007) advocates a form of sceptical invariantism according to which to (strictly speaking) know something requires that one be completely justified in believing it. He suggests that people are indeed completely justified in believing many things, and that, oftentimes, they also know those things. For instance, supposing DeRose were right now in the bank, waiting in line, listening to the bustle, and so on, Davis suggests that DeRose would be completely justified in believing that he is in the bank, and that he would (strictly speaking) know that he is in the bank. Davis goes on to suggest that we know many such propositions about our immediate environment, and much more besides. His answer to this 'Moorean challenge' is therefore that – although sceptical invariantism is true – we *do* know a lot.³⁴

The central objection (at least to my mind) to this kind of suggestion is that – for reasons that the reader can no doubt anticipate – it suffers from many of the problems

³³ Cf. Lewis (1996, p. 549), who writes: “We have all sorts of everyday knowledge, and we have it in abundance. To doubt that would be absurd. At any rate, to doubt it in any serious and lasting way would be absurd; and even philosophical and temporary doubt, under the influence of argument, is more than a little peculiar. It is a Moorean fact that we know a lot. It is one of those things that we know better than we know the premises of any philosophical argument to the contrary.”

³⁴ See Davis (2007, p. 426-430).

that afflicted traditional invariantism. For instance, this moderation seems to lead to (at least) a failure to address judgments concerning ‘lottery style’ cases, a failure to respect the apparent connections between objective chance, epistemic chance, and knowledge, and a failure to respect (MPC).³⁵

Let’s focus on the first-mentioned problem. It is supposed to be an advantage of positions like SSI and contextualism that they can accommodate judgments like the judgment that it oftentimes seems true and appropriate to utter ‘I don’t know that I’ll lose the lottery’ (prior to the draw of a fair lottery, in which you hold a ticket).³⁶ As we saw in Chapter 1, given closure, judgments like these threaten the conclusion that we know very little.

Of course, the sceptical invariantist is – one might have thought – in the pleasing position that she can just accept these judgments, and their sceptical conclusions, at face value. But it is difficult to see how to do this once we moderate the scepticism on offer. To see this, it will be instructive to look at Davis’s (2007) attempt to grapple with these issues. Davis tries to walk a middle path. He suggests that even though DeRose would be completely justified in believing that he is in the bank (and in believing many other things besides), he would *not* be completely justified in believing that he will, for instance, lose the lottery.³⁷ Consequently, Davis suggests that even though DeRose (strictly speaking) knows that he is in the bank (and many other things besides), he does not (strictly speaking) know that he will lose the lottery.

The problem with this kind of suggestion is familiar. The odds that DeRose loses the lottery can be vanishingly small, and presumably equal to those that, for example, he is not in fact in the bank (and, unbeknownst to him, the bank has been converted into a bank-like film set, or he’s hallucinating being in the bank, or whatever). So why is he completely justified in believing that he’s in the bank but *not* completely justified

³⁵ It may well also lead to a failure of principles connecting knowledge and appropriate practical reasoning, and knowledge and assertion, that Davis himself seeks to endorse. For instance, Davis writes (2007, p. 424) : “If S really knows p, then S would seem to have enough evidence to use “p” as a premise in any practical reasoning, not just the reasoning S happens to be engaged in”. He also endorses (2007, p. 405) the following principle: if S knows p, then S has enough justification or evidence for p to assert p in *any* context.

³⁶ See e.g. DeRose (1996), Williamson (2000, ch. 11), and Hawthorne (2004a, ch 1).

³⁷ Davis (2007, esp. p. 414 & p. 434).

in believing that he'll lose the lottery? Davis offers little guidance, and it is very difficult to see how the sceptical invariantist might answer such questions.³⁸

Of course, Davis could just retreat, and maintain that DeRose does (strictly speaking) know that he will lose the lottery as well. But problems would still abound. For example, would it then be true for DeRose to utter '(Strictly speaking,) I know that I will lose the lottery, but I might win'? Such utterances seem straightforwardly infelicitous. Let's therefore suppose that they are not true, and express contradictions. But that is no help, because then it should presumably seem true – in all contexts – for DeRose to utter 'It's not the case that I might win the lottery, but I have a ticket, and the lottery is fair'. This consequence seems just as bad. I shan't dwell on the 'moderation' option further here.

Can the sceptical invariantist hold firm and admit that we do *not* know a lot? Perhaps she can. The first thing to note is that it is scarcely difficult to get a good many speakers to admit that 'You never really know anything'. This is important in at least two respects. First, it reveals that the status of the relevant 'Moorean fact' is not on a par with other prominent 'Moorean facts' in philosophy, such as the 'Moorean fact' – familiar from discussions in metaphysics – that things resemble each other. Try getting speakers to admit the denial of *that*! Second, to the extent that the denial of the relevant 'Moorean fact' constitutes the natural reaction on behalf of many speakers to some gentle probing, it seems that a theory that denies that 'fact' might not be much impugned by so doing.

Of course, if you simply confront someone with the question 'Do we know a lot?', the answer you will typically receive is 'Yes'. But *this* kind of reaction can be accounted for by appeal to just the kind of pragmatic mechanism we have already explored. Such a general question could be considered relevantly similar to something like the question – in the case of 'one hour' – 'Do many activities last one hour?', or – in the case of 'all gone' – to the question, 'Is the phosphorus of some nations all gone?'

³⁸ See Hawthorne (2004a) for relevant discussion.

The natural answer to these questions is also ‘Yes’. However, there is bound to be *some* phosphorus left in the relevant countries, and comparatively few activities last *precisely* one hour (and not a nanosecond more or less). In these cases, it seems that the relevant question is naturally – and unconsciously/implicitly – interpreted as (respectively) ‘Is the phosphorus of some nations sufficiently close to being all gone such that no further mining activity will take place?’ and ‘Do many activities last *about* one hour?’ (or something along these lines). The sceptical invariantist might therefore propose that the question ‘Do we know a lot?’ is naturally – and unconsciously/implicitly – interpreted as asking something like: ‘Are we sufficiently close to knowing a great many things for most purposes?’ or ‘Do we have excellent evidence to believe many things?’ The sceptical invariantist can still accept that the answer to *these* questions is a resounding ‘Yes!’, and propose that that is why people respond as they (typically) do.

However, there is a more subtle worry here. If we rarely use knowledge ascriptions and denials to convey what the sceptical invariantist claims is semantically expressed by them, how do they come to semantically express what we are told that they do? This is the ‘Mysterious Semantic Value’ objection (see Hawthorne 2004a, p. 122-123).

In response, the sceptical invariantist might sensibly draw attention to various features of our use that lend support to the assignment of the sceptic’s semantic value to ‘knows’.

(a) First, the central tenet of the theory – ‘You never really know anything’ – is something that many speakers endorse rather readily.

(b) In a related vein, it has been suggested that the most popular reaction among speakers to the puzzle posed by ‘moderate sceptical arguments’ is the one predicted by the sceptical invariantist. The arguments in question are those like the following.

(1) You don’t know that your car wasn’t towed or stolen

(2) If you don’t know that your car wasn’t towed or stolen, then you don’t know that it’s in the Main Street parking lot.

(C) You don't know that your car is in the Main Street parking lot.

DeRose (2009, p. 179, n. 21) suggests that the most popular reaction to the *prima facie* tension between (C) and the knowledge ascription that might be more typical in everyday contexts – viz. 'You know that your car is in the Main Street parking lot' – is that the above argument reveals that (C) is true (in all contexts), and that the ordinary claim to know is false, but perhaps appropriate. Of course, this is just what the sceptical invariantist claims (and the traditional invariantist, subject sensitive invariantist (and the standard contextualist) denies).

(DeRose does then suggest that the situation might be different with respect to sceptical arguments that play on 'Brain-in-a-Vat'-style scenarios, and that in *that* instance, a contextualist resolution might be more commonly accepted. However, it seems plausible that our reactions to less outlandish scenarios are more important in determining meaning.)

(c) Sceptical invariantism allows us to respect those principles that seem to underlie our ordinary use of 'knows' – principles like (CK), (PK), and (MPC), and those expressing intellectualism, and certain connections between knowledge and action, like (1a). The only theory – at least on the assumption of invariantism – that enables us to respect these principles is the sceptical one. The sceptical invariantist might therefore maintain that our commitment, revealed in ordinary usage, to principles like these plays a crucial role in giving 'knows' the meaning it has.

(d) Finally, on the assumption that knowledge ascriptions and denials express what the sceptical invariantist suggests that they do, we are able to explain our judgments concerning ordinary knowledge ascriptions and denials by appeal to independently motivated pragmatic features of language.

In the absence of a more precise objection, the above considerations seem sufficient to explain how a theory that accounts for our ordinary use of 'knows', whilst claiming that knowledge ascriptions and denials oftentimes do not convey what they semantically express, can be true.

§8.1 Further Objections III: Contextualism.

Recall that the sceptical invariantist explains our judgments to the effect that our ordinary knowledge ascriptions are true and appropriate – for example, our judgment to the effect that DeRose’s knowledge ascription in LOW is true and appropriate – by appeal to the claim that although false, those ascriptions convey truths. However, it might seem to be an advantage of a theory if it can explain our judgments to the effect that some utterance is true and appropriate by taking that judgment *at face value*. That is, it might seem to be an advantage of an account if it not merely explains why we might judge that a particular utterance is true and appropriate, but explains that judgment in terms of the fact that the utterance in question semantically expresses a truth.

In this vein, Brian Weatherson (2003, p. 8) writes:³⁹

If we can explain why we have the mistaken intuition [in the case at hand, and from the sceptical invariantist’s viewpoint, the ‘mistaken intuition’ would be the judgment that DeRose’s utterance of ‘I know that the bank will be open’ in LOW is true and appropriate], that counts for a lot in reducing the damage the counterexample does. Grice did not just assert that the theory on which an ordinary head scratch was voluntary was more systematic than the theory of voluntariness Ryle proposed, he provided an explanation of why it might seem that his theory was wrong in certain cases.

Such remarks might be taken to imply that it would nevertheless be *better* for a theory of the meaning of some expression if it did not encounter ‘mistaken intuitions’ in the first place. In the case at hand, this suggests that it would be better for a theory of ‘knows’ if it could explain our judgments to the effect that various utterances are true and appropriate by appeal to the fact that those utterances semantically express truths. This thought might be captured by a principle like (MSP):

(MSP) In regard to an arbitrary expression, ‘A’, one should – *ceteris paribus* – prefer the theory of the meaning of ‘A’ that explains our judgments in re truth and

³⁹ See also DeRose (2009, p. 50-51).

appropriacy by appeal to the fact that the utterances in question semantically express truths.

This leads us into a final issue: contextualism.⁴⁰

§8.2 Further Objections III: Contextualism Contd.

There is a natural objection to a pragmatic defence of sceptical invariantism that we have not considered thus far. This is the objection that the ‘HIGH-LOW’ contrast in our judgments concerning expressions like ‘all gone’, ‘completely flat’, and ‘two hours’ should be given a *semantic* rather than a pragmatic explanation. In consequence – the thought runs – the fact that our judgements regarding ‘knows’ exhibit a clear resemblance to our judgments regarding those expressions only bolsters the conclusion that *contextualism* about knowledge ascriptions is true.⁴¹

In barest outline, the standard contextualist explanation of the relevant ‘HIGH-LOW’ contrast in the case of ‘all gone’ will take the following form. In *Case I*, Peter’s utterance of ‘The vodka is all gone’ semantically expresses something like *that the vodka is all gone by low standards of precision*. In *Case II*, Peter’s utterance of ‘The vodka’s not all gone’ semantically expresses something like *that the vodka is not all gone by high standards of precision*. In both cases, his utterances semantically express something true, and that accounts for our judgments to the effect that both utterances seem true and appropriate.

The contextualist about ‘all gone’ might suggest that her explanation has the advantage that not only can we *respect* (that is, explain) our judgments to the effect that the relevant utterances are true and appropriate, we can *also* account for those judgments in terms of the fact that the utterances in question semantically express truths. On these grounds – in effect, an appeal to (MSP) – the contextualist might

⁴⁰ Of course, there may be further possible objections to the sceptical invariantist position. I have sought to address some of the most pressing here. One prominent objection (targeted at an ‘exaggeration’ account of the HIGH-LOW contrast) that has not been considered here is centred around our use of ‘literally’ (see Hawthorne 2004a, p. 120). See Davis (2007, p. 410) for a response to this problem.

⁴¹ Davis (2007, p. 415) credits Stewart Cohen with suggesting an objection along these lines.

suggest that her explanation is to be preferred. In consequence, the fact that our judgments regarding expressions like ‘all gone’ exhibit widespread resemblance to those regarding ‘knows’ would seem merely to support the case for contextualism about ‘knows’, and not the case for sceptical invariantism. Or so the thought runs.

There is a more general concern here, one not (at least not directly) tied to the correct semantic treatment of expressions like ‘all gone’. The concern is that adopting a contextualist account of our judgments regarding ‘knows’ enables us not simply to win the advantages of sceptical invariantism, but *also* to explain our judgments to the effect that our everyday knowledge ascriptions are true and appropriate by appeal to the fact that they semantically express truths. If that’s right, then it can seem that even granting that sceptical invariantism represents the most plausible form of invariantism – and even granting that sceptical invariantism is able to respect all the relevant judgments concerning our use of ‘knows’ – it can nevertheless seem that contextualism about ‘knows’ is correct. (This argument also involves implicit appeal to some principle like (MSP).)

The central object of the present work has been to reveal the central pitfalls of recent defences of traditional invariantism and subject sensitive invariantism, and to suggest that the sceptical invariantist is in fact well-placed to explain much of the relevant phenomena. The goal here is not to undertake a thorough comparison of the relative merits and demerits of contextualism and sceptical invariantism. Nevertheless, this work was not undertaken without at least the suspicion that sceptical invariantism is not merely ‘The Best Invariantism’, but also the correct theory. It is therefore important to engage in some preliminary discussion of the grounds on which might one argue for the (sceptical invariantist’s) pragmatic treatment *over* the (contextualist’s) semantic treatment. The following reflections may also give pause to those who take it that the argument of the preceding chapters merely serves to bolster – given the no doubt popular thought that sceptical invariantism is just an untenable capitulation – the case for contextualism.

Let’s first consider some suggestions from the literature. Davis (2007) offers two proposals for how to support the sceptical invariantist position over that of the contextualist. The first is to argue that our judgments concerning ‘knows’ *also* mimic

our judgments in cases for which the relevant form of contextualism is clearly implausible. He offers ‘trumpet’ as a case in point.⁴² Suppose that Peter plays the cornet but not the trumpet. (The cornet is a kind of horn similar to but distinct from the trumpet.) Davis suggests that in certain ‘LOW’ contexts – for example, those in which Peter is talking to a lay person with little knowledge of the relevant instruments – Peter might appropriately utter ‘I play the trumpet’. In other ‘HIGH’ contexts – for example, those in which Peter is in the midst of an interview for a place at a prestigious music school – it might seem *inappropriate* for Peter to utter ‘I play the trumpet’ and true and appropriate for him to utter ‘I don’t play the trumpet (but I do play the cornet)’. In short, we seem to find a similar ‘HIGH-LOW’ contrast with ‘trumpet’.

If the range of judgments we witness concerning what (might at least appear to be⁴³) entirely uncontroversial cases exhibits widespread resemblance to the range of judgments we see with ‘knows’ (and expressions like ‘all gone’), that might seem to lend support to the idea that the sceptical invariantist’s pragmatic treatment of ‘knows’ (and ‘all gone’) is to be preferred to the contextualist’s semantic treatment (at least, on the assumption that there are no clear-cut cases of *contextualism* for which we see similar, widespread resemblance in our judgments).

Such an appeal to cases like ‘trumpet’ strikes me as unlikely to succeed. The central problem is that there seem to be potentially pertinent disanalogies in our judgments concerning these apparently more ‘clear-cut’ cases of pragmatic phenomena, and the case of ‘knows’. For example, it seems that there is a *much* stronger temptation to judge that Peter’s utterance of ‘I play the trumpet’ in the relevant ‘LOW’ context is false (but perhaps appropriate), than there is to judge (at least initially) that DeRose’s knowledge ascription in LOW is false (but perhaps appropriate). There are other disanalogies besides. For instance, it seems plausible that it is a universal (and firmly

⁴² See Davis (2007, p. 417)

⁴³ It seems that one *could* pursue a semantic treatment of the ‘HIGH-LOW’ contrast in the case of ‘trumpet’ as well. On the standard contextualist line, Peter’s utterance in the ‘LOW’ context of ‘I play the trumpet’ would semantically express *Peter plays the trumpet by low standards of precision* and, in the ‘HIGH’ context, *Peter doesn’t play the trumpet by high standards of precision*. (Note that Lewis (1980) can be read as endorsing a kind of contextualist treatment of such ‘loose use’. In MacFarlane’s (2009) terminology, Lewis’s account might be described as a form of ‘non-indexical contextualism’.)

held) belief among musical experts that a cornet is distinct from a trumpet. But it is not a universal (and firmly held) belief among epistemologists (or any other relevant group, for that matter) that contextualism rather than sceptical invariantism is true. The contextualist might leverage such pertinent disanalogies in our judgments to suggest that although a pragmatic explanation of the phenomena associated with ‘trumpet’ is plausible, that does little to support the claim that the sceptical invariantist’s pragmatic account of ‘knows’ is to be preferred to the contextualist’s semantic account.

Davis has another suggestion. He suggests that our retraction judgments in the case of ‘knows’, ‘all gone’, ‘two hours’, and so on, favour a pragmatic treatment, rather than a semantic treatment, of those expressions.⁴⁴ Why might that be?

On the contextualist account, DeRose’s utterance of ‘I know that the bank will be open’ in LOW semantically expresses, and conveys, something like the proposition that DeRose knows that the bank will be open by low standards. That proposition is true. It therefore seems that the contextualist, like the traditional invariantist, has no difficulty accounting for the fact that to DeRose (and his ‘many other speakers’) an utterance of ‘What I said is false’, in the context resembling HIGH and in reference to his earlier knowledge ascription, might seem *false*. However, the reaction of the majority of philosophers who have written on the topic is much more difficult to accommodate. On the contextualist assumption that DeRose’s utterance of ‘I know that the bank will be open’ in LOW semantically expresses *and also* conveys something true, why would it seem true (and perhaps appropriate) for DeRose to utter ‘What I said is false’ in reference to that ascription? It seems that the contextualist might lack an adequate explanation of this fact, and is therefore forced to attribute

⁴⁴ See Davis (2007, p. 415). Davis (2007, p. 415-417) also suggests that our judgments concerning *contradiction* – for example, whether what DeRose said in LOW contradicts what he says in HIGH – favour the sceptical invariantist over the contextualist position. Our retraction judgments have been much more widely reported in the literature, and the dialectical situation is rather similar in the case of our retraction judgments (even down to the recent dispute regarding them (see DeRose 2009, p. 157-159)), so I shall keep the discussion focused on those. Nevertheless, in providing a full defence of her view over that of the contextualist, the sceptical invariantist may well wish to draw on our contradiction judgments as well.

significant error to those – perhaps the majority of speakers – who are inclined to judge in this way.⁴⁵

Given, as we saw in §3.2, that it seems that the sceptical invariantist can explain these judgments *without* attributing such error, they might seem to favour the sceptical invariantist's pragmatic account over the contextualist's semantic account of 'knows'. Parallel remarks apply to the pragmatic and the semantic account of 'all gone'.

In response, the contextualist might point to the fact that we see similar retraction judgments to those we find in the case of 'knows' (and 'all gone') in the case of expressions for which contextualism is the most plausible theory – in particular, in the case of certain gradable adjectives like 'flat' and 'square'.⁴⁶

Let's first look at the relevant 'HIGH-LOW' contrast in the case of 'flat', and the pertinent retraction judgments. Take the 'LOW' case first. Suppose that Peter and Mary are in a laboratory. Peter has just made some tea, and asks Mary where he should put the tea pot. Mary points to the desk 'That's flat. Can you put it there?'. Her utterance seems appropriate and true. Now a 'HIGH' case. Suppose that at some later time Mary is practicing an experiment that she will later undertake with her students, and that involves rolling ball bearings on various surfaces. Peter sees Mary rolling the ball bearings across the desk and remarks: 'Haven't you considered all the small bumps on the surface?'. Mary replies 'Yes, it's not flat. But I don't have anything flatter.' In that case, Mary's utterance also seems true and appropriate.

Suppose Peter then questions Mary: 'But didn't you say earlier 'That's flat'?'

In that case, it can seem true for Mary to reply 'What I said is false', and perhaps appropriate for her to do so, but much more appropriate for her to do so were she to

⁴⁵ Of course, the contextualist might suggest that our judgments can be explained by appeal to the fact that speakers are 'semantically blind'. But I take it that such appeals amount to little more than the (rather lame) claim that speakers have a faulty grasp of the meaning of 'knows'. As noted in the thesis Introduction, the attribution of such error seems to represent a serious objection to a theory (hence, 'significant error'), in the way in which the attribution of a 'mistake' associated with the effects of some general psychological bias, or of some pragmatic feature of language, does not.

⁴⁶ As observed in Cohen (1999, 2004b).

follow up with ‘but obviously I was speaking loosely’ (or something along those lines).

It therefore seems that in the case of ‘flat’, we see similar retraction judgments to those that we found in the parallel extension of the Bank Case.⁴⁷ The concern thus arises that, given that we see similar retraction judgments in the case of expressions for which a contextualist treatment is the most plausible, our retraction judgments cannot be used to support sceptical invariantism over contextualism.

Can the sceptical invariantist rebut this concern? The central thing to note is that this observation in regard to ‘flat’ and ‘square’ does nothing – in and of itself – to undermine the claim that the contextualist has to attribute significant error to speakers with respect to their retraction judgments (be it in the case of ‘flat’ *or* ‘knows’). In consequence, the sceptical invariantist might still reasonably claim that our retraction judgments in the case of ‘knows’ *do* favour the sceptical invariantist treatment over the contextualist one.

Notice that this need not undermine the claim that – ultimately – the contextualist treatment in the case of ‘flat’ (and ‘square’) is to be preferred to some invariantist pragmatic treatment. In particular, it is plausible there might be *other* judgments that favour the contextualist’s semantic treatment in *those* cases. The natural alternative to the contextualist’s semantic treatment of the ‘HIGH-LOW’ contrast in the case of ‘flat’ and ‘square’ is that utterances of the form ‘X is flat’ and ‘X is square’ semantically express, respectively, that X is *geometrically* flat and *geometrically* square, but that such utterances typically pragmatically convey that X is close enough to being geometrically flat (or square) for present conversational purposes.

Notice that such a pragmatic treatment has at least *prima facie* difficulty with the fact that utterances of the form ‘It’s flat but not perfectly flat’ often seem felicitous.⁴⁸ The

⁴⁷ Notice that there is also a temptation to say that Mary’s utterance of ‘What I said is false’ is *false*, especially once we focus on the details of their former situation and the fact that she clearly intended to convey something like *the table is flat by everyday (teapot placing) standards*. The contextualist about ‘knows’ might try to appeal to this fact when addressing the relevant ‘dispute’ concerning our retraction judgments (discussed in §3.2, above).

⁴⁸ Cf. Davis (2007, p. 418).

reason is that, on the pragmatic account, such utterances presumably semantically express something along the lines of *It's geometrically flat, but it's not geometrically flat* – that is, contradictions. It might therefore seem that, in this case, the defender of the pragmatic account has to attribute significant error to speakers. In contrast, it seems that the contextualist about 'flat' can readily account for the felicitousness of such utterances, since she can maintain that they do not (at least not typically) semantically express contradictions, but rather express something entirely consistent, such as *It's flat by everyday standards but not flat by geometric standards*.

If there are sufficient judgments like these – and it transpires that the pragmatic account cannot be plausibly extended to explain them – such judgments might serve to balance out the error that the contextualist has to attribute in regard to our retraction judgments concerning 'flat'. Moreover, it might seem that there are *no* such outweighing considerations in the case of 'knows' (or 'all gone'). For example, it does not seem felicitous to utter such things as 'I know it, but there's a small chance I'm wrong' (or 'It's all gone, but there's a tiny bit left') in the manner in which it seems felicitous to utter 'It's flat, but not perfectly flat'. This kind of observation might ground the truth of the claim that whilst contextualism about 'flat' and 'square' is correct, contextualism about 'knows' (and 'all gone', 'two hours', 'completely flat', and so on) is not.

In any case, the important point is that simply observing that there are other plausibly context sensitive expressions for which we see similar retraction judgments does nothing – at least not by itself – to undermine the claim that the contextualist must attribute significant error in regard to those judgments, and that they therefore favour the relevant pragmatic explanation over the contextualist one.

In the absence of other judgments favouring the contextualist position over that of the sceptical invariantist, our retraction judgments appear to compromise the contextualist's ability to invoke a metasemantic principle like (MSP) in defence of her position. But is that important? After all, the contextualist might respond by noting that she is able to explain our judgments to the effect that our ordinary knowledge ascriptions are true and appropriate in terms of the fact that those ascriptions semantically express truths *to a much greater extent* than is the sceptical invariantist.

She might suggest that this outweighs the negative impact of the fact that she (the contextualist) must attribute significant error in the case of our retraction judgments (supposing she does indeed have to). However, that is to embrace a far more controversial metasemantic claim.⁴⁹ But it does raise an important question: Can the sceptical invariantist find *further* judgments that might support her theory over the contextualist theory?

Here are two (potentially overlapping and non-exhaustive) options.

(i) Argue that the contextualist is unable to account for various of the judgments that have been considered in the course of the critique of traditional and subject sensitive invariantism in Chapters 1-4. For example, it might be suggested that the contextualist cannot resolve the dilemma facing the traditional invariantist centred around our judgments concerning (MPC) presented in Chapter 3, §5.

(ii) Highlight disanalogies between alleged context sensitive expressions and expressions like ‘knows’, ‘all gone’, ‘completely flat’, and ‘two hours’. (Something akin to this strategy (though not with the object of defending sceptical invariantism) is pursued at length in Stanley (2004, 2005a).⁵⁰)

⁴⁹ Of course, the sceptical invariantist may wish to challenge even (the presumably less controversial) (MSP). Notice that in the passage quoted at the close of the previous section in support of (MSP), Weatherson speaks of a ‘mistaken intuition’ and the possibility that a Gricean (or pragmatic) explanation might dispel a large part – but seemingly not all – of the negative impact of having to declare an intuition ‘mistaken’. To the extent that the persuasiveness of the thought that a Gricean/pragmatic explanation would not remove *all* of the relevant negative impact that lies in the attribution of ‘mistake’, the sceptical invariantist might try and push the line that there really is *no mistake at all* in our judgment to the effect that DeRose’s knowledge ascription in LOW is true and appropriate (and in other judgments of a similar kind). For instance, the sceptical invariantist might suggest that our judgment to the effect that DeRose’s knowledge ascription is true and appropriate is nothing more than a judgment concerning the truth or falsity of *what he said*, and there are a number of thoroughly legitimate readings of ‘What he said’ (cf. the discussion in §3.2 above). In particular, it might be suggested that there is a thoroughly legitimate (and from the sceptical invariantist’s viewpoint, it is arguably the most natural) reading of ‘What he said’ on which it picks out what DeRose’s utterance conveyed. But what it conveyed is true. Viewed from this perspective, it might seem that – on the sceptical invariantist’s pragmatic account – there is no ‘mistake’ at all in our judgment to the effect that DeRose’s knowledge ascription in LOW is true and appropriate.

⁵⁰ DeRose (2009) comprises in part (at least the beginnings of) a contextualist attempt to respond to various of those objections. See also DeRose (2008).

Consider (ii). One observation recorded earlier might be pertinent in pursuing this particular strategy. This is the fact that it has been suggested (DeRose 2009, p. 179, n. 21) that the reaction among ordinary speakers to ‘moderate sceptical arguments’ appears to favour sceptical invariantism over contextualism. In a related vein, it is observed in Cohen (2004b, p. 192) that speakers are much more willing to accept the contextualist resolution of ‘sceptical arguments’ focused on ‘flat’ than they are to accept the contextualist resolution of sceptical arguments focused on ‘knows’. These reactions/judgments *also* seem to favour the sceptical invariantist account of ‘knows’ over the contextualist one. In consequence, they might sensibly comprise an additional part of the sceptical invariantist’s case that her theory is superior to that of the (standard) contextualist.

I should stress again that such reflections on how the sceptical invariantist might argue for the superiority of her position to that of the contextualist are rather speculative and programmatic in nature. I should also stress that I by no means wish to endorse the sceptical invariantist position here. The primary object of this final – and rather more speculative – chapter has been to suggest that the sceptical invariantist position represents a *much* more serious contender in these debates than it is often given credit for, and is, at least arguably, ‘The Best Invariantism’. Let me therefore close with a brief summary of what – in that regard – are the pertinent features of the view.

The sceptical invariantist seems well-placed to capture those judgments and arguments – explored in Chapter 3, and partially reviewed in §1, above – that lend strong support to the conclusion that for much of what we ordinarily claim to know, there’s an epistemic chance, for each of us, that we are wrong. In Chapters 3 and 4, we saw that the traditional and subject sensitive invariantist were able to respect such judgments only by forfeiting their ability to explain the infelicitousness of utterances like ‘I know that p, but there’s a chance that not-p’ and ‘I know that p, but it might be that not-p’. However, it seems that the sceptical invariantist is able to explain the infelicity of such conjunctions by endorsing the attractive principles (CK) and (PK), *and* respect the judgments and arguments that supported the conclusion that, for much of what we take ourselves to know, there’s a chance, for each of us, that we are wrong.

(CK) If S knows that p, then there's no epistemic chance, for S, that not-p.

(PK) If S knows that p, then it's epistemically impossible, for S, that not-p.

In a similar vein, the sceptical invariantist is able to endorse (MPC), whilst avoiding the dilemma that we saw in Chapters 3 and 4 to beset traditional and subject-sensitive invariantist endorsement of that principle.

(MPC) If a subject S knows propositions p_1 to p_n and competently deduces q from p_1 to p_n – retaining throughout her knowledge of p_1 to p_n – then S knows q.

The sceptical invariantist is also able to respect the infelicity of pertinent utterances relating knowledge and inquiry, and knowledge and appropriate practical reasoning, such as those of the form 'I know that p, but we need to inquire further' and 'I know that p, but I cannot rely on p in my practical reasoning', by adopting principle (1a).

(1) a. If S knows that p, then S can appropriately rely on p in her practical reasoning.

It appeared that the only way for the non-sceptical invariantist (that is, the traditional or subject-sensitive invariantist) to capture the infelicity of such utterances was to give up on the (attractive) thesis of intellectualism. In Chapter 4, we witnessed the severity of the problems that accompany invariantist attempts to reject that (attractive) thesis.

Finally, in Chapters 1, 2, and 4, it was shown how the traditional and subject-sensitive invariantist has considerable difficulty capturing our ordinary use of 'knows', and, in particular, the HIGH-LOW contrast revealed in examples like the Bank Case. In contrast, it seems that by means of an appeal to pragmatic features of language, the sceptical invariantist is able to respect the HIGH-LOW contrast, and, in turn, to make sense of our ordinary use of 'knows'.

References

- Bach, K. (1994). Conversational Implicature. *Mind & Language*, 9, 124-162.
- _____ (1999). The Myth of Conventional Implicature. *Linguistics and Philosophy*, 22, 327-366.
- Birch, S. and Bloom, P. (2004) Understanding Children's and Adults' Limitations in Mental State Reasoning. *Trends in Cognitive Sciences*, 8, 255-260.
- Blank, H., Musch, J. and Pohl, R. (2007). Hindsight Bias: On Being Wise after the Event. *Social Cognition*, 25, 1-9.
- _____ and Nestler, S. (2007). Cognitive Process Models of Hindsight Bias. *Social Cognition*, 25, 132-146.
- Blome-Tillmann, M. (2008). Conversational Implicature and the Cancellability Test. *Analysis*, 68, 156-160.
- Brown, J. (2005). Adapt or Die: The Death of Invariantism? *The Philosophical Quarterly*, 55, 263-285.
- _____ (2006). Contextualism and Warranted Assertibility Manoeuvres. *Philosophical Studies*, 130, 407-435.
- _____ (2008). Subject Sensitive Invariantism and the Knowledge Norm for Practical Reasoning. *Noûs*, 42, 167-189.
- Campbell, J. and Tesser, A. (1983). Motivational Interpretations of Hindsight bias: An Individual Difference Analysis. *Journal of Personality*, 51, 605-620.
- Cohen, S. (1988). How to Be a Fallibilist. *Philosophical Perspectives*, 2, 91-123.
- _____ (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13, 57-89.
- _____ (2004a). Knowledge, Assertion, and Practical Reasoning. *Philosophical Issues*, 14, 482-491.
- _____ (2004b). Contextualism and Unhappy-Face Solutions: Reply to Schiffer, *Philosophical Studies*, 119, 185-197.
- Dawes, R. (1989). Statistical Criteria for Establishing a Truly False Consensus Effect. *Journal of Experimental Social Psychology*, 25, 1-17.
- Davis, W. (1998). *Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory*. Cambridge: Cambridge University Press.

- _____ (2007). Knowledge Claims and Context: Loose Use. *Philosophical Studies*, 132, 395-438.
- DeRose, K. (1991). Epistemic Possibilities. *The Philosophical Review*, 100, 581-605.
- _____ (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52, 913-929.
- _____ (1995). Solving the Skeptical Problem. *The Philosophical Review*, 104, 1-52.
- _____ (1996). Knowledge, Assertion and Lotteries. *Australasian Journal of Philosophy*, 74, 568-580.
- _____ (1999). Contextualism: An Explanation and Defense. In J. Greco & E. Sosa (Eds.), *The Blackwell Guide to Epistemology* (p. 187-205). Oxford: Blackwell Publishing.
- _____ (2002). Assertion, Knowledge, and Context. *The Philosophical Review*, 111, 167-203.
- _____ (2004). The Problem with Subject Sensitive Invariantism. *Philosophy and Phenomenological Research*, 68, 346-350
- _____ (2006). “Bamboozled by Our Own Words”: Semantic Blindness and Some Arguments Against Contextualism. *Philosophy and Phenomenological Research*, 73, 316-338.
- _____ (2008). Gradable Adjectives: A defence of pluralism. *Australasian Journal of Philosophy*, 86, 141-160.
- _____ (2009). *The Case for Contextualism*. Oxford: Clarendon Press.
- Dodd, D. (2010). Confusion about Concessive Knowledge Attributions. *Synthese*, 172, 381-396.
- Dougherty, T. and Rysiew, P. (2009). Fallibilism, Epistemic Possibility, and Concessive Knowledge Ascriptions. *Philosophy and Phenomenological Research*, 78, 123-132.
- _____ (2011). Clarity about Concessive Knowledge Attributions: reply to Dodd. *Synthese*, 181, 395-403.
- Dorr, C. and Hawthorne, J. (ms.) Embedding Epistemic Modals.
- Egan, A., Hawthorne, J. and Weatherson, B. (2005). Epistemic Modals in Context. In G. Preyer and G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (p. 131-167). Oxford: Clarendon Press.
- Fantl, J. & McGrath, M. (2002). Evidence, Pragmatics, and Justification. *The Philosophical Review*, 111, 67-94.

- _____ (2007). On Pragmatic Encroachment in Epistemology. *Philosophy and Phenomenological Research*, 75, 558-589.
- _____ (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Feldman, R. (2001). Skeptical Problems, Contextualist Solutions. *Philosophical Studies*, 103, 63-85.
- von Fintel, K. and Gillies, A. (2008). CIA Leaks. *The Philosophical Review*, 117, 77-98.
- Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis*, 23, 121-123.
- Grice, P. (1975). Logic and Conversation. In P. Cole and J. Morgan (Eds.), *Syntax and Semantics 3* (p. 41-58). New York: Academic Press. (Reprinted in Grice 1989, p. 22-40).
- _____ (1978). Further Notes on Logic and Conversation. In P. Cole (Ed.), *Syntax and Semantics 9* (p. 113-128). New York: Academic Press. (Reprinted in Grice 1989, p. 41-57).
- _____ (1989). *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Hawthorne, J. (2004a). *Knowledge and Lotteries*. Oxford: Clarendon Press.
- _____ (2004b). Replies. *Philosophical Issues*, 14, 510-523.
- _____ (2005a). Knowledge and Evidence. *Philosophy and Phenomenological Research*, 70, 452-458.
- _____ (2005b). Chance and Counterfactuals. *Philosophy and Phenomenological Research*, 70, 396-405.
- _____ (2005c). The Case for Closure. In M. Steup and E. Sosa (Eds.), *Contemporary Debates in Epistemology* (p. 26-42). Oxford: Blackwell.
- _____ and Stanley, J. (2008). Knowledge and Action. *The Journal of Philosophy*, 105, 571-590.
- _____ and Lasonen-Aarnio, M. (2009). Knowledge and Objective Chance. In P. Greenough and D. Pritchard (Eds.), *Williamson on Knowledge* (p. 92-108). Oxford: Oxford University Press.
- Hazlett, A. (2009). Knowledge and Conversation. *Philosophy and Phenomenological Research*, 78, 591-620.
- Hill, C.S. and Schechter, J. (2007). Hawthorne's Lottery Puzzle and the Nature of Belief. *Philosophical Issues*, 17, 102-122.
- Koriat et al. (1980). Reasons for Confidence. *Journal of Experimental Psychology*, 6, 107-118.

- Lasonen-Aarnio, M. (2008). Single Premise Deduction and Risk. *Philosophical Studies*, 141, 157-173.
- Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Lewis, D. (1980). Index, Context, and Content. In S. Kanger & S. Öhman (Eds.), *Philosophy and Grammar*. Reidel.
- _____ (1994). Humean Supervenience Debugged. *Mind*, 103, 473-490.
- _____ (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74, 549-567.
- MacFarlane, J. (2005). The Assessment Sensitivity of Knowledge Attributions. In T.S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology: Volume 1* (p. 197-233). Oxford: Oxford University Press.
- _____ (2007). Relativism and Disagreement. *Philosophical Studies*, 132, 17-31.
- _____ (2009). Non-Indexical Contextualism. *Synthese*, 166, 231-250.
- _____ (2011). Epistemic Modals Are Assessment Sensitive. In A. Egan and B. Weatherson (Eds.), *Epistemic Modality*. Oxford: Oxford University Press.
- _____ (f/c). Relativism and Knowledge Attributions. In S. Bernecker & D. Pritchard (Eds.), *The Routledge Companion to Epistemology*. London: Routledge.
- _____ (ms.) Varieties of Disagreement.
- Montminy, M. (2009). Contextualism, Invariantism, and Semantic Blindness. *Australasian Journal of Philosophy*, 87, 639-657.
- Musch, J. and Wagner, T. (2007). Did Everybody Know It All Along? A Review of Individual Differences in Hindsight Bias. *Social Cognition*, 25, 64-82.
- Nagel, J. (2008). Knowledge Ascriptions and the Psychological Consequences of Changing Stakes. *Australasian Journal of Philosophy*, 86, 279-294.
- _____ (2010a). Knowledge Ascriptions and the Psychological Consequences of Thinking about Error. *The Philosophical Quarterly*, 60, 286-306.
- _____ (2010b). Epistemic Anxiety and Adaptive Invariantism. *Philosophical Perspectives*, 24, 407-435.
- _____ (2011). The Psychological Basis of the Harman-Vogel Paradox. *Philosophers' Imprint*, 11, 1-28.
- Nickerson, R. (1999). How We Know – and Sometimes Misjudge – What Other People Know: Imputing One's Own Knowledge to Others. *Psychological Bulletin*, 125, 737-759.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Clarendon Press.

- _____ (2010). Contextualism, Scepticism and Warranted Assertibility Maneuvers. In J. K. Campbell, M. O'Rourke & H. Silverstein (Eds.), *Knowledge and Skepticism* (p. 85-104), Cambridge, Mass.: MIT Press.
- Reed, B. (2010). A Defense of Stable Invariantism. *Noûs*, 44, 224-244.
- Richard, M. (2004). Contextualism and Relativism. *Philosophical Studies*, 119, 215-242.
- Royzman, E. et al. (2003). 'I Know, You Know': Epistemic Egocentrism in Children and Adults. *Review of General Psychology*, 7, 38-65.
- Rysiew, P. (2001). The Context-Sensitivity of Knowledge Attributions. *Noûs*, 35, 477-514.
- _____ (2007). Speaking of Knowing. *Noûs*, 41, 627-662.
- Sadock, J. (1978). On Testing for Conversational Implicature. In P. Cole (Ed.), *Syntax and Semantics 9* (p. 281-297). New York: Academic Press. (Reprinted in Davis 1991, p. 365-376).
- Sainsbury, M. (1997). Easy Possibilities. *Philosophy and Phenomenological Research*, 57, 907-919.
- Sanna, L. and Schwartz, N. (2003). Debiasing the Hindsight Bias: The Role of Accessibility Experiences and (Mis)attributions. *Journal of Experimental Social Psychology*, 39, 287-295.
- Schaffer, J. (2004). Skepticism, Contextualism, and Discrimination. *Philosophy and Phenomenological Research*, 69, 138-155.
- _____ (2006). The Irrelevance of the Subject: Against Subject Sensitive Invariantism. *Philosophical Studies*, 127, 87-107.
- Schiffer, S. (1996). Contextualist Solutions to Scepticism. *Proceedings of the Aristotelian Society*, 96, 317-333.
- Sherman, S. et al. (1985). Imagining can Heighten or Lower the Perceived Likelihood of Contracting a Disease. *Personality and Social Psychology Bulletin*, 11, 118-127.
- Slovic, P. and Fischhoff, B. (1977). On the Psychology of Experimental Surprises. *Journal of Experimental Psychology*, 3, 544-551.
- _____ and Lichtenstein, S. (1982) Facts versus Fears: Understanding Perceived Risk. In D. Kahneman, P. Slovic, and A. Tversky (Eds.) *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- Sosa, E. (1999). How to Defeat Opposition to Moore. *Noûs*, 33, 141-153.
- Stanley, J. (2003). On the Linguistic Basis for Contextualism. *Philosophical Studies*, 119, 119-146.
- _____ (2005a). *Knowledge and Practical Interests*. Oxford: Clarendon Press.

- _____ (2005b). Fallibilism and Concessive Knowledge Attributions. *Analysis*, 65, 126-131.
- Stanovich, K. and West, R. (1998). Individual Differences in Rational Thought. *Journal of Experimental Psychology*, 127, 161-188.
- Vogel, J. (1990). Are there Counter-examples to the Closure Principle? In M. Roth and G. Ross (Eds.), *Doubting: Contemporary Perspectives on Skepticism* (p. 13-28). Dordrecht: Kluwer.
- _____ (1999). The New Relevant Alternatives Theory. *Philosophical Perspectives*, 13, 155-180.
- Weiner, M. (2006). Are All Conversational Implicatures Cancellable? *Analysis*, 66, 127-130.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- _____ (2005a). Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge. *The Philosophical Quarterly*, 55, 213-235
- _____ (2005b). Knowledge, Context, and the Agent's Point of View. In G. Preyer and G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (p. 91-114). Oxford: Clarendon Press.
- _____ (2005c). Replies to Commentators. *Philosophy and Phenomenological Research*, 70, 468-491.
- _____ (2009). Reply to John Hawthorne and Maria Lasonen-Aarnio. In P. Greenough and D. Pritchard (Eds.), *Williamson on Knowledge* (p. 313-329). Oxford: Oxford University Press.
- Yalcin, S. (2007). Epistemic Modals. *Mind*, 116, 983-1026.