# A Minimalist Approach to Epistemology

Christoph Friedrich Florian Kelp

Ph.D. Thesis

Department of Philosophy, University of Stirling

16 July 2007

# Acknowledgements

# Abstract

This thesis addresses the problem of the analysis of knowledge. The persistent failure of analyses of knowledge in terms of necessary and sufficient conditions is used to motivate exploring alternative approaches to the analytical problem. In parallel to a similar development in the theory of truth, in which the persistent failure to provide a satisfactory answer to the question as to what the nature of truth is has led to the exploration of deflationary and minimalist approaches to the theory of truth, the prospects for deflationary and minimalist approaches to the theory of knowledge are investigated. While it is argued that deflationary approaches are ultimately unsatisfactory, a minimalist approach to epistemology, which characterises the concept of knowledge by a set of platitudes about knowledge, is defended. The first version of a minimalist framework for the theory of knowledge is developed.

Two more substantive developments of the minimalist framework are discussed. In the first development a safety condition on knowledge is derived from the minimalist framework. Problems for this development are discussed and solved. In the second development, an ability condition is derived from the minimalist framework. Reason is provided to believe that, *arguably*, the ability condition can avoid the problems that beset traditional analyses of knowledge. It is also shown that even if this argument fails, minimalist approaches to epistemology may serve to provide a functional definition of knowledge. Reason is thus provided to believe that minimalist approaches to epistemology can make progress towards addressing the problem of the analysis of knowledge.

# Table of Contents

Foley's thought experiment is unsuccessful; there are lottery and Gettier cases in which the view makes inadequate predictions.

Minimalist theories of knowledge: 'knows that *p*' does signify a robust epistemic property; the concept of knowledge is characterised via a series of platitudes about knowledge: The factivity, belief, anti-luck, informative speech act, assertion, good informant, and closure platitudes.

The factivity platitude: One cannot know falsehoods; the belief platitude: one can only know what one also believes; the good informant platitude: Someone who knows that *p* is typically a good informant concerning *P*.

Facts about our use of the word 'know' in ordinary language and thought point to a commitment to the factivity, belief and good informant platitudes.

The job of the concept of knowledge is to flag good informants; the concept of knowledge is the objectivised version of the concept of good informant; Craig's characterisation of the concept of knowledge provides support for the factivity, belief and good informant platitudes.

The telling/informing platitude: If one informs/tells someone that *p*, then one represents oneself as knowing that *p*.

The telling/informing platitude can be extended to informative speech acts in general; the informative speech act platitude; the knowledge rule of informative speech acts; the informative speech act platitude and the knowledge rule of informative speech acts are mutually derivable.

Facts about our use of the word 'know' in ordinary language and thought point to a commitment to the informative speech act platitude; a Craig-style argument for the informative speech act platitude.

The knowledge rule of assertion: One must assert that *p* only if one knows that *p*.

The default use of declarative sentences is to make assertions; cases of standard uses of declarative sentences that are felicitous while their contents are not known.

The principle *Strongest Constative Speech Act*; speech act strength can be modelled on strength of represented attitude; knowledge is the strongest attitude normally represented in constative speech acts; all standard uses of declarative speech acts are constatives; the default use of declarative sentences is to make claims to knowledge.

The speech acts in the problematic cases are analysed as indirect speech acts.

The closure platitude: knowledge transmits across competent deduction; *prima facie* counterexamples to closure; the *prima facie* counterexamples are less persuasive than they initially appear to be; if the apparent counterexamples are genuine counterexamples, then the distribution or the equivalence principle fails and the informative speech act platitude/knowledge rule breaks down; Hawthorne on spelling out the closure principle in detail.

Facts about our use of the word 'know' in ordinary language and thought point to an implicit commitment to the anti-luck platitude: Cases of unreliably formed belief, lottery cases, Gettier cases.

A safety-based conception of knowledge is developed from the minimalist framework.

Sosa and Pritchard's versions of the safety-principle are too strong; Williamson's version of the safety-principle remedies the defects of Sosa and Pritchard's versions; Williamson's version of the safety-principle is translated into possible worlds terminology.

The safety condition can be derived form a plausible development of the anti-luck platitude and Pritchard's modal conception of luck.

Pritchard's taxonomy of types of epistemic luck: constitutive, capacity, doxastic, evidential, veritic.

Constitutive, capacity, doxastic, and evidential epistemic luck are compatible with knowledge

  Even if the safety condition is the anti-Gettier condition we have not given an
  analysis of knowledge in terms of necessary and sufficient conditions;
  minimalist safety has an advantage over traditional analyses in that it roots out
  the sources of the Gettier problem (luck) and derives the prospective anti-
  Gettier condition via an independently plausible conception of luck.

  Craig, Zagzebski and Floridi's argument to the effect that analyses of
  knowledge in terms of true belief plus X (where the third condition, X, is
  independent of the truth-condition) must fail; minimalist safety does not fall
  prey to this argument as safe belief is factive.

A case of a Gettiered but safe belief; safety isn't the anti-Gettier condition.

  The Gettiered subject lacks knowledge because he does not hit upon the truth
  through the exercise of his relevant cognitive abilities.

  It is consistent with the diagnosis of section 4.1 that the Gettiered subject's
  ignorance is rooted in luck.

  There is a sense of luck that does not require modal instability: luck relative to
  one's relevant abilities; the phenomenon of luck relative to one's relevant
  abilities is pervasive.

  Pritchard's modal conception of luck is slightly weakened; the weakened
  version holds out the hope of explaining the results of the scientific research on
  luck; the argument from the anti-luck platitude to the safety condition can be
  recovered; the weakened version does not license the converse derivation of
  the anti-luck platitude from the safety condition; it avoids commitment to the
  claim that safety fully captures the sense in which knowledge excludes luck;
  the problem posed by the case of Gettiered safe belief can be avoided.

  The ability and/or achievement platitudes allows the minimalist to incorporate
  the diagnosis of the subject's ignorance in the problematic Gettier case into the
  minimalist framework; the achievement platitude entails an ability condition on
  knowledge; the ability condition explains our intuitions in the cases that
  motivated the anti-luck platitude.

The ability condition is consistent with the closure platitude.

The ability condition is consistent with the achievement, ability and anti-luck platitudes.

Two objections are discussed and dismissed.

There are Gettier cases in which the subject satisfies the ability condition; the case of Archie the archer; the case does not establish what it is intended to establish.

Cases of testimonial knowledge are not cases of achievement; the case of Morris; the case does not establish what it is intended to establish.

# Introduction

What is knowledge? Whatever else it may be it is overwhelmingly plausible that knowledge is a kind of true belief. However, it is also widely agreed that true belief alone is not sufficient for knowledge. Additionally, true belief needs to be "tethered", to put it in Plato's words.[1] For the longest time it was taken for granted that the tether that knowledge requires is justification. That is to say that knowledge requires justified true belief. And conversely, it was also taken for granted that if one believes not only truly but also justifiably then one knows what one believes. Knowledge, then, is justified true belief. Let us call this 'the received view'. The received view is what philosophers have for the longest time taken to be the answer to the question as to what knowledge is.

Enter Edmund Gettier. In a short paper (1963), Gettier provided two examples that have since widely been accepted as a decisive refutation of the received view.[2] To be more precise, what these examples show is that the received view does not provide a sufficient condition for knowledge. In one of Gettier's examples a person, Smith, comes by evidence that his only competitor for a certain job, Jones, will get the job (perhaps the secretary has told him that Jones gets the job) and that Jones has ten coins in his pockets (perhaps he has just counted the coins in Jones's pockets). Certainly, these two propositions jointly entail that the person who gets the job has ten coins in his pockets. Suppose Smith sees that this entailment holds. He then has evidence that the person who will get the job has ten coins in his pockets. Suppose Smith forms a belief on the basis of this evidence. Smith's belief is clearly justified. But now suppose that, unbeknownst to Smith, he, *Smith*, gets the job rather than Jones. Suppose, furthermore, that Smith

---

[1] Cf. Plato (1956), p. 154.
[2] Noticeable exceptions are Stephen Hetherington (e.g. in (1998), (1999), (2001)) and Brian Weatherson (2003) who maintain that Gettiered subjects have knowledge. Moreover, Weinberg, Stich and Nichols (2001) claim to have unearthed empirical evidence that calls the reliability of the Gettier intuitions into doubt.

also happens to have ten coins in his pockets. If so, Smith's belief that the person who gets the job has ten coins in his pockets is not only justified but also true. However, Smith certainly does not know what he believes. After all, the evidence on which he bases his belief—that is, for the greatest part, evidence concerning Jones—have little bearing on the facts that make the proposition he believes true. Hence, contrary to the received view, justified true belief is not sufficient for knowledge.[3]

Gettier's paper was the source of a major upheaval in the theory of knowledge. Once it was clear that the received view is false, epistemologists have tried to find out what knowledge consists of instead. The task that epistemologists set for themselves was nothing less than to provide a full-blown analysis of knowledge, to specify a set of conditions that are individually necessary and jointly sufficient for knowledge. However, the task turned out to be much more difficult that it might have initially appeared to be. While a set of conditions that would deal with Gettier's two examples was not hard to come by, it has proved surprisingly hard to find a set of conditions that would deal with the phenomenon in general. Evidence for this claim can be gleaned from the way in which the debate over the analysis of knowledge has developed: In the run for the right analysis of knowledge, a huge number of proposals have been adduced. However, the proposals have tended to fall prey to Gettier-style counterexamples, that is, cases in which the subject satisfies the proposed set of individually necessary and jointly sufficient conditions but, intuitively, does not know (henceforth 'Gettier cases').[4] As a result of this development, some philosophers have become suspicious of the very prospects for the project of providing an analysis of knowledge in terms of necessary and sufficient conditions. Most notably, Timothy Williamson seems to take the upshot

---

[3] Gettier (1963), p. 121.
[4] Some useful surveys over the Gettier debate can be found in Slaght (1977), Shope (1983) and Lycan (forthcoming).

of the post-Gettier debate on the analysis of knowledge to be that no conceptual analysis of knowledge of this form is possible.[5]

However, Gettier's paper did not only initiate the quest for the condition that would deal with Gettier cases, it also spurred a debate over the status of the justification condition that played the role of the tether in the received view. The justification condition at issue in the received view is internalist. Roughly, internalist justification is justification that is in some salient sense internal to the knower's mind.[6] Accordingly, the debate over the status of the justification condition at issue in the received view has become known as the debate between *internalists*, who hold that internalist justification, justification of the kind at issue in the received view, is necessary for knowledge, and externalists, who maintain that at least in some cases knowledge can be had without internalist justification. It is no surprise that the justification condition at issue in the received view was construed along internalist lines. After all, it was generally acknowledged that reflective activity is central to knowledge. Even in Plato's version of the received view, the way in which true belief is tethered so as to be turned into knowledge is by "working out the reason" why it is true.[7] And the idea of working out the reason was understood as centrally involving reflective activity on the part of the knower.[8] It is not hard to see that the internalist justification condition accommodated the sense in which reflective activity is central to knowledge rather nicely.

---

[5] Cf. Williamson (2000), p. 2. The view that knowledge is not analysable is also defended in Craig (1990).

[6] There is disagreement as to how the precise sense in which justification is internal to the knower's mind is best spelled out. Some spell out the sense in which justification is internal in terms of an epistemic constraint: justification requires that the knower be able to come to know the facts that determine her justification by reflection alone. (This version of internalism has been defended, for instance, by Roderick Chisholm (e.g. (1989)) and Laurence BonJour (e.g. (1985)).) As opposed to that, others spell out the relevant sense in terms of a metaphysical constraint: Only mental states can function as justifiers. The most prominent defenders of this version of internalism are Earl Conee and Richard Feldman (e.g. (1985) and (2001)).

[7] Cf. Plato (1956), p. 154.

[8] Notice, by way of evidence for this point, that in his summary of the *Meno* W.K.C Guthrie claims that, according to Plato, "…right opinion … is converted into knowledge only when you have worked out the explanation for yourself and understand the reason it this is true." (Plato (1956), p. 106.)

However, as some of the cases in the post-Gettier literature suggest, whether or not one is in a position to know a certain proposition may depend rather dramatically on facts about the environment. Consider, by way of illustration, the following case due to Alvin Goldman: Henry is driving through the countryside, sees a barn and thereupon comes to believe that he is looking at a barn. His belief is justified. After all, it is based on excellent visual-perceptual evidence. At the same time, it is true: Henry is in fact looking at a barn. Unbeknownst to Henry, however, the barn he is looking at is the only real barn in a field full of barn façades that are so cleverly constructed that they cannot be distinguished from real barns by the naked eye of passers-by who are driving along the road. Intuitively, Henry's justified true belief does not qualify as knowledge. Henry is Gettiered.[9] Now contrast this case with the case of Benji who is also driving through the countryside, sees a barn at the side of the road and thereupon comes to believe that he is looking at a barn. In Benji's case, however, nothing fishy is going on. Intuitively, Benji's justified true belief does qualify as knowledge. What these two cases illustrate is, of course, just how dramatically knowledge may depend on environmental factors. The relation between Henry and Benji and their respective *immediate* environments with which they are in perceptual contact may be exactly the same. What makes the difference with respect to whether or not they know are only facts about the *wider* environment they find themselves in. Since facts about the environment that make the difference between Henry and Benji are not internal to their respective minds the lesson this sort of case teaches is that no strengthening of the internalist justification condition that retains the internalist identity of this condition will solve the Gettier-problem.[10] By the same token, internalist analyses of knowledge will have to supplement the received view with an additional condition in order to deal with Gettier cases. At the same time, externalists may take the lesson this sort of Gettier case teaches to have anti-internalist

---

[9] Goldman (1976), pp. 772-3.
[10] For a more detailed version of this argument see Pritchard (2001).

impact. For instance, they may argue that the insight that can be gleaned from this sort of Gettier case is that the tethering of knowledge is ensured not by internalist justification but rather by some other condition—*viz.* the one that deals with Gettier cases. If so, however, externalists may worry whether internalist justification is really requisite for knowledge. After all, why should we expect knowledge to be tethered in more than one way?

While none of these brief reflections will settle the debate between internalists and externalists[11], these reflections on the ramifications of Gettier's refutation of the received view suggest that, contrary to what many epistemologists appear to take for granted, the project of analysing knowledge need not be completed even if the condition that successfully deals with Gettier cases were discovered. Before we could lay claim to having closed the book on the analysis of knowledge we would also have to settle the debate between internalists and externalists about knowledge. (And, more obviously, settling the debate between internalism and externalism about knowledge would not thereby allow us to close the book on the analysis of knowledge since we need not and in all likelihood will not have determined the condition that deals with Gettier cases.) This is not to say that in either settling the debate between internalists and externalists or identifying the anti-Gettier condition we would not have made significant progress in the analytical project. However, it remains the case that we will not have completed the project of identifying a set of individually necessary and jointly sufficient conditions, unless we have addressed both issues in a satisfactory fashion.

It is noteworthy that, in consequence, each issue may turn out to be an insuperable obstacle to the analytical project. On the one hand, Williamson's suspicion

---

[11] On the contrary, the debate is still going on. For instance, Conee and Feldman (in Conee and Feldman (2004)) as well as Matthias Steup (in Steup (2004)) defend rather classical versions of internalism. However, see McDowell (1995), (2002) and Neta and Pritchard (forthcoming) for some non-standard defences of internalism. For recent defences of externalist conceptions of knowledge see Pritchard (2005), Williamson (2000), and Sosa (1999). For more on the internalism/externalism debate see also BonJour (1985), BonJour and Sosa (2003), Feldman (2005), Goldman (1999), Greco (2005), and Kornblith (2001).

may turn out to be true: There may just not be a condition that if conjoined with true belief will explain why Gettiered subjects lack knowledge in all Gettier cases. On the other hand, it may be that the debate between internalists and externalists is just irresolvable. In either case, the prospects for identifying a set of individually necessary and jointly sufficient conditions for knowledge will be dim.

Now it seems to me that the development of the Gettier literature—the series of proposed analyses of knowledge and Gettier cases refuting them—on the one hand, and the continuing debate between internalists and externalists on the other may give us reason to be somewhat pessimistic about the prospects for the analytical project. However, having reason to be pessimistic is one thing. Moving, as Williamson does, from the facts that give us reason to be pessimistic to the assumption that the analytical project must be unsuccessful is yet another. True, history has so far not played in our favour. However, that does not mean that it will continue to do so. For all we know, the solution to the Gettier problem may eventually be found and the debate between internalists and externalists may eventually be resolved. Accordingly, if we wanted to explore alternative paths to the project of analysing knowledge in terms of necessary and sufficient conditions, it would be desirable if the alternative were consistent with the possibility of an analysis of knowledge so conceived.

The aim of this thesis will be to explore two such alternatives. The inspiration for this project comes from a similar development in the theory of truth. The project of answering the question as to what truth is was long taken (and, by many, continues to be taken) to be the project of identifying a property that all and only true propositions have in common that explains why they are true. The literature on truth witnesses a variety of proposals concerning the relevant property: To name just a few of them, it has been argued that the relevant property is causal correspondence with the facts, that it is coherence with some designated set of propositions or justifiable belief under ideal

circumstances. Each view, however, faces problems in that the metaphysics of the various truth properties advocated does not mesh with the metaphysics of certain classes of apparently truth-apt statements. In reaction to these problems, some theorists of truth have explored alternative approaches to answering the question as to what truth is. Rather than attempting to identify a *property* that all and only true propositions have in common, the primary goal of these approaches is to provide a characterisation of the *concept* of truth. The question about the truth property is only of secondary importance, it is addressed only once the characterisation of the concept is in play. The two approaches to the theory of truth that I have in mind here are, of course, deflationary theories of truth and minimalist theories of truth (à la Crispin Wright). In this thesis, then, I will explore the prospects for deflationary and minimalist approaches to epistemology construed in a parallel fashion to deflationary and minimalist approaches to the theory of truth. I will argue that the prospects for deflationary approaches to epistemology are dim. However, I do think that a minimalist approach to epistemology is viable and I will defend one version of such an approach. The minimalist approach I will defend is consistent with there being a full-blown analysis of knowledge. In fact, I will argue that the minimalist framework I develop allows us to make significant progress with respect to at least one of the issues that need to be addressed in the traditional analytical project—*viz.* the Gettier problem—in that it enables us to trace the sources of the problem and tackle it at its roots. I will also argue that the results the minimalist framework allows us to secure with respect to the Gettier problem in conjunction with some independently plausible further assumptions can be used to develop a substantial condition on knowledge that, *arguably*, deals with Gettier cases. At the same time, this argument is independent of the basic framework. For that reason, even if it turned out to be unsound, and it turned out that there could be no substantial

anti-Gettier condition, the minimalist approach would still be illuminating the concept of knowledge.

I will begin with a short chapter on the source of my inspiration for this thesis in which I discuss the credentials of deflationary and minimalist theories of truth. In chapter II, I will construe deflationary approaches to epistemology in a parallel way to deflationary approaches to the theory of truth. Two conceptions of knowledge, one due to Crispin Sartwell, the other originally due to Richard Foley and then developed by Duncan Pritchard, are identified as satisfying the criteria of deflationary approaches to epistemology so construed. Both views are discussed and dismissed as unsatisfactory. In chapter III, I outline the first version of a minimalist theory of knowledge construed in a parallel way to Wright's minimalist theory of truth. The concept of knowledge is characterised by a series of platitudes, that is, of highly intuitive principles, relating knowledge to truth, belief, good informants, informative speech acts, assertion, competent deduction, and luck. Support for the platitudes is provided, first, by facts about our use of the word 'know' in ordinary language and thought that point to an implicit commitment to the platitudes, and, second, by theoretical arguments. With the basic minimalist framework in play, I go on in chapter IV to devise two arguments that develop a non-minimalist condition on knowledge—*viz.* the so-called 'safety condition'—from the minimalist framework. The first argument, due to Pritchard, employs a modal conception of luck to derive the safety condition from the platitude relating knowledge and luck, while the second argument exploits some of the considerations Edward Craig adduces to support the claim that the concept of knowledge is the objectivised version of the concept of good informant in order to derive the safety condition. It is also argued that the safety condition is consistent with the other platitudes of the minimalist framework so that there is strong reason to believe that the safety condition is indeed a substantive necessary condition for knowledge.

Chapter V provides a closer scrutiny of the safety condition as developed from the minimalist framework. It is argued that we may expect the safety condition so developed to explain the features of our ordinary use of the word 'know' that motivated the platitude relating knowledge and luck. Although there is *prima facie* reason to believe that that safety condition meets the expectations, the reasons turn out to be *ultima facie* defeated. A case of one of the types for which the platitude relating knowledge and luck promises to explain our unwillingness to apply the word 'know' is provided in which, at the same time, the subject has a safe belief. It is argued that the case is an instance of a more general phenomenon that provides us with excellent reason to believe that Pritchard's modal conception of luck fails. A weakened version of the modal conception of luck is introduced that relieves the safety condition from the expectation to explain all the features of our ordinary language use of the word 'know' that motivated the platitude relating knowledge and luck. At the same time, the weakened version of the modal conception of luck still allows the argument from that platitude to the safety condition to go through. A diagnosis of the subject's ignorance is given by appeal to the idea the subject does not hit upon the truth through the exercise of his relevant cognitive abilities. In chapter VI I argue that this diagnosis can be fitted into the minimalist framework by adopting a platitude relating knowledge and cognitive achievements. It is argued that achievements require success through a set of relevant abilities. An ability condition on knowledge is derived. Since this ability condition requires exactly what, according to the earlier diagnosis, the subject in the problematic Gettier case lacks the diagnosis can thus be fitted into the minimalist framework. It is also argued that success through the exercise of a set of relevant abilities excludes luck relative to one's relevant performance. Reason is provided to believe that the type of luck excluded by the ability condition is the type of luck at issue in the anti-luck platitude and that therefore the ability condition explains all the phenomena that

motivated the anti-luck platitude. The ability condition is shown to be consistent with the other platitudes of the minimalist framework.

# I Deflationism and Minimalism in the Theory of Truth

In this chapter I will look at deflationary and minimalist approaches to the theory of truth. More specifically, I will outline the credentials of such theories and look at arguments that have been made to support them. This chapter provides a basis for the following chapters in which I will attempt to construe parallel deflationary and minimalist approaches to the theory of knowledge and explore their prospects.

## 1  Deflationism in the Theory of Truth

Deflationism in the theory of truth is often characterised as the view that truth has no nature. A related and more precise way of characterising deflationary theories of truth is as the view that the predicate 'is true' does not signify a robust property: There is nothing that all sentences/statements to which the predicate 'is true' applies have in common that explains why the predicate 'is true' applies to them.[1] According to deflationists, the classical approach to the theory of truth, which starts from the question 'What is truth?'—where this question is understood along the lines of 'What is the nature of truth?'—is the wrong way into the subject matter.

But how ought we to start a philosophical inquiry into truth if not by asking the question as to what its nature is? If we disallow this question about truth is not the basis for a proper philosophical inquiry into truth undermined? The deflationists' answer here is, one won't be surprised to be informed, 'no'. The question that deflationists want to take as a starting point for a philosophical inquiry into truth is: What is the role of the concept of truth and the truth predicate expressing it in our ordinary language and

---

[1] Cf. Stoljar (1997).

thought? In this way deflationists aim to provide a characterisation of our concept of truth.[2]

It is noteworthy that if deflationists succeed in providing a characterisation of our concept of truth by specifying the role the concept of truth and the truth predicate expressing it play in our ordinary thought and language use without invoking a robust nature of truth, then there is a *prima facie* case for deflationary theories of truth. After all, first, if truth has a more substantive role to play than the one it plays in our ordinary thinking and language use that has been captured by deflationary theories, then that certainly needs argument. At the same time, second, deflationary theories of truth are metaphysically more lightweight in the sense that they incur fewer metaphysical commitments than their inflationary competitors. And parsimony enjoins us, all else equal, to favour the more lightweight theory. So, if deflationary theories do succeed in characterising the role of the concept of truth and the truth predicate expressing it in our ordinary thinking and language use, the burden of proof is on defenders of robust theories of truth to identify what further role truth has to play and why deflationary theories won't be able to explain how it could do so.

Let me pause here for a brief methodological reflection. What the last paragraph suggests is that deflationists can appeal to the metaphysically lightweight character of their theories of truth in order to shift the burden of proof onto the shoulders of their inflationist competitors. In this way deflationists gain what Pritchard has called a *dialectical* advantage over inflationists. By the same token, it can be seen that deflationary theories, even if, eventually, they do not work, have an important dialectical role to play. By considering deflationary theories and uncovering why they do not work, we also uncover the reasons why we need an inflationary theory.[3]

---

[2] Cf. Armour-Garb and Beall (2005), p. 2.
[3] Cf. Pritchard (2004*a*), pp. 103-4.

Apart from the dialectical advantage over their inflationary competitors, however, deflationists also claim to be able to resolve a central problem that besets inflationary theories of truth. It seems that the claims that inflationists make about the nature of truth will bring along commitments to the metaphysics of discourses with truth-apt sentences/statements. For instance, if one defends a causal correspondence theory of truth the metaphysics of discourses with truth-apt sentences/statements had better be such as to allow there to be a causal correspondence relation between (components of) sentences/statements of such discourses and the entities they countenance. However, for each inflationary theory there are some discourses the sentences/statements of which are *prima facie* truth-apt while, at the same time, the nature of the entities they countenance appear not to mesh with the nature of truth countenanced by the inflationary theory.[4] In the case of the causal correspondence theory, for instance, one such discourse is mathematical discourse. There is strong *prima facie* reason to believe that mathematical sentences/statements are truth-apt. At the same time, it seems that the entities it countenances—such as, for instance, numbers—do not stand in a relation of causal correspondence to anything. Deflationary theories of truth can easily avoid these problems. There is simply no question of how the nature or metaphysics of truth matches with the metaphysics a certain discourse with apparently truth-apt sentences/statements because truth, according to deflationists, has no nature. So, another advantage of deflationism over its inflationary competitors is that it can dissolve certain metaphysical problems that beset inflationary theories.[5]

Whether or not deflationists will be able to claim any of these advantages over their inflationist competitors depends, as I have already indicated, on whether or not they succeed in providing an adequate characterisation of the concept of truth—that is,

---

[4] This has becomes known as the problem of scope. (Cf. Lynch (2001*a*), pp. 723-4 and Wright (1994), pp. 7-12.)
[5] Cf. Wright (1994), pp. 12-13.

on whether or not they succeed in explaining what role the concept of truth and the truth predicate expressing it play in our ordinary thought and language use. How do we judge whether a given deflationary theory of truth has provided an at least *prima facie* adequate characterisation of the concept of truth? I suggest that in order to do so deflationists must at the very least explain:

(a) what our understanding of the concept of truth consists in;
(b) how our understanding of the concept of truth makes it possible for us to ascribe truth in thinking and language use in the way that we do;
(c) why we have a concept of truth in the first place.

For the purposes of this thesis, there is no need get into the details of how the deflationists' explanations are supposed to work. Rather a rough outline shall suffice. A crucial role in the deflationists' explanations is played by the so-called equivalence schema (I will let '<A>' stand for a name of a statement/sentence and 'A' for the statement/sentence named by <A>):

(ES)    <A> is true iff A.

Deflationists maintain that instances of the equivalence schema are conceptually and explanatorily fundamental to the concept of truth. That means that the concept of truth cannot be defined in terms of more basic concepts, that there can be no further question as to why the equivalences hold.[6] Some deflationists even go so far as to construe instances of the equivalence schema as implicitly defining the concept of truth.[7]

Deflationists appeal to instances of the equivalence schema to provide an account of our understanding of the concept of truth. They claim that our understanding of the concept of truth consists in our tendency to accept instances of the equivalence schema. They use this account of our understanding of the concept of truth to explain

---

[6] Cf. Armour-Garb and Beall, pp. 2-3.
[7] Cf. Horwich (1990).

how we can make truth ascriptions in ordinary language and thought in the way we do. It appears to be commonly agreed among theorists of truth that the difficult types of truth ascriptions for deflationists—that is, the types in which truth is not merely ascribed to the quote-name of a sentence/statement—are indirect assertions, endorsements, and generalisations. Deflationists argue that instances of the equivalence schema are all that is needed in order to explain how such truth ascriptions are possible. We may then expect that such truth ascriptions can be made by subjects who have a tendency to accept instances of the equivalence schema and in this way, according to deflationists, understand the concept of truth. Finally deflationists point out that we could not make some indirect assertions, endorsements and generalisations unless we had the concept of truth. In this way they explain why we have the concept of truth in the first place.[8]

If all this is correct, deflationists will be able to discharge the explanatory burden we imposed on them: They can explain (a) what our understanding of the concept of truth consists in in terms of our tendency to accept instances of the equivalence schema; (b) how we can make the truth ascriptions we make in ordinary language and thought in terms of their account of our understanding of the concept of truth; and (c) why we have a concept of truth in the first place by pointing out that there are certain things we say in ordinary language and think in ordinary thought that we could not say or think unless we had a concept of truth. In this way, if the above explanations work in the way envisaged, deflationists may be said to have provided an at least *prima facie* adequate characterisation of the conceptual role of truth in our ordinary language and thought. There is, then a case to be made for a deflationary approach to the theory of truth.

---

[8] My outline her proceeds rather quickly. A more detailed account of how the deflationists' arguments work can be found in Armour-Garb and Beall (2005), pp. 3-6.

## 2  Minimalism in the Theory of Truth

There are at least two theories of truth that go under the label 'minimalism'. On the one hand, Paul Horwich refers to his version of deflationism as 'minimalism' and, on the other hand, Crispin Wright also labels his inflationary theory of truth 'minimalism'. Since I am interested in minimalism as a non-deflationary theory, rather than in being able to distinguish between different kinds of deflationary theory, I will adopt Wright's use of the term 'minimalism'.

Wright convinces himself that deflationary theories of truth must fail. In a nutshell, he argues that truth registers a distinct norm for assertion and that therefore truth has an explanatory role to play which deflationary theories cannot account for.[9] At the same time, Wright takes seriously the metaphysical problems that beset traditional inflationary theories of truth and traces these problems back to the fact that the metaphysical commitments they incur are too strong. What he tries to do in reaction to these problems is to develop a conception of truth that, whilst being inflationary, does not incur the problematic metaphysical commitments of traditional inflationary theories.

At first glance, it may seem as if Wright's project is hopeless. Recall that there are two questions a theory of truth may address: The first question concerns the nature of truth, while the second one concerns the role of the concept of truth and the truth predicate expressing it. Deflationists regard the first question as misguided—they want to approach the theory of truth by answering the second question. If deflationary theories fail, and, hence, if truth does have a nature, then it seems as if we must be able to say something by way of response to the first question. However, if we say something in response to the first question and nail down the nature of truth, are we not

---

[9] Wright (1994), pp. 12-24.

bound to confront the same metaphysical problems as traditional inflationary theories of truth? And if so, is Wright's project not hopeless?

Fortunately, the answer to this question is 'no'. Wright's trick to avoid the problems of traditional inflationary theories is that he lets the answer to the first question about truth—that is, the question concerning its nature—be determined by the answer he gives to the second question—that is, the one concerning the concept of truth. More specifically, Wright can be understood as characterising the concept of truth by a set of platitudes about truth.[10] Platitudes are, according to Wright, "very general, very intuitive principles"[11] which specify the conceptual role of truth and in this way pin down the concept of truth. Wright offers the following as platitudes about truth:

> That to assert is to present as true;
> That any truth-apt content has a significant negation which is likewise truth-apt;
> That to be true is to correspond to the facts;
> That a statement may be justified without being true, and vice versa…[12]

If a predicate satisfies Wright's platitudes, it plays the conceptual role of truth and for that reason qualifies as a truth predicate. At the same time, since Wright's proposal is non-deflationary, the resulting truth predicate will also signify a robust property. In this way truth does have a nature, *viz.* the one encapsulated in the robust property signified by the truth predicate.

This approach to truth leaves Wright in a position that is comfortable in at least the following ways: First, by going inflationary, Wright can avoid what he takes to be an insuperable problem for deflationary theories of truth. Second, Wright can adopt for his own purposes both arguments that deflationary theorists have adduced in support of their theory. Since whether or not a predicate is a truth predicate depends only on

---

[10] The idea that the platitudes characterise the conceptual role of truth is even clearer in Michael Lynch's development of Wright's approach to truth (cf. Lynch (2005)).
[11] Wright (1994), p. 34.
[12] Ibid. Note that Wright thinks that at the end of the day only the first and the second of these platitudes will be needed to characterise the concept of truth since the other platitudes as well as the equivalence schema are validly derivable from them (cf. pp. 34-35, 61-64).

whether or not it satisfies the platitudes, and since the truth property is the property signified by the truth predicate, the metaphysical commitments that come along with the truth property depend, ultimately, on platitudes about truth. Since, plausibly, such a strategy will not lead one to metaphysical excess but rather to comparatively lightweight and well-motivated metaphysical commitments, Wright can claim a dialectical advantage over traditional inflationary theories of truth. Moreover, Wright can hope to solve the metaphysical problems that beset traditional inflationary theories. The key to Wright's solution is that his minimalist approach to truth is consistent with *pluralism* about truth.[13] After all, different predicates may satisfy the platitudes about truth in different discourses. If so, different properties will be signified by these predicates. Truth, then, may have multiple natures. In this way, the fact that the metaphysics of a certain property does not mesh with the metaphysics of a given apparently truth-apt discourse need not count against that property's being signified by the truth predicate in some other discourse. For instance, while the fact that the property of causal correspondence appears not to mesh with the metaphysics of mathematical discourse can be regarded as evidence that causal correspondence is not the property signified by the truth predicate in mathematical discourse, it need not count against the idea that this property is signified by the truth predicate in some other discourse—for instance, in the discourse about middle-sized dry goods. So, there is hope that Wright's minimalist approach to truth can avoid the metaphysical problems that traditional inflationary theories encounter.

---

[13] Cf. Wright (1994), p. 38.

# II Epistemic Deflationism

Having outlined how deflationary and minimalist theories are implemented in the theory of truth, I would now like to turn to the theory of knowledge. What I will try to do is to construe deflationary and minimalist approaches to the theory of knowledge in a way that parallels deflationary and minimalist approaches to the theory of truth. More specifically, I will first outline a way of implementing a deflationary thesis in the theory of knowledge and argue that the prospects for it are rather dim. Note that even if the prospects for deflationary theories of knowledge are dim, there is still something to be learned from the discussion. Recall that we found that deflationary theories, even if they do not work, have an important dialectical role to play in that they help us to identify the reasons for accepting inflationary theories. In this way, by considering and dismissing deflationary theories, we can ensure that our cherished inflationary theories are not simply based on dogma and false preconception. Having dealt with deflationary approaches to the theory of knowledge, I will then, in the next chapter, turn to minimalist approaches to epistemology. In Wrightian spirit, I will provide a set of platitudes about knowledge which fix the concept of knowledge. First things first, however, let me begin with a discussion of deflationary approaches to the theory of knowledge.

Recall that deflationism in the theory of truth was characterised as the view that the predicate 'is true' does not signify a robust property in the sense that there is nothing that all sentences/statements to which 'is true' applies have in common that explains why 'is true' applies to them. Accordingly, deflationism in the theory of knowledge construed in a parallel way can be characterised the view that the predicate 'knows that $p$' does not signify a robust property in the sense that there is nothing that all subjects to which the predicate 'knows that $p$' applies have in common that explains why 'knows

that *p*' applies to them. It seems to me, however, that the prospects for a deflationary theory of knowledge so construed are rather dim. After all, it is very plausible that all subjects to which the predicate 'knows that *p*' applies also have a true belief that *p* and that the fact that they have a true belief that *p* in part explains why the predicate 'knows that *p*' applies to them.[1]

However, even if we do not construe deflationism in the theory of knowledge along those lines, we may be able to characterise a more restricted type of deflationism about knowledge. According to this restricted type, 'knows that *p*' does not signify a robust *epistemic* property in the sense that there is nothing *epistemic*—that is to say, nothing over and above true belief such as, for instance, justification—that all subjects to which 'knows that *p*' applies have in common that explains why 'knows that *p*' applies to them. Yet another way of characterising this type of deflationism about knowledge is to say that knowledge does not have an epistemic nature.[2]

---

[1] The factivity of knowledge is usually taken as so overwhelmingly plausible that it is hardly ever discussed. Notable exceptions are Keith DeRose (2002, pp. 7-8) who argues against the idea that truth might just be a conversational implicature of knowledge claims rather than a necessary condition for knowledge and Kelp and Pritchard (forthcoming) who discuss the prospects for a weakened version of the factivity principle to block the so-called paradox of knowability. As regards the belief condition the consensus is slightly weaker. Colin Radford (1966), for instance, provides what he takes to be counterexamples to the belief condition. However, these counterexamples are of dubious persuasiveness. Since space is limited I will not discuss them here. Rather, I refer the reader to Williamson ((2000), p. 42) and Pritchard ((2004*a*), pp. 105-6) for further discussion. I will also argue that a certain argument due to Craig provides a rationale for the factivity and belief conditions (chapter III, section 1.2).

[2] It would seem that if a deflationary theory of knowledge of this kind is true, the debate over the status of the internalist justification condition at issue in the received view will immediately be settled: If a deflationary theory of knowledge is true, then so is externalism about knowledge. Now a critic may concede that if a deflationary theory of knowledge is true, then internalist justification does not enter into the *analysis* of knowledge. After all, it is plausible that each part of the analysis contributes towards the explanation of why the subject knows. Since, by the lights of deflationary approaches to knowledge so construed, internalist justification has no such part to play it does not enter into the analysis. However, the critic may point out that by the lights of at least some deflationary theories of knowledge, internalist justification may still be *necessary* for knowledge. After all, the conditions on knowledge the deflationist countenances may be such that one could not satisfy them without also having internalist justification for the proposition known. (One such deflationary theory would be a theory according to which one knows that *p* if and only if one has maximally accurate and comprehensive true beliefs (cf. Foley's Sally case, pp. 38-9).) Moreover since in the introduction the debate between internalists and externalists was characterised as a debate over whether or not internalist justification is necessary for knowledge, it is not the case that deflationism about knowledge entails externalism. By way of response to this objection, notice that it seems no less correct to describe the debate between internalists and externalists about knowledge as a debate over whether the internalist justification condition enters into the analysis of knowledge. In consequence, it seems also correct to say that if a deflationary theory of knowledge is true, then so is externalism.

Recent epistemological literature witnesses at least two conceptions of knowledge that can be understood as deflationary in this restricted sense. The first one is due to Sartwell: It analyses knowledge that *p* as true belief that *p*. The second one has first been proposed by Foley and then been developed by Pritchard: Foley and Pritchard construe knowledge that *p* as true belief that *p* that is embedded in further true beliefs about the neighbourhood of *P*.[3] Both views reject the idea that 'knows that *p*' must also signify an epistemic property of the subject—such as the property of justifiably believing that *p*, of having arrived at the belief that *p* via a reliable method and the like—that explains why 'knows that *p*' applies to the subject. For that reason they are naturally understood as deflationary theories of knowledge in this restricted sense.

Just as in case of deflationary theories of truth we may expect deflationary theories of knowledge to provide at the very least a *prima facie* adequate characterisation of the concept of knowledge. In order to be *prima facie* adequate a characterisation of the concept of knowledge must at least be able to explain the following:

(a*)   what our understanding of the concept of knowledge consists in;
(b*)   how our understanding of the concept of knowledge makes it possible for us to ascribe knowledge in thinking and language use in the way that we do;
(c*)   why we have a concept of knowledge in the first place.

If either Sartwell's or Foley/Pritchard's deflationary conception of knowledge is to enjoy *prima facie* plausibility, then it must be able to explain (a*) to (c*). I take it that neither conception will have great difficulties in explaining what our understanding of the concept of knowledge consists in. Both Sartwell and Foley/Pritchard can claim that our understanding of the concept of knowledge consists in an (implicit) grasp of the biconditionals offered by their accounts. The problems with their views that I will identify lie, rather, in a mismatch between the knowledge ascriptions we make in

---

[3] It is noteworthy, however, that Pritchard is merely exploring the view rather than endorsing it.

ordinary thought and language use and the conditions that, according to these deflationary theories, we grasp in our understanding of the concept of knowledge: If our understanding of the concept of knowledge really consisted in what these deflationary theories claim it to consist in, the knowledge predicate expressing the concept of knowledge would be deployed in a different way than it actually is. In this way, I will attempt to show that both deflationary theories fail to explain how our understanding of the concept of knowledge makes it possible for us to ascribe knowledge in thinking and language use in the way that we do in a fully satisfactory way, that, therefore, they do not provide a *prima facie* adequate characterisation of the concept of knowledge, and that, therefore, these deflationary theories of knowledge fail. Leaving the story about the relation between understanding and knowledge ascriptions aside, what I will be doing can also be described in much simpler terms: I will attempt to identify *counterexamples* to the proposed conceptions of knowledge—cases in which we would say that the subject is ignorant while the theory at issue predicts that she knows (or, conversely, cases in which the theory predicts that the subject is ignorant while we would say that she knows). And, having explained how the relation between understanding and knowledge ascriptions is supposed to work once, I will henceforth use the simpler terminology to describe what I am doing whenever possible.

In what follows I will outline both Sartwell's and Foley/Pritchard's conceptions of knowledge in more detail and provide my arguments against them along the lines just indicated. Let me start with Sartwell.

# 1 Sartwell's Epistemic Deflationism: Knowledge is Merely True Belief

The most straightforward way to spell out a theory according to which there is nothing epistemic that explains why 'knows that *p*' applies to all subjects to which it does apply is by analysing knowledge that *p* as true belief that *p*. On this proposal, then, one knows that *p* if and only if one truly believes that *p*. This conception of knowledge has been advocated by Sartwell. Sartwell's argument for it has two parts: a negative one which is supposed to establish that there is no pre-analytic commitment to a distinction between true belief and knowledge. The second part of Sartwell's argument is positive: Sartwell attempts to establish that knowledge could not be anything but true belief.

## 1.1 An Outline and Critique of Sartwell's Negative Argument

Let us turn to the negative part of Sartwell's argument first. Sartwell aims to establish that there is no pre-analytic commitment to a distinction between true belief and knowledge and hence no such commitment to an epistemic condition on knowledge either.[4] That is to say, as far as our intuitions about knowledge and our ordinary

---

[4] A few remarks by way of clarification are in order. To begin with in order to bring his point home, Sartwell argues that there is no pre-analytic commitment to a *justification* condition on knowledge. For his argument to establish that there is no pre-analytic commitment to a distinction between true belief and knowledge, a number of assumptions must be in place. First, since no matter how the term 'justification' is interpreted, it is plausible that a justification condition is an epistemic condition on knowledge, it must be assumed that there are no further non-epistemic conditions on knowledge. Although this assumption is far from uncontroversial (see Foley and Pritchard's deflationary approach to epistemology as well as Hawthorne (2004) and Stanley (2005) who argue that there is condition concerning practical interest on knowledge), for present purposes, I will let Sartwell have it. Second, and more importantly, however, it must be assumed that the term 'justification' is used as a placeholder for any epistemic condition that may be requisite for knowledge. Otherwise, even if Sartwell managed to show that there is no pre-analytic commitment to a justification condition on knowledge, he would not thereby have established that there is no such commitment to a distinction between knowledge and true belief.

Unfortunately, however, Sartwell's use the term 'justification' cannot always be understood in this way. For there is one part of his discussion in which Sartwell's uses of 'justification' in a sense that presupposes that justification is something that is determined by facts that are—at least typically—accessible to the subject when she has justification—*viz.* when he construes justification as what he calls a criterion for knowledge. Sartwell's idea here is that when a subject makes a claim to knowledge her claim may be challenged by asking how she knows. If she is unable to cite facts in support of her claim—

language use of the word 'know' are concerned there is no such commitment. To begin with, Sartwell is aware of the fact that there are cases that provide *prima facie* reason to believe that, contrary to his claim, there is such a commitment. The cases he discusses include cases in which the subject acquires a true belief that *p* via freak methods, such as dreaming *P* up, guessing *P*, or reading *P* from the tea leaves, cases in which the subject is conditioned to assent to *P* in the right circumstances, and cases in which the subject acts as if she believed *P*.[5]

Sartwell's treatment of these cases is twofold. His first step in dealing with them is to beef up the notion of belief: First, beliefs must always be integrated into a web of beliefs and can never stand in isolation. Second, belief requires serious commitment on the part of the subject. And, third, belief requires further knowledge. Sartwell deals with some of the *prima facie* counterexamples by pointing out that as they are described they violate at least some of those conditions on belief. At the same time, he claims that if the cases are described in such a way that the subjects count as properly believing, it is less clear that our intuition that they are not cases of knowledge remains.[6]

Sartwell does concede that if the cases are described in such a way that the subject does properly believe the proposition at issue, then she also knows. That is not to say, however, that he has run out of resources for explaining any remaining intuitions

---

or, in other words, if she is unable to produce *justification*—then we are entitled to deny that she knows: The subject fails to satisfy the criterion for knowledge. However, if subjects could have justification whilst frequently failing to be able to access the facts that determine her justification when challenged, then justification would certainly be a poor criterion for knowledge. (At least it would be a poor criterion for establishing the inappropriateness of a knowledge claim—never mind for warranting the denial of such a claim. Since, as we will see in due course, Sartwell needs the criterion to work in this way, the point that justification would then be a poor criterion for knowledge remains.)

However, the two uses of 'justification' to which Sartwell is committed are importantly distinct. In particular, there are some candidate epistemic conditions on knowledge such that the facts that determine whether the subject satisfies them may not typically be accessible by the subject when she satisfies them. At the same time, in order to show that there is no pre-analytic commitment to a distinction between knowledge and true belief, Sartwell needs to show that there is no pre-analytic commitment to such epistemic conditions either. Hence, the two uses of 'justification' are importantly distinct. In order to demarcate the distinction, I will use 'placeholder-justification' or 'P-justification' for the first use of 'justification' to which Sartwell is committed and 'criterion-justification' or 'C-justification' for the second such use.

[5] Cf. Sartwell (1991), pp. 158-8.
[6] Cf. Sartwell (1991), pp. 158-9.

that the subject does not know. A second resource for explaining away our intuitions in such cases can be found in the role Sartwell assigns to justification—that is, in the present context, *C-justification*.[7] According to Sartwell, C-justification is a *criterion* for knowledge rather than a necessary condition in much the same way in which yielding a certain taste when bitten is a criterion (but not a necessary condition) for being gold. If we are in doubt as to whether a certain subject has knowledge, we ask for (C-) justification. If she provides C-justification her claim to knowledge is acceptable—just as a claim that a certain object is gold is acceptable if it yields the right taste when bitten. If, on the other hand, the subject cannot provide C-justification, then we are entitled to deny that she knows—just as we are entitled to deny that an allegedly golden object is gold when it does not yield the right taste when bitten.[8]

A natural way of explaining away any remaining intuitions of ignorance in the problematic cases is the following: The subjects in such cases lack C-justification for their beliefs. Since C-justification is requisite for a claim to knowledge to be appropriately sustainable, that means that these subjects cannot appropriately sustain their claims to knowledge. That is to say, however, that, while the subjects meet the conceptual conditions for knowledge, their claims to knowledge are pragmatically defective. However, our intuitions about knowledge are indiscriminate between semantic and pragmatic sources. So, any remaining intuitions that the subjects in the *prima facie* counterexamples do not know can be explained by appeal to the fact that their claims to knowledge are pragmatically defective.[9]

To complete his argument that there is no pre-analytic no commitment to a distinction between true belief and knowledge, Sartwell adduces cases in which we would—or at least some of us might—say that the subject knows although she has no

---

[7] I introduced the term 'C-justification' in footnote 4 (p. 23).
[8] Cf. Sartwell (1991), pp. 161-2.
[9] Sartwell does not make this point in so many words. However, the general gist of it can be found in Sartwell (1991), p. 161.

justification—that is, now, *P-justification*—for her belief. These cases include cases in which people come to truly believe mathematical propositions in "a flash of insight"[10], cases of true religious belief and cases in which people truly believe a proposition, while no P-justification is available to them, then find out that their belief is (P-) justified and go on to claim that they knew all along.[11]

In view of the fact that the *prima facie* evidence in favour of a condition on knowledge over and above true belief is far from decisive in conjunction with the *prima facie* evidence against such a condition Sartwell concludes that there is no pre-analytic commitment to a distinction between knowledge and true belief. The negative argument is completed.

Before moving on to the positive part of Sartwell's argument, let me say something about the negative part. To begin with, notice that, as regards the *prima facie* problematic cases for this theory, Sartwell discusses only cases that can plausibly be described as cases of *unjustified*[12] true belief. However, the threat for his theory comes not only and, I submit, not even mainly from cases of unjustified true belief. Rather, the main threat comes from cases in which subjects have (in some sense of 'justified') justified true beliefs that, intuitively, do not qualify as knowledge. After all, if these cases show that not even *justified* true belief is sufficient for knowledge, then they certainly also show that true belief is not so sufficient either.

Now, we have already encountered one type of case in which subjects have (in some sense of 'justified') justified true beliefs that, intuitively, do not qualify as knowledge, *viz.* Gettier cases. For instance, recall the case of Henry who drives through the countryside, sees the only real barn in a field full of barn façades and comes to believe that he is facing a barn. Henry has a true belief that is justified (in some sense of

---

[10] Sartwell (1991), p. 160.
[11] Cf. Ibid.
[12] It is plausible that the relevant subjects' beliefs are neither C- nor P-justified.

'justified')—and hence a true belief—that, intuitively, does not qualify as knowledge. Let us ask whether either of Sartwell's strategies will serve to explain away our intuition of ignorance in this case. To begin with, let us ask whether Sartwell may be able to maintain that Henry does not really believe that he is facing a barn. There is excellent reason to believe that he cannot. After all, Henry's belief may be as well integrated into his web of beliefs as any other perceptual belief he has. Moreover, he may be as strongly committed to his belief as to any other of his perceptual beliefs. Thus, if Henry has any perceptual beliefs at all, then, in the present case, he also believes that he is facing a barn. Moreover, it seems that Henry also satisfies Sartwell's criterion for knowledge. After all, if Henry were to make a claim to knowledge and were to be challenged by being asked how he knows, Henry may defend his knowledge claim by adducing the justification he has for his belief—*viz.* his perceptual experiences as of a barn and the relevant background knowledge about the looks of barns, that barns can be found in fields and so on.[13] So, by Sartwell's lights, his knowledge claim remains in good standing. We therefore have a first type of case in which, by the lights of the true belief account, subjects know, while Sartwell's strategies of explaining away our intuitions of ignorance fail.

There is yet another type case that poses a problem for Sartwell, one that revolves around lotteries. There is a firm intuition—indeed one of the firmest intuitions we have about knowledge—that we cannot know, in advance of the drawing, that a

---

[13] That is to say, Henry's belief is C-justified. Now one might wonder whether Sartwell could beef up the notion of C-justification such that no Gettiered subject is C-justified and Henry fails to satisfy the criterion for knowledge after all. It seems to me that the prospects for such a strategy are dim. To begin with, it is plausible that C-justification is an internalist kind of justification. For otherwise, it is hard to see how Sartwell will be able to ensure that the facts that determine C-justification are typically accessible to the subject in the way they need to be in order for C-justification to be a useful criterion for knowledge. Recall, however, that it was argued in the introduction that no internalist justification condition that retains its internalist identity can serve as the anti-Gettier condition. Given that all this is so, however, C-justification cannot be beefed up in such a way that no Gettiered subject is C-justified whilst, at the same time, retaining its internalist identity.

One might also think that Henry's claim to knowledge is pragmatically defective *despite* his response to the challenge. If so, so much the worse for Sartwell: What it indicates is that something is wrong with his criterion for knowledge.

given ticket in a fair lottery won't win. Does Sartwell's true belief account manage do justice to this intuition? In order to find out let us look at a more specific example. Suppose that Smith has bought a ticket in a state lottery in which six numbered balls are drawn from an urn with forty-nine numbered balls. The chance of his winning the jackpot is 1/13983816.[14] I believe that Smith will not win the jackpot. My belief is well integrated into my overall web of beliefs—and indeed my knowledge—about lotteries, probability calculus, Smith's financial situation after the drawing and so on. I am strongly committed to what I believe. I am prepared to bet a significant amount of money on it, to act on it if necessary etc. My belief that Smith will not win the jackpot clearly satisfies the additional criteria Sartwell imposes on belief. So, by Sartwell's lights, I believe that Smith won't win the jackpot. Add that Smith's ticket is indeed not the one that wins the jackpot. My belief is true and hence, by Sartwell's lights, it qualifies as knowledge. By the same token, Sartwell's first strategy of explaining away an intuition of ignorance fails.

How about Sartwell's second strategy of explaining away such an intuition? Again, the prospects for success here are very dim indeed. After all, the odds against winning provide excellent (C-) justification for my belief that Smith's ticket won't win the lottery. If I were to make a claim to knowledge and were to be challenged, I could adduce the relevant (C-) justification. Thus, by Sartwell's lights, my knowledge claim remains in good standing. We have a second type of case in which the true belief account predicts that the subject knows, while, at the same time, Sartwell's strategies of explaining away our intuition of ignorance fail.

Given that this is so, any appeal that Sartwell's negative argument may retain now rests on the cases in which, on the face of it, the subject knows whilst not P-justifiably believing what he knows. If these cases can be shown to be *ultima facie*

---

[14] That is, provided that the balls are not replaced once drawn.

unconvincing, then Sartwell's negative argument fails. Now, it seems to me that there is excellent reason to believe that Sartwell's cases are *ultima facie* unconvincing. To see why this is so, recall, first, that a subject will be P-justified if he satisfies any candidate epistemic condition for knowledge (that is, not just if he has C-justification). However, it is consistent with Sartwell's description of these cases that the subjects satisfy various epistemic conditions on knowledge. Thus, the mathematician may be *entitled* by rationalist intuition to his belief, the religious person's belief may be *caused* by God and the person who later finds out that his belief is justified may have formed his earlier belief via a *reliable method* (without his knowing it). In fact, it is consistent with Sartwell's description of these cases that the subjects satisfy some reliability condition on knowledge. At the same time, it is plausible that if the cases are re-described in such a way that it is clear that their beliefs do not satisfy any reliability condition, our intuition that the subjects know evaporates. For instance, if the mathematician believes the mathematical proposition because his horoscope predicted that he will come to believe a true mathematical proposition, intuitively, he does not know.[15] Hence, it is not only consistent with Sartwell's description of the cases but also plausible that the subjects satisfy some reliability condition. Since reliability conditions are candidate epistemic conditions for knowledge reliability conditions count as P-justification conditions. Thus it may be denied that Sartwell's cases are cases in which the subject knows without having P-justification.

In conclusion, then, there are cases—*viz.* Gettier and lottery cases—in which the true belief account predicts knowledge, while Sartwell fails to explain away an intuition of ignorance. Thus, there remain some *prima facie* counterexamples to the true belief account. At the same time, the cases that, according to Sartwell, provide *prima facie*

---

[15] This will be so even if the mathematician has a sufficiently tight web of further belief and commitment to the proposition to qualify as knowing. If Sartwell suggests my claim is false, for instance, because he claims that we have the intuition that the mathematician knows in these circumstances, it seems to me that he is simply mistaken.

evidence that one can know just in virtue of believing truly are, *ultima facie*, unconvincing. Even if they show that their protagonists know whilst lacking a certain type of justification, they do not provide reason to believe that one can know just in virtue of believing truly. Given that all this is so, there is excellent reason to believe that Sartwell's negative argument fails, that there is a pre-analytic commitment to a distinction between true belief and knowledge and that the true belief account is false.

## 1.2  An Outline and Critique of Sartwell's Positive Argument

Before advancing to the last conclusion, however, let us look at the positive part of Sartwell's argument. If Sartwell can succeed in establishing that knowledge cannot be anything but true belief, we will be facing a puzzle and might have to review the intuitions that suggested otherwise, however firm they may be. There are two arguments that can be reconstructed from Sartwell's discussion. Here they are:

*Argument 1*
  (1)   There is one and only one goal of inquiry with regard to particular propositions.
  (2)   True belief—that is, *merely* true belief—is the goal of inquiry with regard to particular propositions.[16]
  (3)   Knowledge is the goal of inquiry with regard to particular propositions.[17]
  (4)   Hence, knowledge is merely true belief.[18]

*Argument 2*
  (5)   There is one and only one goal of inquiry with regard to particular propositions.
  (6)   Justification—that is, again, P-justification—is merely instrumental to achieving the goal of inquiry with regard to particular propositions.[19]
  (7)   Knowledge is the goal of inquiry with regard to particular propositions.
  (8)   If something is merely instrumental to achieving a goal then it cannot be built into the description of the goal.[20]
  (9)   So, P-justification cannot be built into the description of knowledge.
  (10)  So, knowledge does not require P-justification.

---

[16] Cf. Sartwell (1992), pp. 172-3.
[17] Cf. Ibid., p. 175.
[18] Cf. Ibid., p. 173.
[19] Cf. Ibid., p. 174.
[20] Cf. Ibid.

Now should either of these arguments bother those who reject that knowledge is just true belief? The answer is 'no'. Let us turn to the first argument first. The obvious complaint that an advocate of the view that knowledge is not just true belief would file here is that (2) and (3) are not both true. Either the goal of inquiry with regard to particular propositions[21] is not merely true belief or it is not knowledge. Indeed, in view of the independent evidence against the true belief account provided the last subsection, he can regard the argument from (1) to (4) as a *reductio* on one of its premises. So, if Sartwell's first argument is to have any bite he had better provide good support for its premises. For present purposes, I will grant Sartwell his first premise which is, by the way, implicit in his discussion and for that reason does not receive explicit support. Let us look at what he has to say in support of the second premise—that true belief is the goal of inquiry with regard to particular propositions. First, Sartwell claims that it is universally agreed.[22] However, since he is aware that it is also universally agreed that knowledge is not merely true belief, he thinks that it stands in need of independent support. Here is how he ventures to get it:

(11) We cannot specify what it is to have P-justification without reference to the notion of truth.[23]

(12) "This indicates that the purpose [i.e. goal, CK] of inquiry can be formulated without reference to the notions of ground or reason [or P-justification, CK]."[24]

(13) Hence, everyone who accepts (11) is committed to truth's being the goal of inquiry.[25]

However, this argument is invalid. Certainly, if truth is the goal of inquiry, then (11) indicates what (12) takes it to indicate. That truth is the goal of inquiry can, on pain of circularity, not be assumed, however. On the other hand, barring this additional assumption, (12) begs the question against a defender of the view that knowledge is the goal of inquiry and that knowledge entails P-justification. Either way, then, the

---

[21] For stylistic reasons I may henceforth omit the qualification 'with regard to particular propositions'.
[22] Cf. Sartwell (1992), p. 172.
[23] Cf. Ibid., p. 175.
[24] Cf. Ibid., p. 176.
[25] A variation of this argument can already be found in Sartwell (1991), p. 161.

argument does not establish (2). What about the idea that it is widely agreed that (merely) true belief is the goal of inquiry? It is, indeed, rather widely agreed that the goal of inquiry is *acquiring* true beliefs and *avoiding* false ones. However, acquiring true beliefs and avoiding false ones is a *general* goal of inquiry. From this it does not follow that true belief is the goal of inquiry *with respect to particular propositions*. So, even the common agreement among philosophers to which Sartwell appeals does not lend support to (2). For that reason a defender of the view that knowledge is not merely true belief may regard Sartwell's argument from (1) to (4) in conjunction with last subsection's independent evidence against its conclusion as a *reductio* on (2).

What about premise (3)—that knowledge is the goal of inquiry with regard to particular propositions? Can Sartwell at least convince us that we must hold on to (3)? Again, the answer seems to be 'no'. Sartwell adduces two reasons in support of (3). First, he says that if knowledge is not the goal of inquiry, then it is not clear why knowledge is of any importance to us. In particular, if knowledge is not the goal of inquiry, then, claims Sartwell, 'knowledge' is merely "a technical term with a stipulated definition."[26] Second, he claims that if knowledge is not the goal of inquiry, then it will no longer "be central to epistemology, since it no longer represents our epistemic goal."[27]

Both reasons appear to be spurious: First, even if knowledge is not the goal of inquiry, it may be of importance to us in much the same way as insurance is of importance to us. Knowledge gives us stability in our beliefs that may be of importance to us in much the same way as the stability to our wellbeing that is provided by health insurance. At the same time, insurance is certainly instrumental to our overarching goals. Similarly, knowledge may be instrumental to the overarching goal of inquiry. Neither insurance nor knowledge thereby loses its importance. The same analogy also

---

[26] Sartwell (1992), p. 175.
[27] Ibid.

suggests that even if knowledge is not the goal of inquiry, 'knowledge' is not thereby reduced to a technical term with a merely stipulated definition: Just as 'insurance' is not reduced to a technical term just because insurance is instrumental to our overarching goals, 'knowledge' need not be so reduced either. The first consideration Sartwell adduces in support of (3) thus fails. The second one does not fare better. Since Sartwell acknowledges that justification (of any sort) is not our epistemic goal it seems that, by parity of reasoning, he is committed to the claim that justification (of any sort) is not central to epistemology—a claim I take to be obviously false. So, Sartwell also fails to provide satisfactory support for (3). Defenders of the view that knowledge is not merely true belief can thus view his argument in conjunction with last subsection's evidence against its conclusion as a *reductio* on (3) as well.

Let us then move on to the second argument. As a first observation I would like to point out that the argument will be successful only on the assumption that there is no further non-epistemic condition on knowledge.[28] Since Sartwell aims to establish that knowledge could not be anything but true belief, it not at all obvious that he is entitled to this assumption. However, it is far form clear that the argument will be successful even given this assumption. Let us ask how plausible it is for those who accept that knowledge entails P-justification to accept all of (5), (6), (7), and (8) at the same time. It would seem that it is not very plausible. In particular, it is not clear why someone who accepts that knowledge entails P-justification should also accept both (6)—that P-justification is merely instrumental to achieving the goal of inquiry with regard to particular propositions—and (7)—that knowledge is the goal of inquiry with regard to particular propositions. He can just take the independent evidence against the true belief account provided in the last subsection in conjunction with Sartwell's argument as a *reductio* on the claim that (6) and (7) are both true. So, everything depends on whether

---

[28] This was already pointed out earlier on in footnote 4, p. 23.

or not Sartwell manages to supply independent and compelling support for both (6) and (7). However, Sartwell does very little by way of defence of these statements. We have already looked into the considerations Sartwell adduces in support of the claim that knowledge is the goal of inquiry—here premise (7)—and found them unconvincing. So, someone who rejects the true belief account can comfortably take his argument from (5) to (10) in conjunction with last subsection's evidence against its conclusion as a *reductio* on (7).

What about premise (6)? Given that knowledge rather than true belief is the goal of inquiry Sartwell has done nothing to support (6). At best, Sartwell has argued that P-justification is instrumental to truth. That does establish (6), of course, if (2) is true, that is, if truth (true belief) is the goal of inquiry. If Sartwell uses (2) to establish (6), however, the argument from (5) to (10) then collapses into the argument from (1) to (4) and can be criticised as outlined above. On the other hand, barring the assumption of (2), (6) begs the question against the defender of the view that knowledge is the goal of inquiry and that knowledge entails P-justification. Since Sartwell fails to provide support for (6), Sartwell's argument in conjunction with last subsection's evidence against its conclusion can be regarded as a *reductio* on (6) as well.

In conclusion, neither of Sartwell's positive arguments for the true belief account is compelling. At the same time, the last subsection provided independent evidence against the true belief account. In view of these facts, I take it that Sartwell's two positive arguments do not establish the true belief account. Rather, they are *reductios* on at least one of their premises. Since the true belief account remains unsatisfactory, there is reason for us to reject it: Knowing *P* is not equivalent with truly believing *P*.

## 2  Foley/Pritchard's Epistemic Deflationism: Knowledge is True Belief Embedded in Further True Beliefs

It is premature, however, to conclude that because Sartwell's true belief analysis fails, we may suppose that knowledge cannot be analysed in terms of true belief only. For there is another proposal according to which knowledge is analysed in terms of true belief only, while, at the same time, knowing *P* is not just truly believing *P*. According to this proposal, knowing *P* consists in truly believing *P* and having further true beliefs about propositions that are in the neighbourhood of *P*. This proposal has first been advanced by Foley (1996), (2002) and has then been developed further by Pritchard (2004*a*). Since Foley and Pritchard's conception of knowledge clearly satisfies the criteria for deflationary approaches to epistemology outlined above, I suggest testing whether it fares better than Sartwell's true belief account.

### 2.1  Foley/Pritchard's Epistemic Deflationism: The Credentials

Foley gives different specifications as to how the view is to be understood in detail. In the early paper he offers the following biconditional:

> One knows that P if and only if one has accurate enough beliefs about P and also comprehensive enough beliefs about P.[29]

In the later talk he is more precise on what 'accurate and comprehensive enough' amounts to:

> To say that S knows that P is to say that S believes that P, P is true, and there is no significant, relevant lacuna in S's information about P or closely related matters.[30]

---

[29] Foley (1996), p. 93.
[30] Foley (2002), p. 4.

Since the second statement of the view can be understood as saying in more detail what the idea of accurate and comprehensive enough true belief in the first statement amounts to, the second statement can be read as a more refined version of the first one. Accordingly, I will use both statements of the view as my purposes require.[31]

Foley does not give precise criteria for what makes for a significant lacuna. However, he makes two claims that are notable in this regard. First, he points out that how accurate and comprehensive one's beliefs about $P$ must be in order to be accurate and comprehensive enough varies with context.[32] In support of this claim, Foley cites examples that suggest that there are broader contexts in which one must have a rather wide-ranging set of true beliefs about $P$ and related matters in order to for the sentence 'one knows that $p$' to turn out true. On the other hand, there are narrower contexts. According to Foley, in some contexts, instances of mere true belief that $p$ can be enough to make a sentence of the form 'S knows that $p$' true. The second claim Foley makes about significant lacunae that is notable consists in some rough guidelines as to which truths are likely to constitute a significant lacuna in the subject's information about the situation. Included are:

> Truths that are conceptually linked with [the proposition believed, CK] … truths that have played (or would have played if S had believed them) a role in shaping what S believes … [and, CK] truths that are likely to have important behavioural consequences.[33]

---

[31] Of course, the second statement of the view can be understood as a refinement of the first one only if a lacuna in one's information about $P$ or closely related matters consists in—or at least is necessarily accompanied by—a lack of belief about some proposition in the neighbourhood of $P$. Since in the 2002 talk Foley himself shifts between the two statements (cf. for instance Foley (2002), p.3 for a statement akin to the first one and p.4 for the refined statement), it is safe to assume that, according to Foley, lacunae in one's information about $P$ or related matters do consist in—are necessarily accompanied by— such a lack of belief. Accordingly, I will also allow myself to say things like "there is no significant, relevant lacuna in S's *web of beliefs* about $P$."

[32] Those who are familiar with recent debate about contextualism in the theory of knowledge will be inclined to inquire exactly how Foley envisages implementing his contextualist thesis. Foley himself does not address the issue. In his development of the view, Pritchard ((2004*a*), pp. 13-16) suggests implementing the contextualist thesis along the lines of Michael Williams's subject contextualism (cf. Williams (1996)). Since the problems I will raise for the view stand regardless of how exactly the contextualist thesis is implemented, however, there is no need to settle the issue within the scope of this thesis.

[33] Foley (2002), p. 10.

Foley gives two arguments in favour of his conception of knowledge. First, he claims that it succeeds in explaining why subjects in Gettier cases lack knowledge. Second, he adduces a thought-experiment suggesting that nothing other than true belief embedded in further true beliefs could be constitutive of knowledge.

## 2.2   Foley/Pritchard's Deflationism and Gettier Cases

Let's start with Gettier cases. Recall that Gettier cases are cases in which a subject has a justified true belief, for instance, but does not know. We have also seen that since Gettier's paper, epistemologists have been trying to answer the question as to what needs to be added to true belief in order to solve the Gettier problem. Foley's proposal in this regard is, of course: further true beliefs that are both accurate enough and comprehensive enough. Or, in other words, there must be no significant lacuna in the subject's information about the proposition truly believed or related matters. In order to solve the problem posed by Gettier cases, then, Foley has to show that a Gettiered subject's belief is never accurate and comprehensive enough, that there is always a significant lacuna in her information about the proposition believed or about related matters.

How does he pull this off? To begin with, Foley views Gettier cases as little epistemic stories told by a narrator. He claims that an invariable and crucial feature of these stories is that the narrator tells them in such a way that there is an aspect, a truth about the situation that the Gettiered subject is unaware of. At the same time, the story is told in such a way as to convince the audience that the aspect of which the Gettiered subject is unaware constitutes a significant lacuna in his information about the situation. In this way, the narrator convinces the audience to deny that the Gettiered subject knows. In the case of Henry in Barn Façade County, for instance, the narrator of the

story, I, tell the story in such a way that Henry is unaware that the field is full of barn façades. At the same time, I convince my audience, you, that this aspect is a significant aspect of the situation. By the same token I convince you that Henry is unaware of a significant aspect of the situation which in turn leads you to deny that Henry knows.[34]

## 2.3   Sally: A Subject with Maximally Accurate and Comprehensive Beliefs

Foley's case does not rest solely on his promised solution to Gettier cases. Rather, he adduces an argument that appears to establish that nothing other than true belief that *p* backed up by further true beliefs about the neighbourhood of *P* could be constitutive of knowledge that *p*. Foley asks us to consider the limiting case of his theory:

> Imagine that Sally's beliefs are as accurate and comprehensive as is humanly possible. She has true beliefs about the basic laws of the universe, and in terms of these basic

---

[34] Cf. Foley (2002), p. 4. An alternative—and, it seems to me, better since more clearly *semantic*—explanation of the Gettiered subject's ignorance proceeds along the following lines: First we combine Foley's definition of knowledge and his contextualist thesis in the following way:

(FC)    'S knows that *p*' is true in context c iff S believes a set of truths, $T_1 \ldots T_n$, relevant in c.

With FC in play we reconstruct Foley's argument in the following way: Since the narrator a) *introduces* a truth, $T_i$, that is b) *significant* he thereby creates a context in which $T_i$ is relevant. Moreover, the narrator sets up the story in such a way that the Gettiered subject, *g*, does not believe $T_i$. Hence, the narrator creates a context in which *g* does not believe all relevant truths. Hence, he creates a context in which the sentence '*g* knows that *p*' is false.

One might object at this point that the intended sense of 'significant' at issue in the above reconstruction of Foley's argument is equivalent to the intended sense of 'relevant' at issue in FC. Notice, however, that Foley has (at least the beginning of) a story of what sorts of truths are significant (cf. p. 36). Foley may then be understood as claiming that the fact that the narrator a) *introduces* a truth, $T_i$, that b) *satisfies some of his criteria of significance* creates a context in which $T_i$ is relevant. The explanation of the Gettiered subject's ignorance can then proceed as rehearsed.

Finally, one might object that what the present explanation of the subject's ignorance suggests is that we do not really have to appeal to the figure of the narrator. After all, it would seem plausible that—at least in some cases—a truth can be relevant without being introduced by a narrator. If so, it is no longer clear why there should not be some Gettier cases in which there is no narrator introducing a truth of which the Gettiered subject is unaware. If there are such cases, however, Foley's explanation of why subjects in Gettier cases lack knowledge does not hold with full generality. In response to this objection, Foley may venture to argue that the truths that Gettiered subjects fail to believe, which explain their ignorance, must always be introduced by a narrator since, otherwise, we will not manage to construct a *Gettier* case: Gettier cases are in a salient sense *abnormal*. Since there is a presumption of normality, the abnormality must be introduced to get the case off the ground. Alternatively, however, he can also drop the idea of the narrator and explain Gettiered subjects' ignorance simply in terms of FC and a list of criteria of relevance.

laws, she can explain everything that has happened, is happening, and will happen. She can explain the origin of the universe, how life came into existence on earth, the mechanisms by which cells age, when the sun will die, and so on. She can even explain how it is that she came to have all this information. She has as close to a perfectly accurate and comprehensive belief system as a human can have.[35]

Foley then asks us to consider a proposition about the mechanism by which cells age and submits that if anybody knows that proposition, Sally does. It seems to me, however, that this restriction is unnecessary. Foley might as well maintain that since Sally has maximally accurate and comprehensive true beliefs, that she has as tight a web of true beliefs as is humanly possible, for virtually any proposition, *P*, Sally knows that *p* if anybody does. Foley then asks us to consider theories of knowledge according to which knowing is not just a matter of having sufficiently accurate and comprehensive true beliefs but according to which knowing depends on having a true belief that also has a further feature *F*. Foley discusses externalist proposals for *F* such as belief formed by a reliable method, belief formed via a properly functioning cognitive apparatus, belief caused by the facts in an appropriate way and belief with certain modal features on the one hand, and internalist proposals such as non-defectively justified belief and indefeasibly justified belief on the other. As regards externalist proposals, Foley points out that Sally might lack all of the proposed candidates for *F*. Sally might have acquired her beliefs in freakish ways such that her belief forming methods are unreliable, her beliefs are not caused by the facts, do not derive from the exercise of properly functioning cognitive faculties and are modally instable. Since, intuitively, Sally has knowledge despite not having either of these candidate features, none of the candidate features is necessary for knowledge.[36] As regards internalist proposals, Foley points out that, while Sally is most likely to have them, the fact that she has them is a consequence of the fact that she has maximally accurate and comprehensive true beliefs. Sally's

---

[35] Foley (2002), p. 13.
[36] Cf. Ibid., pp. 13-15.

knowledge is constituted by her having maximally accurate and comprehensive true beliefs. As a by-product, she may have various kinds of internalist justification.[37]

It is implicit in Foley's discussion of the various internalist and externalist proposals concerning the additional feature *F* that is required for knowledge besides true belief that his argument generalises: For any proposed additional feature *F* either Sally need not have it or it is merely a consequence of the fact that she has maximally accurate and comprehensive true beliefs. In either case it is not a constituent of knowledge. Foley's thought experiment appears to show that knowledge could not be anything but true beliefs plus further true beliefs. It completes the case for his theory.

## 2.4   A Critique of Foley/Pritchard's Deflationism

Having outlined the credentials of and the arguments in support of Foley/Pritchard's conception of knowledge, I will now argue that it is not viable. To be more precise, there are two problems that Foley/Pritchard's conception of knowledge encounters. First, the thought-experiment revolving around Sally on which the second half of the case for the position rests does not provide the support needed. Second, Foley/Pritchard's conception of knowledge fails to give adequate predictions in some lottery and Gettier cases so that the first half of the case for it is undermined, too.

To see why Foley's thought-experiment does not support his theory consider, first, the following variation of Foley's Sally case:

> Kelly is a compulsive gambler. One day she decides to play a game: She considers propositions and flips a coin to 'determine' whether they are true or false. She will do this until she has arrived at a contradiction. As it so happens Kelly's coin gives her consistent results for a long period of time. Kelly is somewhat mystified by this fact and suspects that her coin might tell her the truth. So she decides to 'test' the coin: She flips it to 'determine' whether the coin has so far 'told' her only the truth. The coin's 'answer' is positive. Kelly gets more excited and decides to play on. After a long time

---

[37] Cf. Foley (2002), p. 16.

she has still not arrived at a contradiction. Then she does the 'test' again. Again it comes out positive. Kelly is now at the verge of believing everything the coin 'tells' her. After another while without contradiction, Kelly has succumbed to the coin: She believes everything it has so far 'told' her. She frequently applies the 'test' and is always confirmed in her beliefs. Now she really wants to know it. She asks the coin whether it will continue to 'tell' her the truth. The coin's 'answer' is again positive. Kelly plays on and arrives at as complete a system of beliefs as is humanly possible. Furthermore, coincidence has it that the coin has always 'told' her the truth. Thus, Kelly has arrived at a maximally accurate and comprehensive system of beliefs (including beliefs about how she has arrived at her beliefs, *viz.* by flipping a completely normal coin, and about the epistemic status of her beliefs such as that they are a result of an incredible coincidence).[38]

Does Kelly's maximally accurate and comprehensive system of beliefs qualify as knowledge? It seems to me that the answer is clearly 'no'. The way Kelly arrived at her beliefs resembles a series of right guesses. However, a series of right guesses does not make knowledge. Guesses are too lucky to qualify as knowledge. Similarly, the beliefs she arrives at via her game are too lucky to qualify as knowledge.

What my version of the example suggests is, of course, that someone's beliefs have to have some other feature or features besides being maximally accurate and comprehensive in order to qualify as knowledge. How come, then, that upon being presented with Foley's version of the case some of us may have had the intuition that they don't? It is not easy to make out exactly what drove such intuitions in Foley's version of the case. However, it seems to me that part of it may be that Foley is largely silent about how exactly Sally came by her beliefs. For all Foley has said, Sally may have built up her belief system step by step, testing her beliefs as she goes along. If so, then it is quite plausible that her beliefs qualify as knowledge. However, it is rather implausible that her beliefs are *merely* true beliefs. Or, alternatively, she may have acquired her system of beliefs all at once whilst having had it long enough to acquire

---

[38] Note that we may suppose that Kelly does not at any point have a false belief that her coin has special powers. The idea is that she gets sucked into believing what the coin 'tells' her. She also believes that the coin only ever gives her the truth. At the same time, we may suppose that she remains agnostic about whether it is a special coin. She is so focused on the game that she does not even consider these things. We may also suppose that she forms the beliefs about the nature of the coin and the epistemic status of her beliefs at the end of her game and is 'told' that it is a normal coin and that everything that has happened is a huge coincidence.

good inductive evidence that all of her beliefs are true. If so, then perhaps she knows that all of her beliefs are true in which case, it is plausible to assume, all of her beliefs also qualify as knowledge. Again, however, it is rather implausible that Sally's beliefs are *merely* true beliefs. After all, the crucial belief—that all of her beliefs are true—is supported by good inductive evidence. On the other hand, if, for instance, Foley had described the case in such a way that Sally wakes up from a coma and finds herself with a very tight web of beliefs about nearly all aspects of the world, including, for instance, that she has acquired those beliefs by being, simultaneously, struck by lightning and exposed to radiation as well as electric shocks, it is much less clear that we would judge, intuitively, that Sally's beliefs qualify as knowledge even if, in fact, all of her beliefs are true. One indication of the fact that we would not judge in this way is that if Sally were to assert any of the things she believes and were to defend her assertion against a challenge how *she* came by such knowledge by pointing out that she has a maximally accurate and comprehensive system of beliefs which she acquired by being struck by lightning etc., we would—at least initially, that is, until they have been tested—treat her as out of her wits rather than take her seriously.

The second and presumably greatest problem that arises for Foley's proposal is posed by some lottery and Gettier cases. Recall that, intuitively, we do not know, in advance, that given tickets in fair lotteries won't win. Suppose, Jones truly believes that ticket 3 won't win fair lottery L. By Foley's lights, that means that there must be some proposition in the neighbourhood of the proposition believed by Jones that he fails to truly believe. For if there is no such proposition, then it follows from Foley's account that Jones knows that ticket 3 won't win L—contrary to our intuitions. Now, what is the proposition that Jones fails to truly believe? The obvious answer seems to be that Jones fails to believe of the winning ticket that it will win. That is the lacuna in his web of beliefs that prevents him from knowing what he truly believes—*viz.*, that ticket 3 will

lose. On second thought, however, it is far from clear that this must be so, that Jones must fail to believe of the winning ticket that it will win. And it is not clear that if he doesn't, there must be some other lacuna in his web of beliefs about the situation. For, suppose Jones entertains beliefs about the nature of lotteries in general and about L in particular: Jones believes there are different kinds of lottery, that the chance of winning depends on the number of tickets, that most lotteries are fair etc.; that in L there are 10000 tickets, let us say, that there is one and only one winner, that the chance of winning is 1 in 10000, that all the tickets have been sold, how much the tickets cost, when L will be held etc. Furthermore, suppose he also has exacting beliefs about the outcome of L: He believes of exactly one ticket—let us say the one he owns—that it will win and of all the other tickets that they will lose. In fact, we may assume that Jones has as tight a web of beliefs about the lottery situation as is humanly possible. Given that this is so, there is no lacuna in his web of beliefs about the lottery situation. Now suppose Jones's ticket is the winner and that none of Jones's beliefs about lotteries in general and about L are false. If so, Jones's beliefs about the lottery situation are true. Since there is no lacuna in Jones's beliefs about the lottery situation and since all of Jones's beliefs about it are true, we may expect Foley's theory to predict that Jones's beliefs qualify as knowledge: Jones knows, in advance of the drawing, which ticket will win and which tickets will lose—contrary to our intuitions.

Before moving on, I would like to point out that the same line of thought can be used to show that Foley will not be able to provide a satisfactory analysis of all Gettier cases. Recall, for instance, the case of Henry in Barn Façade County: Henry sees the only real barn in a field full of barn façades and comes to believe that he is facing a barn. As a matter of fact, his belief is true. But it does not qualify as knowledge. According to Foley/Pritchard's deflationary conception of knowledge, the reason for this is that Henry lacks a belief about an important fact about the situation—*viz.*, that the

field is otherwise full of barn façades. But now suppose that Henry, somewhat irrationally, entertains true beliefs about the fact that the field is otherwise full of barn façades and that he is facing the only real barn around.[39] The lacuna in Henry's web of beliefs is closed. However, our intuition that Henry's belief that he is facing a barn does not qualify as knowledge lingers. Foley/Pritchard's conception of knowledge does not succeed in providing a fully satisfactory analysis of Gettier cases either.

One might try to rescue Foley/Pritchard's conception of knowledge from these objections by pointing out that Jones's belief that his ticket will win the lottery and Henry's belief that he is facing the only real barn in a field that is otherwise full of barn façades are *irrational*. However, the mere fact that Jones and Henry's beliefs are irrational does not, by Foley/Pritchard's lights, prevent them from being knowledge. After all, all that is needed for knowledge is that the belief be true and embedded in a sufficiently tight web of further true beliefs about the situation. Of course, Foley and Pritchard could place an additional non-irrationality condition on knowledge. In that case, however, they have moved away form the deflationist proposal that there is no epistemic property—such as justification, reliability, rationality and the like—that is constitutive of knowledge.

So what Foley and Pritchard (*qua* deflationists) would have to say in response to the problem posed by these cases is that the irrationality of Jones and Henry's beliefs gives rise to an important lacuna in their respective webs of belief. And indeed one may be tempted to say that since (some of) Jones and Henry's beliefs are irrational Jones and Henry will certainly not believe (of those beliefs) that they are irrational. And they will also not believe that it is irrational for them to believe the propositions they irrationally believe. If they don't have these beliefs then there will be a lacuna in Jones and Henry's

---

[39] Perhaps someone told him that with one exception there are only barn façades in the field. However, Henry is convinced that the barn he is now facing is the only one which is not a façade.

webs of belief about the situation that Foley and Pritchard might exploit to explain why they do not know after all.

I would like to make two comments about this way of rescuing Foley/Pritchard's conception of knowledge from the last objections: First, the beliefs that Henry and Jones lack and that a defender of Foley and Pritchard wants to exploit in order to rescue their theory from these objections are beliefs about the epistemic status of some of Henry and Jones's beliefs (whether they are irrational) and some propositions (whether it is irrational to believe them) respectively. It is a lot less clear, however, that mere lack of those beliefs should make for a lacuna that is important to whether or not Jones and Henry's beliefs qualify as knowledge. Consider children and animal knowledge. Children and animals may be incapable of having beliefs about the epistemic status of beliefs or propositions—for instance if they lack the relevant concepts of epistemic status. Yet, we are often willing to attribute knowledge to them. Here lack of such beliefs is often irrelevant to attributions of knowledge. Given that this is so, it is far from clear that Jones and Henry's lack of belief about the epistemic status of some beliefs or propositions should constitute a lacuna in their belief system that is relevant to whether or not their beliefs qualify as knowledge.

My second comment on the present attempt to rescue Foley/Pritchard's proposal is that it seems to me that at least in some cases the subjects may even truly believe that the belief at issue is irrational/that it is irrational for them to believe the proposition at issue. For instance, I may point out to Jones that given the odds it is irrational for him to believe that his ticket will win and that in consequence he ought not believe that he will win. It seems to me that Jones may believe everything I have pointed out to him whilst, at the same time, holding onto his belief that his ticket will win. (Maybe he paid so much for the ticket that giving up this belief would cause him to have severe depression so that he hangs on to his belief as a matter of self-preservation.) If so, then the defence

of Foley/Pritchard's conception of knowledge fails even more dramatically in that the alleged lacuna can once again be closed without altering our intuitions.

In summary, then, we have looked at two conceptions of knowledge that can naturally be understood as deflationary: Sartwell's proposal according to which knowing $P$ just is truly believing $P$ on the one hand and Foley/Pritchard's proposal according to which knowing $P$ is having a true belief that $p$ that is embedded in a sufficiently tight web of true beliefs in propositions about $P$ and related matters on the other. However, we found both theories unsatisfactory. Strikingly, both theories ran into problems with lottery cases. There were at least some cases in which both theories predicted—contrary to our intuitions—that subjects know, in advance, that given tickets in a fair lottery won't win.[40] This is no accident. What it points to is that true belief (even if embedded in further true beliefs) may be too lucky to qualify as knowledge. Given that this is so, we need a further condition that gives true belief a stability that rules out this kind of luck. It is highly plausible that a condition that gives true belief the right kind of stability needed for knowledge will be epistemic in nature.

---

[40] In fact, both theories even predicted, at least in some cases, that subjects know, in advance, that given tickets in fair lotteries *will win*.

# III  Epistemic Minimalism

We have reason to believe that deflationary approaches to the theory of knowledge do not work. Since we characterised such approaches as approaches according to which there is nothing epistemic that explains why 'knows that *p*' applies to all subjects to which it does apply, we have reason to believe that there is something epistemic that explains why 'knows that *p*' applies to all subjects to which it applies. Of course, we may now plunge right into speculations as to what the epistemic property or properties in question might be. In doing so, however, we would make ourselves vulnerable to advance just another proposal in the post-Gettier literature on the analysis of knowledge. For that reason I propose to opt for a different strategy. In parallel to Wright's minimalist approach to the theory of truth, we may approach the theory of knowledge in a minimalist way: we may attempt to pin down the concept of knowledge by a set of platitudes about knowledge. This approach will have similar advantages as Wright's approach to the theory of truth. First, since we will have arrived at any resulting commitments about the nature of knowledge via a set of platitudes, these commitments will be comparatively lightweight and well-motivated. Moreover, just as Wright can hope to avoid some of the metaphysical problems that traditional inflationary theories of truth encounter we may hope to avoid some of the problems that have beset the various proposals in the post-Gettier literature that attempted to spell out a set of individually necessary and jointly sufficient conditions for knowledge.

In his minimalist approach to the theory of truth, Wright proposes to pin down the concept of truth by a set of platitudes about truth. Similarly, in the minimalist approach to the theory of knowledge, I suggest pinning down the concept of knowledge by a set of platitudes about knowledge. It seems to me that the following are platitudes about knowledge:

The Factivity Platitude: One cannot know falsehoods.

The Belief Platitude: One can know only what one also believes.

The Good Informant Platitude: Someone who knows *P* is typically—that is, provided that he is willing to part the information, has not lost his credibility etc.—a good informant concerning *P*.

The Telling/Informing Platitude: If one tells/informs someone that *p* then one represents oneself as knowing that *p*.

The Closure Platitude: One knows the conclusions of one's competently executed deductions provided that one knows the premises and comes to believe the conclusions on the basis of one's deductions.

The Anti-Luck Platitude: Knowledge excludes luck.

Following Wright, we might be inclined to say that these platitudes pin down the concept of knowledge such that any predicate that satisfies the platitudes is what we may call a knowledge predicate. Knowledge could then be understood as the property signified by the knowledge predicate. We must, however, be careful in what we may expect the above platitudes to do for us. In particular, we may expect them to fix the concept of knowledge in the relevant sense only if they constitute a *complete* list of platitudes about knowledge (or, at least, a list of platitudes from which all other platitudes are derivable). For if there is some platitude missing, there is no guarantee that any predicate that satisfies the platitudes will be a knowledge predicate. After all, a predicate may satisfy all the abovementioned platitudes but fail to satisfy the missing platitude in which case it will not be a knowledge predicate. At the same time, I do not want to lay claim to having provided a complete list of platitudes about knowledge. Accordingly, it will—for now at least[1]—be wise to aim for something weaker than the perfect parallel to what Wright envisages for the concept of truth. Rather than expecting the above platitudes to completely fix the concept of knowledge, we may expect them to *partly* fix it. If so, then, even if it need not be the case that every predicate that satisfies the platitudes will be a knowledge predicate and signify the property of knowledge, it

---

[1] I will sketch how the minimalist approach can be used to arrive at a knowledge predicate, which may then be thought of as signifying the property of knowledge, in the conclusion of this thesis.

will still be the case that no predicate will be a knowledge predicate and no condition will qualify as a condition on knowledge unless it satisfies these platitudes.

Let me also be clear that I do not take a platitude to be something that most folk believe. Folk, I take it, often do not have beliefs at this level of abstraction. However, I do think that these platitudes are implicit in the way in which we, including folk, deploy the word 'knows' in our ordinary language and thought. Given that this is so, it should be possible to point to intuitions and facts about usage that provide the relevant support. This is what I will do in the following sections. Apart from supporting the platitudes by pointing to intuitions and facts about usage that point to a commitment to the platitudes I will, where possible, provide further arguments in support of them. Let me begin with the factivity, belief and good informant platitudes.

## 1  The Factivity, Belief and Good Informant Platitudes

### 1.1  Support from Intuitions and Facts about Usage

There are certain facts about our use of the word 'know' in ordinary language and thought that point to a commitment to the factivity, belief and good informant platitudes. Let us first look at the factivity platitude—that is, that one cannot know falsehoods. To begin with, suppose that someone believes, completely unjustifiably and irrationally, that two plus two is five. He has been told many times that two plus two is four not five but he sticks to his belief. Intuitively, his false belief does not qualify as knowledge. However, it is not just unjustified and irrational false beliefs that, intuitively, do not count as knowledge. For suppose John believes that there are no black swans. Intuitively, John's belief does not qualify as knowledge—no matter how good his inductive evidence. Or suppose Mary believes on the basis of otherwise highly reliable testimony that Amsterdam is the capital of the Netherlands. Again, intuitively,

Mary's belief, no matter how well justified, does not qualify as knowledge. At the same time, many *true* beliefs that Mary acquires by equally reliably testimony will, intuitively, qualify as knowledge. Intuitions like these point to a commitment to the factivity platitude.

Perhaps the best evidence for the belief platitude—one only knows what one also believes—is that, intuitively, people who suspend belief about a given question are not knowers. For instance, suppose Green suspends belief about whether Hynes committed the murder. Intuitively, Green does not know whether Hynes committed the murder. Moreover, this will be so, it seems to me, even if Green has evidence that is strong enough to give him knowledge that Hynes did the deed. So long as Green does not commit himself to what the evidence indicates, so long as he does not form a belief on the basis of it, he does not count as knowing. At the same time, if Green forms a true belief that it was Hynes on the basis of the evidence he has, then he will also count as knowing that it was Hynes. Further support for the belief platitude is provided by the fact that we can undermine someone's testimony that S knew at *t* that *P* by providing evidence that at *t* S did not even believe that *p*. For instance, if Tom testifies that Jack knew that Dave was at the scene of the crime at the time of the deed, we can undermine Tom's testimony by producing evidence that Jack did not even believe that Dave was at the scene of the crime at the time of the deed—for instance by producing phone bills that show that Jack tried to phone Dave in various other locations at the time of the deed.[2] The intuition that those who suspend belief do not know but, provided that they also have the right evidence, will know as soon as they form a true belief on the basis of their evidence on the one hand and the fact that we can undermine testimony that S knows at *t* that *p* by providing evidence that at *t* S does not believe that *p* on the other point to a commitment to the belief platitude.

---

[2] This is a slight variation of a case due to Pritchard (2004*a*), p. 106.

Let us then turn to the good informant platitude according to which someone who knows that *p* typically is a good informant concerning *P*. Most obviously, when someone inquires into *P* and you know that *p*, then you will feel competent to answer that *p* is true. For instance, if someone asks you when the concert is on tonight and you have just read in the newspaper that it is on at eight you will feel competent to provide the corresponding answer. Moreover, if you don't know whether *p* but you know that someone else knows—or can easily come to know or at least is likely to know—whether *p*, you will naturally recommend that other person to the inquirer as an informant. For instance, if you don't know when the concert is on but you know that one of your common friends bought a ticket, you will naturally recommend your common friend as an informant to the inquirer. Both facts about ordinary usage point to a commitment to the good informant platitude.

## 1.2   A Craig-Style Argument for the Factivity, Belief and Good Informant Platitudes

There is a more elaborate argument that provides support for all three platitudes which is due to Edward Craig (1990). In a nutshell, Craig argues that the need of communities of language users to flag good informants pressures these communities to introduce the concept of knowledge. It is the very job of the concept of knowledge to flag good informants. If Craig's argument is sound, then we may expect there to be a connection between knowledge and good informants of the kind stated by the good informant platitude. Knowers must typically be good informants because flagging good informants is precisely the job that the concept of knowledge does for us. While it is not hard to see how Craig's argument, if successful, supports the good informant platitude, the case is not so straightforward with respect to the factivity and belief platitudes. In order to

understand how Craig's argument supports the latter platitudes it will be necessary to look at (at least some of) the details.[3]

Since Craig's methodology is somewhat unconventional a few remarks by way of clarification appear in order. Craig proposes to start with a *prima facie* plausible hypothesis about the concept of knowledge concerning the job it does for us and then attempts to develop the conditions that would govern the application of a concept that does this job for us. The idea then is that if the concept developed can retain the main features of the concept of knowledge as ordinarily understood—such as, for instance, approximate its intuitive extension and intension—then there is evidence that the hypothesis is true, that the concept of knowledge indeed does the job the hypothesis says it does.[4]

Now one might be worried that even if Craig manages to show that the concept of knowledge does a certain job *for us*, there is not much by way of conceptual insight to be gleaned here—after all, the concept may do a different job for other people. In order to avoid this objection, Craig chooses as the setting for his argument a so-called *state-of-nature* scenario. In the present case, that means that Craig envisages the situation of linguistic communities that do not yet have the concept of knowledge and asks what sorts of needs could force them to introduce this concept. The hypothesis is intended as a hypothesis about precisely this situation. It is intended to capture the job a concept would do that is introduced in response to those needs. In this way since the hypothesis concerns the needs that pressured linguistic communities with a concept of knowledge to introduce this concept in general, the insights that can be gleaned from Craig's argument will be equally general—they will be true of any linguistic

---

[3] The ensuing outline of Craig's argument may appear to be unjustifiably detailed. However, some of Craig's considerations will be of considerable importance later on in the thesis. For that reason I ask the reader to bear with me for now and to allow me to spend a bit more time on developing Craig's argument than would otherwise have been necessary.
[4] Cf. Craig (1990), p. 2.

community with a concept of knowledge rather than just of our linguistic community. Accordingly, given that they are true of any linguistic community with a concept of knowledge, we may also expect them to be conceptual.[5]

With these remarks about methodology and scope in play let us now move on to the actual argument. As I have already indicated, the hypothesis that Craig devises is that the job of the concept of knowledge is to flag good informants. The starting point of Craig's argument in support of this hypothesis is a state of nature scenario—a scenario featuring arbitrary linguistic community before the introduction of the concept of knowledge. Craig highlights three features of the state of nature scenario. First, the members of the linguistic community need true beliefs about their environment in order to succeed in their actions. Second, it will be highly advantageous to sometimes obtain such true beliefs from other members of the community—that is to say, from *informants*. Third, since that is so, such communities will need concepts to evaluate informants. The hypothesis can be understood as characterising the concept of knowledge as a concept that serves one such evaluative function, namely evaluation as a good informant.[6]

In support of the hypothesis Craig goes on to develop the application conditions of the concept of good informant and argues that the application conditions of that concept approximate the application conditions of the concept of knowledge as ordinarily understood. How does he pull this off? To begin with, Craig envisages an inquirer who wants to find out whether $p$ is true but so far does not have a belief as to whether $p$. A good informant with respect to $P$ will be an agent the inquirer will want to acquire a belief that $p$ (or that not-$p$) from. The idea here is that we can get at the application conditions of the concept of good informant if we ask ourselves what general properties an inquirer will want an agent from whom he wants to acquire a

---

[5] Cf. Ibid., pp. 9-10.
[6] Cf. Ibid., p. 11.

belief that *p* (or that not-*p*) to have. Craig starts out with two general properties that the inquirer will want his informant to have. The inquirer will want his informant to be such that

    (a) either *P* and he believes that *p* or not-*P* and he believes that not-*p*[7], and
    (b) if he tells me that *p*, I shall thereupon believe that *p*.[8]

These two general properties of a good informant about *P* point to the idea that a good informant about *P* is a person who also has a true belief about *P*. Already the property stated in (a) indicates this. The property stated in (b) reinforces this indication. For it seems that if the prospective informant does not himself believe that *p* then he will often fail to tell the inquirer that *p* is true in a way that exhibits sufficient conviction to make him, the inquirer, believe that *p* upon being told.[9]

However, true belief is not all that the inquirer who wants to acquire a belief about *P* will want his prospective informant to have. After all, someone may have come by his true belief by flipping a coin and believing *P* if it comes out heads and not-*P* if it comes out tails. And such a person would not be the kind of person whom the inquirer would want to acquire a belief from. Rather, Craig points out that the inquirer will want his prospective informant to be such that:

---

[7] Ibid., p. 12.
[8] Cf. Ibid., p. 13.
[9] Cf. Ibid. Notice, however, that Craig thinks that there is reason to waver over whether belief is actually a necessary condition for being a good informant and hence whether it is a necessary condition for knowledge. He considers apparent counterexamples to the belief condition provided by Radford (1966) and maintains that, on the one hand, there are certain aspects that the subjects in those examples have in common with prototypical good informants/knowers, while, on the other hand, there are certain clear differences between the subjects and prototypical good informants/knowers. While the similarities between these subjects and the prototypical good informants/knowers enjoin us to regard these subjects as good informants/knowers too, the differences drive us in the opposite direction. In consequence we waver over what to make of these examples. Craig welcomes this result. He takes it to indicate that the traditional approach to the theory of knowledge, which attempts to analyse knowledge in terms of necessary and sufficient conditions, is bound to fail. At the same time, Craig also takes it to show that his approach has an edge over traditional analyses since it is not committed to the format of necessary and sufficient conditions and can instead explain—along the lines just outlined—why our intuitions waver (Cf. Ibid. pp. 13-14). Again, I am not at all convinced by Radford's alleged counterexamples to the belief condition on knowledge (cf. p. 20, n. 1). Accordingly, I take it that there is no reason for us to be in doubt about whether knowledge really requires belief. Moreover, it is noteworthy that Craig can allow this much without giving up his thesis that it is best not to conduct an inquiry into the concept or the nature of knowledge by trying to identify necessary and sufficient conditions for knowledge.

(c) He has a "detectable property … which correlates well with being right about p".[10]

Let's look at the part of 'correlating well with being right about *P*'. Craig argues that this part of the property has a modal feature. There are many possibilities that might, for all the inquirer knows, be actual possibilities—possibilities that the inquirer does not know not to obtain. Craig labels these possibilities 'open possibilities'. He maintains that the inquirer will want his prospective informant to be right across the range of nearby and open possibilities. After all, for all the inquirer knows, each of them is actual. If he wants his resulting belief to be better than a guess, the informant had better be right in all worlds that are nearby and open possibilities. Accordingly, the property mentioned in (c) that is both detectable and correlates well with being right about *P* can be understood as having a modal feature. It involves being right across a range of possible worlds.[11] As regards the precise nature of the property mentioned in (c)—or in other words, as regards the answer to the question: What is X? in 'S has a detectable property that correlates well with being right about *p* iff X'—Craig claims that there is no answer. All that we can say about the nature of this property is that it must correlate reliably and lawfully with being right about *P*.[12] Over and above that, however, the property is gerrymandered. Craig maintains that "[t]here could be almost as many different answers [to the question as to what the precise nature of the property at issue is, CK] as there are types of thing that the inquirer might want to know about."[13] For instance, someone may have the relevant property with respect to the question whether you may park your car where you have just parked it in virtue of being a traffic warden.

---

[10] Craig (1990), p. 18.
[11] Cf. Ibid., pp. 21-22.
[12] Presumably, the modal feature spells out part of what the notion of reliable and lawful correlation at issue here amounts to.
[13] Ibid., p. 27.

At the same time, someone might have the relevant property with respect to the question of how many peaches there are in the fruit bowl in virtue of standing right next to it.

Craig goes on to explain various intuitions we have about the extension and intension of the concept of knowledge as ordinarily understood in terms of the concept of good informant so construed. Or, to be more precise, Craig argues that, given the above plausible remarks about the concept of good informant, if his initial hypothesis is true—that is, if the job of the concept of knowledge is to flag good informants—then we could be expected to have exactly these intuitions about its extension and intension. If his argument is successful, then that will provide crucial evidence that the initial hypothesis is indeed true. Craig covers a lot of ground here, more than I can hope to reproduce even in rough outline.[14] For that reason I will skip this part of Craig's argument and move on to a problem that Craig's hypothesis encounters. The reason I discuss this problem is twofold. First, the way Craig solves it constitutes a crucial part in his account of the concept of knowledge. At the same time, second, some of the ideas that Craig introduces in his discussion will be of importance later on in the thesis.

The problem that arises for Craig's hypothesis consists in the existence of cases that suggest that the relation between knowers and good informants is less intimate than one may think Craig's hypothesis requires. Here are a few such cases: Suppose Luigi has seen where the body is buried but he won't tell. Luigi knows where the body is but he is not a good informant about it. Or suppose that Matilda sees that the house is on fire and announces that it is on fire but no one believes her because she has made false announcements to that effect too many times. Matilda knows that the house is on fire but she is not a good informant either. Finally, the milkman may be able to solve a difficult mathematical problem but one would not consider asking him about it. Again,

---

[14] The most relevant material can be found in Craig (1990), chs. V, VII, VIII, XII and XIII.

he may know the solution to the problem but he is not a good informant.[15] Given that there are cases such that, intuitively, someone knows a proposition, *P*, while, at the same time, he is not a good informant about *P*, the pressure is on Craig's hypothesis. After all, if it is the job of the concept of knowledge to flag good informants, then it would seem that these cases suggest that the concept of knowledge does its job rather poorly.

Craig proposes to deal with the pressure such cases exert on his hypothesis by introducing the idea of the objectivisation of a concept. There is a subjective dimension to the concept of good informant about *P* as a person who has a true belief that *p* with a detectable property that correlates well with being right about *P*—*viz.*, that the property be *detectable*. In the first instance, someone is detectable by me, say, in the relevant sense, if he is accessible to me here and now, if I can recognise him as likely to be right about the proposition at issue and if the channels of communication are open between us. There is yet another dimension of subjectivity—*viz.*, that the informant be as likely to be right about the proposition at issue as my concerns require.[16] Now, in a community of inquirers and informants with a practice of passing on information between one another, there is reason to assign less importance to the subjective feature of detectability of good informants. Whether or not a prospective informant is accessible to me here and now, whether I can recognise him as having a property that correlates well with being right about the proposition at issue and whether the channels of communication are open between us will be of less importance since the prospective informant may be accessible to someone else at some other time, someone else may recognise that he has the property in question and the channels of communication may be open between him and someone else. The standards for detectability will be

---

[15] Cf. Craig (1990), p. 82.
[16] Cf. Ibid., p. 85.

relaxed.[17] Similarly, whether or not a prospective informant is as likely to be right about the proposition at issue as my concerns require will also be of less importance. After all, if, for instance, I want to recommend that person as an informant to someone else, what matters are not my concerns but his. Note that in contrast with the standards for detectability, objectivisation will drive the standards for concerns up. When one recommends an informant to someone else, for instance, one does not always know what the other person's concerns are. For that reason the recommended informant had better be suitable for anyone's concerns—no matter how stringent they may be. So, as in this respect, the demands on good informants will be tightened rather than relaxed.[18] In a community of inquirers with a practice of passing on information, then, there is reason to filter out the subjective dimensions of the concept of good informant, to objectivise that concept. The concept of knowledge, according to Craig, does exactly that. It is the objectivised version of the concept of good informant.[19]

Given that this is so, it can easily be seen that the apparently problematic cases are no longer problematic. After all, the reason why the subjects in these cases are not good informants is that they fail to satisfy some of the more subjective conditions of good informants. Thus, Luigi is not a good informant because the channels of communication between him and us are not open, while Matilda and the milkman are not good informants because we cannot recognise them as being likely to be right about the propositions at issue. Since, according to Craig, the concept of knowledge is the objectived version of the concept of good informant and since objectivisation filters out the subjective dimensions of the concept of good informant, the examples do not pose a serious threat to Craig's hypothesis after all—the pressure evaporates.

---

[17] Cf. Craig (1990), pp. 87-90.
[18] Cf. Ibid., p. 91.
[19] Cf. Ibid.

In conclusion, Craig offers us a hypothesis about what the concept of knowledge does for us—*viz.*, that it serves to flag good informants. This hypothesis, argues Craig, explains why we have several intuitions about the concept of knowledge. The fact that it does explain these intuitions and the fact that the hypothesis is a *prima facie* plausible hypothesis about the needs that could pressure a linguistic community to introduce the concept of knowledge confirm the hypothesis. At the same time, the concept of knowledge is not just identical with the concept of good informant. Rather it is the objectivised version of the latter concept.

Of course, if it is the very job of the concept of knowledge to flag good informants, then there is excellent reason to believe that there is a connection between knowers and good informants of the kind stated in the good informant platitude. But now recall that earlier on I have claimed that the minimalist can also extract support for the factivity and belief platitudes from Craig's arguments. How can he do so? To begin with, if the concept of knowledge is the objectivised version of the concept of good informant, then we may expect that the concept of knowledge will retain the application conditions of the concept of good informant that remain after the objectivisation of the latter concept. Now, recall that Craig has identified three conditions for the application of the concept of good informant: A truth condition, a belief condition and a more complex third condition. As Craig points out, however, the truth condition and the belief condition survive the objectivisation of the concept of good informant.[20] That is to say, however, that they will also be application conditions for the concept of knowledge. Since these two conditions are exactly the ones at issue in the factivity and

---

[20] Cf. Craig (1990), p. 90.

belief platitudes Craig's argument can be used to extract support for these platitudes as well.[21,22]

## 2 The Telling/Informing Platitude

### 2.1 The Telling/Informing Platitude, The Informative Speech Act Platitude, and the Knowledge Rule of Informative Speech Acts

Recall that the telling/informing platitude takes the form: If one tells/informs someone that $p$, then one represents oneself as knowing that $p$. It is plausible that the telling/informing platitude generalises to a class of speech acts that Kent Bach and Robert Harnish have aptly labelled 'informatives'.[23] According to Bach and Harnish, the class of informatives comprises speech acts associated with verbs such as 'advise', 'announce', 'apprise', 'disclose', 'inform', 'insist', 'notify', 'point out', 'report', 'reveal', 'tell' and 'testify'. The telling/informing platitude can accordingly be modified in the following way:

---

[21] As regards the complex third application condition of the concept of good informant, I will argue in chapter IV, section 3 that Craig's line of thought can be developed to motivate a safety condition on knowledge.

[22] Now one might wonder whether I am entitled to use Craig's argument in order to support any platitude of the minimalist framework. After all, what Craig does is to provide a characterisation of the concept of knowledge. That is, the aim of his argument is the same as the aim of the minimalist approach to the theory of knowledge. Given that this is so, should the two not be viewed as rivals over the correct characterisation of the concept of knowledge? And if so, should I not venture to establish that my approach has an edge over Craig's? It is not clear that I must or even that I should. After all what both approaches try to do in the first instance is to *illuminate* the concept of knowledge. However, it is plausible that something can be illuminated from different vantage points. For instance, a novel can be illuminated from a psychoanalytical vantage point, from a feminist vantage point and so on. Moreover, it is also far from clear that I will not be entitled to use the results of Craig's characterisation for my purposes: Just as the feminist may draw on the work of the psychoanalyst to improve her *feminist* illumination of the novel, it seems to me as though there is nothing intrinsically problematic with my using the results of Craig's characterisation to improve my own characterisation of the concept of knowledge. What is of some importance, however, is that the two vantage points be sufficiently different. It seems to me, however, as far as Craig's and my characterisations of the concept of knowledge are concerned, there is good reason to believe that the two are sufficiently different: While Craig characterises the concept of knowledge by tracing its origin, I characterise it by presently held platitudes. Moreover, the two approaches have different virtues: While Craig clearly provides a story as to what job the concept of knowledge does for us, the minimalist approach, I submit, can be used to develop a more fine-grained characterisation of the concept of knowledge and for that reason will be able to offer more fine-grained solutions to some of the relevant problems (most notably to the Gettier problem).

[23] Cf. Bach and Harnish (1979), p. 42.

<u>The Informative Speech Act Platitude</u>: If one performs any speech act of the class of informatives, then one thereby represents oneself as knowing its content.

A related platitude is the stronger assertion platitude according to which if one asserts that $p$ one thereby represents oneself as knowing that $p$. Since the weaker informative speech act platitude is less controversial while it will serve my purposes perfectly well, I will first develop arguments that support it. Afterwards, I will look at the prospects for defending the stronger and more controversial assertion platitude.

Since the informative speech act platitude is weaker than the assertion platitude, one need not be surprised to find out that the facts about usage that support the assertion platitude will also support the informative speech act platitude. These facts are relatively well understood from the work on assertion by Peter Unger, Michael Slote, Timothy Williamson and Keith DeRose.[24] It will be useful, however, to rehearse them in the present context. Before I begin, I would like point to a consequence of the informative speech act platitude in conjunction with a further uncontroversial assumption. If in performing an informative speech act that $p$ one represents oneself as knowing that $p$, and if—here comes the further assumption—one ought not represent oneself falsely, it follows that one ought not perform an informative speech act that $p$ unless one knows that $p$. This consequence of the informative speech act platitude may thus be labelled 'the knowledge rule of informative speech acts'[25]:

*Knowledge Rule of Informative Speech Acts*    One must: perform an informative speech act to the effect that $p$ only if one knows that $p$.[26]

Since the assumption that is needed to derive the knowledge rule of informative speech acts (henceforth simply 'knowledge rule') is so uncontroversial, I take it that if the

---

[24] Cf. Unger (1975), ch.6, Slote (1979), DeRose (1991), (2002) and Williamson (1996), (2000), ch.11.
[25] My terminology is adapted from Williamson's (2000) who labels the corresponding rule of assertion 'the knowledge rule of assertion'.
[26] Cf. Williamson (2000), p. 243.

informative speech act platitude is true, then so is the knowledge rule. It is also plausible that if the knowledge rule holds—that is, if one must perform an informative speech act that *p* only if one knows that *p*—then upon performing an informative speech act that *p* one will represent oneself as knowing that *p*.[27] In consequence, the informative speech act platitude and the knowledge rule may be regarded, to put it in DeRose's terms, "as two sides of the same coin."[28] Accordingly, I will use the informative speech act platitude and the knowledge rule to make the points I have to make as my purposes require. Let me then turn to the facts about usage that support the informative speech act platitude/the knowledge rule.

## 2.2   Support for the Informative Speech Act Platitude/Knowledge Rule

First, the fact that informative speech acts of the statements of the form 'I'm telling you/informing you/disclosing to you/revealing to you etc. that *p* but I don't know that *p*' are infelicitous but not inconsistent supports the informative speech act platitude.[29] For instance, it would be extremely odd if I uttered: "I am telling you: There is milk left; but I don't know that there is." The informative speech act platitude has a neat explanation as to why such informatives are infelicitous whilst being perfectly consistent. By telling you that the first conjunct is true I represent myself as knowing that there is milk left. At the same time, by telling you that the second conjunct is true I explicitly deny—and hence also represent it not to be the case—that I know that there is milk left. In this

---

[27] Cf. DeRose (2002), p. 179-80.

[28] Cf. Ibid., p. 180.

[29] It may be argued that this is not correct for all informative speech acts. For instance, it seems plausible that one cannot disclose that *p* unless one knows that *p*. Rather, when one does not know that *p* it may only seem to one as if one discloses that *p*, while in fact one does not. If this is correct, then knowledge is not just a (constitutive) rule of disclosing. It is a necessary condition. However, it is not plausible that knowledge is a necessary condition for the performance of all informative speech acts. For instance, one may announce or report that *p* without knowing that *p*. Still it remains true that announcements/reports of the form "*p* but I don't know that *p*" are infelicitous.

way, I represent a contradiction. More generally speaking, by the lights of the informative speech act platitude, if one performs an informative speech act that has as its content a statement of the form '*p* but I don't know that *p*' one represents a contradiction. For that reason, informative speech acts with contents of the form '*p* but I don't know that *p*' are infelicitous. At the same time, since the contradiction is only represented, the contents of the speech acts performed are perfectly consistent.[30]

Second, a normally legitimate challenge to an informative speech act is to ask how the utterer knows the content of the informative speech act he has performed. For instance, if my friend tells me that Holmes is in Italy, I may challenge his speech act by asking him how he knows that. If it were not the case that in telling me that Holmes is in Italy he represented himself as knowing that he is, he could, at least in some cases, just reply: "I never said I knew that Holmes is in Italy." His speech act of telling me that Holmes is in Italy could remain in good standing although he has nothing to say in response to the challenge. However, that does not seem right. My friend's reply is inappropriate and his inability to respond to the challenge does undermine the speech act's good standing.[31] A related datum that also supports the informative speech act platitude is that when people cannot appropriately respond to a challenge to their informative speech acts, they will typically retreat to something they know and thus represent themselves truly. For instance, supposing that my friend knows that Holmes is travelling through France and Italy and that he is at a stage of his journey where it is rather likely that he is in Italy, may retreat to telling me that it is *rather likely* that Holmes is in Italy. Again, the informative speech act platitude/knowledge rule provides a neat explanation of conversational patterns like these.[32]

---

[30] Cf. Williamson (2000), pp. 253-4.
[31] Cf. Ibid., pp. 252-3.
[32] Cf. Unger (1975), pp. 264.

Third, and most importantly, however, the informative speech act platitude/knowledge rule explains why it is always inappropriate to perform an informative speech act to the effect that a given ticket in a fair lottery will lose (or, for that matter, that it will win). Suppose you have bought a ticket in the state lottery. The winner of the lottery has been drawn thirty minutes ago. You missed the drawing and are now sitting in front to the TV waiting for the winning numbers to be announced on the news. I walk in and, knowing that you are waiting for the numbers to be announced, I tell you that your numbers did not win. Suppose that I did not hear the original announcement of the winning numbers. Rather, the evidence that supports the belief represented in my speech act is merely the probabilistic evidence against winning. Intuitively, my speech act of telling you is inappropriate. The informative speech act platitude/knowledge rule provides a simple and elegant explanation of this intuition: I must not tell you that your ticket did not win the lottery unless I know that it didn't. However, one cannot know merely on the basis of the probabilistic evidence against winning that a given ticket in a fair lottery will not/did not win. Consequently, I do not know that your ticket did not win the lottery. In telling you that your ticket did not win the lottery I violate the knowledge rule. In this way, the informative speech act platitude/knowledge rule explains the inappropriateness of my informative speech act.[33]

This line of thought can be generalised: Lottery cases can be used to show that merely probabilistic evidence for a proposition, *P*—that is, evidence for *P* such that the probability of *P* given the evidence is smaller than 1—is never good enough to render an informative speech act to the effect that *p* appropriate. For suppose that one has evidence for *P* such that given one's evidence the probability of *P*'s being true is x (where x<1) and that one's evidence is strong enough to render one's informative speech act that *p* appropriate. It is plausible that if such evidence is strong enough to

---

[33] Cf. Williamson (2000), pp. 247-8 for an argument that the conversational impropriety cannot be explained along Gricean lines.

render one's informative speech act that *p* appropriate, then if one has even stronger probabilistic evidence for another proposition, *Q*, one's informative speech to the effect that *q* is true will also be appropriate. However, for any degree, x, of probabilistic evidence smaller than 1, there will be a (fair) lottery, L, such that the probability that a given ticket in L loses is higher than x. Hence, if probabilistic evidence of degree x renders some informative speech act appropriate, an informative speech act to the effect that a given ticket in L will lose is appropriate, too. However, informative speech acts to the effect that given tickets in fair lotteries will lose are never appropriate. Hence, no degree of probabilistic evidence smaller than 1 will be strong enough to render an informative speech act appropriate.[34] The informative speech act platitude/knowledge rule can avoid this difficulty. Knowledge is factive in the sense that knowledge of a proposition entails the truth of the proposition known. Given that this is so, the probability of a known proposition given one's knowledge is 1. Accordingly, if one must perform an informative speech act to the effect that *p* only if one knows *P*, one must do so only if one is in a state relative to which the probability of *P* is 1. The fact that the odds against winning imaginary lotteries can be indefinitely high (short of 1) conjoined with the fact that informative speech acts to the effect that a given ticket in such a lottery will not win are bound to be inappropriate have no force against the knowledge rule.

However, there is yet another argument that can be made in support of the knowledge rule which exploits the relations between knowers and good informants outlined in the previous section. Consider a state of nature scenario of the kind envisaged by Craig. We have an inquirer who wants to acquire a belief about *P* from another member of his linguistic community. The question Craig is interested in is: What features does the inquirer want his prospective informant to have? Craig's answer

---

[34] Cf. Williamson (2000), p. 251.

essentially is that he will want her to be (something close to) a knower and to have some features that make her suitable for him, the inquirer, as an informant. But now consider the following question: If the interactions between inquirers and informants were to be enshrined into a general practice of passing on information, what would such a practice look like? Well, to begin with, it is plausible that the features that inquirers *want* their informants to have will be turned into *demands* on informants. Moreover, it is also plausible that the subjective features that individual inquirers want their informants to have will be of rather little importance to the demands placed upon informants by the practice. After all, if all the features that may be subjectively important to individual inquirers were enshrined in the demands imposed on informants by the practice, the demands would in all likelihood be too hard to be met and the practice would be likely to be useless. The features that will go into the practice will rather abstract from the subjective features individual inquires may want their informants to have. As Craig has argued, what remains once we have abstracted from the subjective features individual inquirers want their informants to have is knowledge. For that reason, if the interactions between inquirers and informants were to be enshrined in a practice of passing on information, then one demand on informants would be that they know the propositions they pass on as information. Now it is undeniable that we have a practice of passing on information. Moreover, it is plausible that the paradigmatic speech acts via which information is passed on between informants and inquirers are informative speech acts. Accordingly, if—as is plausible—the practice of passing on information originated form interactions between inquirers and informants of the sort Craig describes and if Craig's thesis that the concept of knowledge is what remains after objectivisation of the concept of good informant is correct—as we have seen there is reason to believe it is—then we must expect that a demand on the practice with informative speech acts will be that the speaker knows the content of his speech act. Or, in other words, we must expect that

one demand governing our practice with informative speech acts takes the following form: Don't perform an informative speech act with content *P* unless you know that *p*. This, of course, is just another version of the knowledge rule of informative speech acts. In this way, the knowledge rule of informative speech acts is a natural extension of the sorts of considerations that supported the good informant platitude. The good informant and the informative speech act platitude provide a first taste of how deeply interwoven the platitudes are.

## 3  *Interlude: The Knowledge Rule of* Assertion

In this section I would like to look at the arguments that Williamson et al. have given in support of the knowledge rule of assertion and suggest ways in which the knowledge rule can be strengthened to cover assertions as well.

As I have already anticipated earlier the data Williamson et al. take to support the knowledge rule of assertion are parallel to the data that support the informative speech act platitude/knowledge rule of informative speech acts. Assertions of Moorean and lottery statements are infelicitous: One cannot felicitously utter: "I assert that *p* but I don't know that *p*" or "I assert that your ticket won't win the lottery." And one cannot felicitously utter, flat-out: "*p* but I don't know that *p*" or "Your ticket won't win the lottery." Moreover, one can typically challenge the speech act performed by a flat-out utterance of "*p*" or an utterance of "I assert that *p*" by asking: "How do you know?" If the utterer cannot respond to this challenge, his speech act will not remain in good standing: He will have to withdraw it or weaken it in suitable ways. These considerations support the knowledge rule of assertion because utterances of sentences prefixed with 'I assert that' and flat-out utterances of declarative sentences are paradigmatic ways of making assertions. Since the knowledge rule of assertion will

explain all of these data as neatly and elegantly as the knowledge rule of informative speech acts explains the corresponding data for informative speech acts, the infelicity of assertions of Moorean and lottery statements as well as conversational patterns support the knowledge rule of assertion.

Now, why might one think that the knowledge rule of assertion is more controversial than the knowledge rule of informative speech acts? After all, the data seem to support the former just as well as the latter. So, where is the problem? The problem for the knowledge rule of assertion is rooted in a further fact about assertions—*viz.*, that "the default use of declarative sentences is to make assertions".[35] In what follows I will first outline how the problem can be generated and then propose a solution on behalf of the defender of the knowledge rule of assertion.

## 3.1   A Problem for the Knowledge Rule of Assertion

To bring out the problem consider, for instance, utterances of the following sentences:

(a) Blair's decision to take sides with Bush on the war on Iraq was a bad one (uttered by Doris Schröder in a conversation with Gerhard on the day after Blair has decided to go to war in Iraq).
(b) Frank Lloyd Wright is the better architect than Mies van der Rohe (uttered by Roger Scruton in a conversation with Jeremy Paxman).
(c) The French will not attack before nightfall (uttered by Captain Aubrey upon seeing the French manoeuvre).[36]
(d) This is the work of Professor Moriarty (uttered by Sherlock Holmes upon entering the scene of the crime).[37]

All of these utterances are standard uses of declarative sentences. If it is true that the default use of declarative sentences is to make assertions, however, then, since (a)-(d) are standard uses of declarative sentences, there is *prima facie* reason to believe that the speech acts performed in uttering them are assertions. Moreover, it seems that,

---

[35] Cf. Williamson (2000), p. 258.
[36] Cf. Weiner (2005), p. 231.
[37] Cf. Ibid.

intuitively, all of these speech acts will be felicitous. However, it also seems that, at the same time, utterers of (a)-(d) might not know the contents of their utterances. Whether Blair's decision was a bad one may, at the time of Doris Schröder's utterance, be something that only time can tell. Similarly, it is hard to see what sort of evidence could be strong enough to turn one's belief that Wright is the better architect than van der Rohe (even if true and held by Roger Scruton) into knowledge. It may be evident that looking at manoeuvres will not give Captain Aubrey the evidence needed for knowing the content of his utterance and it may also be evident that Holmes's first impression of the scene of the crime won't do either. Given that this is so, these cases pose a problem for the knowledge rule of assertion: Intuitively, the speakers felicitously perform speech acts that, on the face of it, are assertions, while, at the same time, they do not know their contents. The same point can be put rather nicely once it is appreciated that just as the knowledge rule of informative speech acts and the informative speech act platitude are two sides of the same coin so are the knowledge rule of assertion and the corresponding assertion platitude: If one asserts that $p$ then one represents oneself as knowing that $p$. While, on the face of it, utterances of (a)-(d) have to be construed as assertions, the speakers can hardly be said to represent themselves as knowing their contents. Accordingly, if it is correct that utterances of such sentences are to be taken at face value and construed as assertions, then the assertion platitude—and hence the knowledge rule of assertion—do not hold.[38]

## 3.2  The Knowledge Rule of Assertion Defended

How can the advocate of the knowledge rule of assertion respond to this problem? It seems to me that the following may be an appealing shot at a solution. First, the

---

[38] Weiner (2005) pushes this objection.

defender of the knowledge rule of assertion may point out that the relevant data—Moorean and lottery statements as well as conversational patterns—provide *prima facie* evidence for the knowledge rule of assertion. Second, he may argue that the default use of declarative sentences is to make claims to knowledge. If the defender of the knowledge rule manages to establish this much, the objection will no longer go through. After all, since the utterances in the problematic cases are not claims to knowledge, they will not be counted as default uses of declarative sentences. Accordingly, objectors to the knowledge rule of assertion cannot appeal to the fact that the default use of declarative utterances is to make assertions to support the claim that the problematic utterances are assertions. Finally and third, he may then provide an alternative explanation of how such utterances of declarative sentences can be appropriate even if their contents are not known.

### 3.2.1   The Solution Part 1: The Default Use of Declarative Sentences is to Make Claims to Knowledge

Let me now turn to the details of the argument. How can the defender of the knowledge rule of assertion make his second point—*viz.*, that the default use of declarative sentences is to make claims to knowledge? I would suggest an argument that exploits general principles of communication. As DeRose points out, the following principle holds for contents of assertions: "[W]hen you're in a position to assert either of two things, then, other things being equal, if you assert either of them, you should assert the stronger."[39] It is plausible that this principle generalises to constative speech acts—one of the broadest classes of speech acts that comprises the spectrum of speech acts from

---

[39] DeRose (2002), p. 175.

informatives over expressions of opinion, conjectures and speculations to guesses—in the following way:

| *Strongest Constative Speech Act (SCSA)* | When one is in a position to perform either of two constative speech acts, then, other things being equal, if one performs either of them, one should perform the stronger. |
|---|---|

The data that support DeRose's principle governing contents of assertion and SCSA are parallel. As DeRose points out in asserting, for instance, that it is *possible* that *p* when one knows that *p,* and therefore could have appropriately asserted that *p*, one will be misleading the hearer. One has made a weaker statement than one could have made.[40] In the same way, in performing a conjecture when one could have performed an informative speech act, for instance, one will be misleading the hearer, too. If I could have just told you that Holmes is in Italy (because he has just called me from Rome) it will be misleading if I merely *conjectured*, for instance, that he is. Further data that support SCSA are the fact that the speech acts performed by utterances of the following sentences are, intuitively, odd: "I am of the opinion that the cat is in the living room. I can see it sitting on the sofa." "I conjecture that Holmes is in Italy. He just called me from Rome." SCSA gives a neat explanation of the oddity of the speech acts performed by these utterances. The warrants that, in the second half of my utterance, I give my hearer to understand I have for the contents of the speech acts performed in the first half are stronger than the warrants needed for such speech acts to be appropriate. They would warrant stronger speech acts than the ones performed. Accordingly, by SCSA, I should have performed a stronger speech act—one that reflects the strength of my warrant. It is, then, no surprise that it is odd if I don't.

Now, an assumption that underlies the above line of thought is that the strength of a given constative speech act can be modelled on the strength of the attitude—such as

---

[40] Cf. DeRose (2002), pp. 174-75.

justified belief of varying degrees, knowledge etc.—represented in performing it. However, that assumption appears to be independently plausible. Bach and Harnish, for instance, individuate speech acts in general by the illocutionary intents of the speakers and argue that "illocutionary intents correspond to expressed attitudes."[41] If their suggestion is along the right lines, a natural consequence is that speech act strength can be modelled on the strength of the attitude represented in performing them.[42]

Now suppose SCSA holds. Suppose also that a certain form of words is normally used to perform a number of speech acts of differing strengths—all of them constatives. In that case the mere utterance of this form of words underdetermines the type of speech act performed by them. Suppose, furthermore, that in a particular case it is not otherwise clear—for instance, from context—which speech act has been performed by an utterance of this form of words. If so, then, for all the hearer knows, the speaker has performed any of the types of speech act normally performed in so many words. But the hearer will presume that the speaker conforms with general principles of conversational cooperation. Since all of the speech acts normally performed in so many words are, *ex hypothesi*, constatives, that means that the hearer will presume that the speaker conforms with SCSA and that, therefore, he has performed the strongest speech act he could have performed. By the same token, if, in such a case, the speaker uses the form of words at issue to perform a weaker speech act, he will be misleading his audience. Since he ought not mislead his audience, in such a case, he ought not use this form of words to perform a speech act other than the strongest speech act normally performed in uttering so many words. More generally

---

[41] Cf. Bach and Harnish (1979), p. 39.

[42] True, Bach and Harnish's conception of speech acts rejects an approach to speech acts in terms of constitutive rules that, one may think, I have so far taken for granted. However, even if one wants to stick to a constitutive rules approach to speech acts it is plausible that the different attitudes that, according to Bach and Harnish, correspond to the speakers illocutionary intents will figure in the constitutive rules governing the speech acts and will be represented when the relevant speech act is performed. Again, there appear to be reasonable prospects to model the strength of the speech acts on the strength of the attitudes that figure in the constitutive rules.

speaking, then, in case a certain form of words is normally used to perform speech acts of differing strengths all of which are constatives, then, unless it is otherwise clear—for instance, from context—that one is not performing the strongest speech act normally performed in so many words, one ought not use this form of words to perform a speech act other than the strongest speech act normally performed in so many words. Since it must not generally be presumed that it is otherwise clear that a speech act weaker than the strongest speech act normally performed in uttering so many words is being performed when it is performed, we may say that, in case a certain form of words is normally used to perform a number of speech acts of differing strengths all of which are constatives, *all else equal*, one ought not use this form of words to perform a speech act other than the strongest speech act normally performed in utterances of this form of words. Given that this is so, it is plausible that, in said case, the default speech act performed by an utterance of this form of words is the strongest speech act normally performed by such an utterance.[43] Given that speech act strength can be modelled on the strength of the represented attitudes, however, that is to say that, in said case, the default speech act performed by an utterance of a certain form of words is the one in which the strongest attitude is represented.

Now, I take it that all the speech acts normally performed by utterances of declarative sentences are constatives. Given that this is so, what the defender of the knowledge rule of assertion needs in order to employ these thoughts for his purposes is, of course, an argument to the effect that knowledge is the strongest attitude normally

---

[43] Of course, even if, in the circumstances at issue, the default speech act performed by an utterance of a certain form of words is the strongest speech act normally so performed that does not mean that we may not often use that form of words to perform speech acts weaker than the strongest. After all, if the all else equal proviso fails, one can felicitously perform weaker speech acts by this form of words. And the all else equal proviso may fail rather often. For instance, it may often be the case that it is clear from context that the speech act performed is weaker than the strongest speech act normally performed in so many words. In this way it may be argued that although, in case a certain form of words is normally used to perform a number of constative speech acts of differing strengths, the default speech act performed in so many words is the strongest speech act normally performed in so many words, the same form of words can be normally used to perform weaker speech acts as well.

represented in utterances of declarative sentences. Here are two considerations providing support for this claim. First, knowledge is factive in the sense that if one knows that $p$, then $P$ is also true. At the same time, the other attitudes that are plausible candidates to be attitudes that are represented in standard uses of declarative sentences are justified belief of varying degrees. Now, according to our intuitive understanding, justified belief is not factive. At the same time, belief certainly aims at truth and justification is truth-conducive. In this way, justification is a good guide to but does not guarantee the truth of the proposition believed. As opposed to that, knowledge is not only a good guide to truth but also guarantees the truth of the proposition known. Given that this is so, we have some *prima facie* evidence that knowledge is the stronger attitude than justified belief.[44] Again, given that justified beliefs of varying degrees are the most promising alternative candidates for attitudes normally represented in the speech acts performed by utterances of declarative sentences, this would be a first step towards arguing that knowledge is the strongest attitude normally represented in utterances of declarative sentences.

A second step towards that end would again focus on the factivity of knowledge. Being in the state of knowledge guarantees the truth of the proposition known. Accordingly, if by uttering a certain form of words one represents oneself as knowing a certain proposition, one represents oneself as being in a state that guarantees the truth of the proposition expressed by one's utterance. Now, I take it that constative speech acts aim at truth.[45] If so, then in performing a constative speech act in which one represents oneself as knowing its content one represents oneself as being in a state that secures the

---

[44] Of course, in order to argue that knowledge is strictly logically stronger than justified belief (intuitively understood) one would have to show that knowledge entails justified belief (so understood). One way of doing this is by adopting the corresponding platitude. Although it seems to me that there is something to be said for such a platitude I will leave it an open question as to whether it should enter the minimalist framework. Accordingly, I will stick with *prima facie* evidence that knowledge is stronger than justified belief just provided.

[45] This claim can be supported by some of Bach and Harnish's work on speech acts. According to Bach and Harnish, belief is the attitude expressed in all constative speech acts (Bach and Harnish (1979), p. 41). Provided, additionally, that belief aims at truth, it is also plausible that constatives aim at truth.

aim of the speech act performed. Moreover, as I have already pointed out, I also take it that the speech acts performed by normal uses of declarative sentences will all belong to the class of constatives and thus aim at truth. Given that this is so, one might want to ask: What interest could we have in having a normal use of declarative sentences by means of which the speaker represents himself as being in a stronger state than one that guarantees the truth of the proposition expressed by his utterance? It is not obvious that we could have such an interest. If we don't have such an interest, however, why should we bother with introducing such a use in the first place? Again, it is not clear that we should. The point here is, of course, that given that knowledge is factive, given that knowledge is an attitude normally represented in utterances of declarative sentences, and given that the speech acts performed by normal uses of declarative sentences all belong to the class of constatives, there is no obvious need to have a normal use of declarative sentences in which a state stronger than knowledge is represented. Knowledge already gives us everything we need. In this way, I suggest, the defender of the knowledge rule can argue that knowledge is the strongest attitude normally represented in the speech acts performed by utterances of declarative sentences. If knowledge is the strongest attitude normally so represented, however, then it is also the attitude represented in default uses of declarative sentences. Or, in other words, the default use of declarative sentences is to make claims to knowledge. Given that the problematic utterances (a)-(d) are not claims to knowledge, they are not default uses of declarative sentences. So, the fact that the default use of declarative sentences is to make assertions does not serve to support the idea that the utterances in (a)-(d) are assertions. The objection won't go through.

### 3.2.2 The Solution Part 2: The Speech Acts Performed in the Problematic Cases are Indirect Speech Acts

What remains to be done in order to fully allay the worries raised by the objection is to tell a story as to how in the problematic cases the speech acts by performed by utterances of declarative sentences with unknown contents can be felicitous. A promising strategy to pursue, and a strategy that fits well with what I have said so far is to argue that the relevant utterances of (a)-(d) are *indirect* speech acts. Indirect speech acts are speech acts such that one performs one speech act by way of performing another. For instance, while an utterance of "Can you reach the salt?" is in the first instance a *question* concerning the hearer's basic physical abilities, it is often made as a *request* to pass the salt on to one.[46] There are certain conditions that must obtain if an indirect speech act is to be felicitous as an act of communication. Most importantly, the hearer has to be able to figure out that the speaker is performing an indirect speech act. The resources the hearer has for this task are his background knowledge of the situation, an (implicit) grasp of the different kinds of speech acts there are, and the presumption that the speaker conforms with principles of conversational cooperation.[47] In the case I just mentioned, the hearer will be able to draw on his background knowledge that he is sitting at the table with the speaker having dinner, that salt is something that people are likely to wish to add to their meal during the course of a dinner, that they have not been talking about his basic physical abilities etc. In conjunction with the presumption that the speaker conforms with general principles of conversational cooperation, this background knowledge renders it unlikely that the speech act performed is to be taken at face value, that is, as a question concerning the hearer's basic physical abilities.[48] By the same token, it is likely that the speaker tries to communicate something else. It is on

---

[46] Cf. Searle (1975), p. 65.
[47] Cf. Ibid., p. 63.
[48] Cf. Ibid.

the hearer to figure out what it is that he wants to communicate. Although there are stories to be told as to how the hearer figures out which indirect speech act has been performed[49], they need not concern us here. Suffice it to say that hearers like us exhibit a tendency to correctly figure out which indirect speech act has been performed.

If the fact that some speech acts performed by utterances of declarative sentences with unknown contents are felicitous is to be explained by the idea that the speech acts performed are indirect, then we may expect there to be a parallel story that renders it unlikely that the speech act performed is an assertion. Now, it seems to me that such a story is available. In all of the problematic cases, facts about the conversational context make it clear that the speech act performed is not a claim to knowledge—and hence, by the lights of defenders of the knowledge rule of assertion, not an assertion. In the first problematic case, for instance, it is clear from context that only time (and hence not Doris Schröder) can tell whether Blair's decision to take sides with Bush on the war in Iraq was a bad one. Accordingly, given the presumption of conformance with general principles of conversational cooperation, it is unlikely that Doris's utterance of a simple declarative sentence with that content is to be taken at face value—that is, as a claim to knowledge. By the same token, it is likely that the speech act performed by Doris is of a different nature. It is on the hearer, Gerhard, to figure out what that nature is. In general we are good at figuring indirectness out. In the present case, for instance, Gerhard will most likely take the utterances as an expression of political opinion or, given his position as the chancellor of Germany at the time, as an advice not to take sides with Bush on the issue. The defender of the knowledge rule of assertion may tell a similar story for the other problematic cases. The indirect speech act performed in case (b) may be an indirect expression of architecture-historical opinion, while the one performed in case (c) is, presumably, an indirect prediction and the one

---

[49] Cf. Searle (1975), p. 63.

performed in case (d) an indirect retrodiction. All of these speech acts can be felicitously performed because the epistemic constraints governing them are weaker than knowledge.[50]

It is noteworthy that the explanation in terms of indirection of the felicity of the speech acts performed in utterances of (a)-(d) fits well with the initial argument that the default use of declarative sentences is to make claims to knowledge. Recall that what that argument suggested was that, all else equal, one ought not utter a declarative sentence unless one knows its content. The all else equal proviso will fail just in case it is otherwise clear that the claim made by such an utterance is not a claim to knowledge. Given that this is so, however, an explanation of the felicity of utterances of declarative sentences with unknown contents in terms of indirection will be appealing. After all, if it is clear to the hearer that no claim to knowledge is made, then the presumption of conformance with principles of conversational cooperation renders it likely that some other claim is made. It is on the hearer to figure out what sort of claim other than a claim to knowledge the speaker has made. This phenomenon, however, is characteristic of indirect speech acts. In this way, the two parts of the treatment of the problematic cases I have suggested fit well together.

Finally, I would like to test the explanation of the problematic cases in terms of indirection against a range of further cases. The first set of cases comprises variations of the problematic cases. As opposed to the problematic cases as construed by the objector, these variations are construed in such a way that it isn't otherwise clear that the contents of the utterances of the relevant declarative sentences are not known. If the explanation I have suggested is along the right lines, we must expect to find that, intuitively, the speech acts performed by the relevant utterances will not be felicitous—that is, of

---

[50] This is supported by the fact that these speech acts can be felicitous even if the sentences uttered are Moorean- or lottery sentences. The speech acts performed by uttering: "I predict/retrodict/opine that $p$ but I don't know that $p$" can all be felicitous as well as the ones performed by uttering: "I predict/retrodict/opine that your ticket won't/didn't win the lottery."

course, unless their contents are known. Let us turn to the cases. Suppose you are talking to a historian of architecture who is one of the great authorities in the field (and you know that to be the case). You yourself don't know much about architecture. In particular, you have never heard of Wright or van der Rohe (and the historian of architecture knows that to be the case). Suppose the historian of architecture utters (b)— "Wright is the better architect than van der Rohe"—although he doesn't know the proposition (b) expresses. In this case, intuitively, the speech act performed is inappropriate. You will naturally take the historian of architecture to make a claim to knowledge. After all, it is plausible that there are some cases in which it can be known that one architect is the better than another and you have been given no reason to believe that this is not the case with Wright and van der Rohe. Since the speech act performed is a constative you will presume that the historian of architecture conforms with SCSA and performs the strongest speech act he could have performed, which is a claim to knowledge. Accordingly, his utterance will be misleading you and his speech act will be infelicitous. Similarly, suppose that Captain Aubrey is talking to someone who does not have the slightest clue about warfare. He doesn't know what exactly a manoeuvre is, that from ways of manoeuvring it can be predicted but typically not known what move the enemy will make next etc. Suppose Aubrey utters (c): "The French will attack upon nightfall." It would seem that a perfectly natural reaction of the layman interlocutor would be to ask how Aubrey knows that they will. Again, I take it that this indicates that Aubrey's speech act is understood as a claim to knowledge. And his speech act will not remain in good standing if he cannot provide an appropriate response to the challenge (which, since, *ex hypothesi*, he doesn't know the content of his speech act, he cannot).[51] Similar scenarios can be constructed for the other utterances.

---

[51] It is noteworthy that Aubrey may give some response that will satisfy his audience. For instance, Aubrey may answer: "I know it from the way they manoeuvre." Given that he doesn't actually know it,

These data further support the explanation I have suggested. In variations of the problematic cases in which it is not otherwise clear that the speech acts performed by utterances of the relevant sentences aren't claims to knowledge, they will be naturally taken as claims to knowledge. Accordingly, they can be appropriate only if their contents are in fact known. Given that the default use of declarative sentences is to make claims to knowledge and given that indirection requires that there is reason to believe that the default speech act is not the speech act actually performed this is exactly what we may expect.

On the other hand, the explanation I have suggested will also be supported if there are cases in which the speech acts performed by utterances of lottery or Moorean sentences can be felicitous where it is otherwise clear that the speech acts performed are not claims to knowledge. Although I find it hard to produce Moorean sentences that satisfy these criteria[52] it seems to me that there are lottery sentences that do. For instance, as Williamson notices, the speech acts performed by utterances of lottery sentences can be felicitous if the sentences are uttered in a "special jocular tone".[53] Here the special tone makes it clear that the speech act performed is not a claim to knowledge. Given the presumption of conformity with principles of conversational cooperation there is reason for the hearer to try to figure out which speech act has been performed instead. Plausible suggestions, it seems to me, are, for instance, that the speech act performed is advice not to count on one's ticket winning or not to invest to much emotional energy etc.

---

this is not an adequate response to the initial challenge. The hearer is satisfied because his epistemic position is too impoverished to assess Aubrey's backup claim.

[52] The difficulty of constructing cases in which speech acts performed by utterances of Moorean sentences are felicitous may be grounded in the fact that if it is otherwise clear that one doesn't know their content, then in saying that one does not know them one will be flouting Grice's Maxim of Quantity (cf. Grice (1991), p. 26) in which case, of course, the speech acts performed will, again, be infelicitous.

[53] Cf. Williamson (2000), p. 246.

In conclusion, there is reason to believe that not only informative speech acts but also assertions are governed by the knowledge rule. Since the knowledge rule of assertion and the assertion platitude are "two sides of the same coin", we may thus be tempted to extent the list of platitudes by the following additional platitude:

> The Assertion Platitude: If one asserts that $p$, then one thereby represents oneself as knowing that $p$.[54]

## 4 The Closure Platitude

According to the closure platitude, one knows the conclusions of competently executed deductions provided that one knows the premises and comes to believe the conclusions on the basis of the deductions. Just like the other platitudes, the closure platitude is supported by facts about our use of the word 'know' in ordinary language and thought. Most obviously, we credit mathematicians and logicians with knowledge of theorems they have deduced from known axioms and on that basis have come to believe them. However, there are also less specialist examples: If, for instance, Kim acquires testimonial knowledge that she has exactly three aunts and two uncles that belong to her mother's family, she deduces and on that basis comes to believe that her mother has more sisters than brothers, intuitively, she knows what she has thus come to believe. Or

---

[54] The reader who is familiar with Williamson's arguments concerning the relation between assertion and knowledge may wonder what becomes of Williamson's knowledge *account* of assertion, that is, the further thesis that the knowledge rule is the *unique* rule governing assertion and that it is *individuating* for assertion (cf. Williamson (2000), p. 241-2). After all, if the arguments presented in the last two sections are sound, then not only assertion but more generally informative speech acts are governed by the knowledge rule. If so, then it would seem that the knowledge account of assertion must be false. However, either of the following two lines of thought may make progress towards reconciling the arguments presented above and Williamson's idea of a knowledge account: First, Williamson could be construed as defending a knowledge account of informative speech acts. However, Williamson construes the class of informative speech acts in a slightly different way than I have proposed: The speech act of assertion is also a member of the class of informative speech acts. Alternatively, Williamson can be construed as holding on to the idea of a knowledge account of assertion but reject the idea that assertion is on a par with the types of speech act that fall in the category of informatives. Picking up an idea that surfaces in other areas of his work, he may maintain that assertion is the most general speech act governed by knowledge. On this interpretation, he may construe speech types such as telling, informing, disclosing etc. as different ways of making assertions in a similar way as seeing that $p$, hearing that $p$ etc. are different ways of knowing that $p$ (cf. Williamson (2000), pp. 33-41).

if I am informed that Chelsea scored in the 111th minute and deduce and thereby come to believe that the match has gone to extra time, intuitively, my belief qualifies as knowledge as well.

There are, however, some considerations that put pressure on the idea that the closure platitude holds. Or, to be more precise, there are *prima facie* counterexamples to closure. Most famously, Fred Dretske has proposed the following counterexample: Suppose I am at the zoo. I am looking at the pen that is clearly marked as the zebra-pen. I see what looks like a zebra to me and come to believe that the animal I am looking at is a zebra. My belief is true. Intuitively, my belief qualifies as knowledge. At the same time, intuitively, I do not know that the animal I am looking at is not a mule that is cleverly disguised to look like a zebra. After all, it would seem that in order to know *that* I would have to be able to tell zebras from mules so disguised which, unfortunately, I can't. However, the fact that the animal I am looking at is a zebra entails that it is not a cleverly disguised mule. I may be aware of the entailment and I may even have competently deduced the consequent from the antecedent. All that does not change our intuitions, first, that I know that the animal I am looking at is a zebra and, second, that I do not know that it is not a cleverly disguised mule. Hence, the case constitutes a *prima facie* counterexample to closure.[55]

It is noteworthy, however, that, on closer inspection, the *prima facie* counterexamples are much less plausible than they may initially appear to be. Thus Jonathan Vogel (1990) and Pritchard (2007*b*) have argued that Dretske-style cases do not constitute genuine counterexamples to the closure principle in roughly the following way: First, they point out that the subjects possess a wide range of background knowledge that constitutes evidence favouring the proposition deduced over its negation. In the zebra case, for instance, I have a wide range of background knowledge

---

[55] Cf. Dretske (1970), pp. 1014-16.

including knowledge that the pen is clearly marked as the zebra-pen, that zoos typically exhibit genuine specimen, that it would be pointless to deceive customers and so on. This background knowledge provides evidence that favours the proposition that the animal I am looking at is a zebra over the proposition that it is a cleverly disguised mule. Given, second, the additional plausible assumption that one can know in virtue of believing truly on the basis of such favouring evidence, it is, on second thought, much less plausible that one does not know the conclusions of one's deductions in such cases. So, given that I have the right background knowledge, I may know that the animal I am looking at is not a cleverly disguised mule after all. On the other hand, third, if one competently deduces a proposition from a set of others and finds that one has no evidence favouring the proposition deduced over its negation, then it is also no longer clear that one knows all of one's premises. After all, one is now aware of an error possibility to the conjunction of the premises that one's evidence does not speak to in any way. However, it is not at all clear—indeed it is rather implausible—that a belief that is held in the face of error possibilities that one's evidence does not address at all is still in good enough standing to qualify as knowledge.[56] Given that all this is correct, however, the alleged counterexamples to closure turn out to be less convincing than they may initially have appeared to be.

Moreover, as John Hawthorne has recently argued, accepting the *prima facie* counterexamples to closure as genuine counterexamples incurs significant further theoretical costs. It is plausible that if closure fails, then there will be counterexamples of the kind Drestke adduces in virtue of which it fails. Moreover, it is plausible that, if

---

[56] One might be inclined to object that such a line will commit its defenders to the implausible claim that becoming aware of sceptical error-possibilities undermines every piece of ordinary empirical knowledge one may antecedently have had. In response to this objection it may be noted that scepticism poses a problem for every theory of knowledge and that, for that reason, sceptical counterexamples to closure will in and of themselves not provide sufficient motivation for denying closure. It is also noteworthy that, presumably for this very reason, the examples that Dretske takes to be counterexamples to closure are not sceptical ones.

there are such counterexamples to closure, we can know that closure fails in virtue of these counterexamples. As Hawthorne points out, however, given that this is so, we will have to give up one of two further highly intuitive principles:

> *The Equivalence Principle*: If one knows *a priori* (with certainty) that P is equivalent to Q and knows P, and competently deduces Q from P … [and thereby comes to believe Q, CK] one knows Q.[57]

> *Distribution*: If one knows the conjunction of P and Q, then as long as one is able to deduce P, one is in a position to know P (and as long as one is able to deduce Q, one is in a position to know Q).[58]

Suppose, in the situation Dretske describes, one knows that x is a zebra. The proposition that x is a zebra is *a priori* equivalent to the proposition that x is a zebra and ~x is a non-zebra cleverly disguised to look like a zebra. Suppose one knows this to be the case, deduces the latter proposition from the former and thereby comes to believe it. By the equivalence principle, one now knows that x is a zebra and ~x is a non-zebra cleverly disguised to look like a zebra. If one is able, as many of us will be, to deduce the second conjunct from the conjunction, then, by the distribution principle, one is in a position to know that ~x is a non-zebra cleverly disguised to look like a zebra. However, the intuition the *prima facie* counterexamples exploits is precisely that one cannot know such a proposition in the kind of situation one is in. Accordingly, if we take these examples to be genuine counterexamples to closure, then we will also have to give up either the equivalence or the distribution principle. And that would mean to have to give up at least one further highly intuitive principle.[59]

---

[57] Hawthorne (2005), p. 31

[58] Ibid. Since Hawthorne's Distribution principle is entailed by the stronger distribution principle—necessarily, if one knows that *p* and *q*, then one knows that *p* and one knows that *q*—the stronger distribution principle fails if the weaker does.

[59] Cf. Hawthorne (2005), p. 32. It might be objected that those who take Dretske-style examples as genuine counterexamples to closure may drop the equivalence principle without further ado. However, as Hawthorne points out ((2005), p. 31), it is not clear that the considerations that drive our intuitions in the counterexamples to closure serve to put pressure on the equivalence principle. What drives our intuitions in the counterexamples to closure is that while we are able to distinguish situations in which the antecedent of our deduction is true from at least some situations in which it is false, we are unable to distinguish situations in which the consequent of the deduction is true from any situations in which it is

The second unpleasant consequence Hawthorne notices is that if Dretske-style counterexamples are genuine counterexamples and we can know closure to be false in virtue of some of them, then the informative speech act platitude/knowledge rule of informative speech acts will break down. Suppose Dretske's zebra case does constitute a genuine counterexample to closure and I know this to be the case. Suppose we are at the zoo and you utter: "Tell me: Is the animal we are looking at a zebra?" to which I reply affirmatively. Then you ask me to tell you whether I know that that entails that it's not a cleverly disguised mule. Again my answer is: "Yes". Finally, you ask me whether I also acknowledge that the animal is not a cleverly disguised mule. Knowing that I don't know this to be the case and not wanting to violate the knowledge rule I reply: "No I don't acknowledge that." At the same time, since I know both that the animal is a zebra and that its being a zebra entails its not being a cleverly disguised mule, I am unwilling to withdraw any of my earlier claims. As Hawthorne points out I resemble "perfectly Lewis Carroll's Tortoise, that familiar object of ridicule who was perfectly willing to accept the premises of a *modus ponens* argument but was unwilling to accept the conclusion."[60]

In this way, then, there are a number of reasons to resist the thought that Dretske-style *prima facie* counterexamples are genuine counterexamples to closure. Since it is plausible that if closure fails then it fails in virtue of some Dretske-style counterexample, these considerations provide us with some reason to believe that closure does not fail.

---

false. As regards equivalent propositions, however, the situations in which the antecedent and the consequent are true and false respectively are exactly the same. Consequently, we will be able to distinguish situations in which one of them is true from at least some situations in which it is false just in case we are able to distinguish situations in which the other one is true from at least some situations in which it is false. So, it seems that those who take Dretske-style counterexamples as genuine counterexamples cannot drop the Equivalence Principle without further ado—that is, not for the reasons Dretske provides.

[60] Hawthorne (2005), p. 32.

It is noteworthy that among theorists who accept a closure principle for knowledge, there are diverging views as to how to best state the principle in detail. The version I favour is akin to Williamson's version of the closure principle, which he aptly labels 'intuitive closure'.[61] John Hawthorne has suggested improving on the present version of the closure principle in two ways: First, by distinguishing between two principles—single-premise closure, that is, a closure principle for deductions with only one premise, on the one hand, and multi-premise closure, that is, a closure principle for deductions with more than one premise, on the other. However, the distinction appears to be of dubious value. If we don't get multi-premise closure, then we won't be able to explain how so simple deductions as Kim's can result in knowledge. After all, Kim's deduction rests on (at least) four premises: (1) that she has three aunts, (2) that she has two uncles, (3) that all of them belong to the family of her mother, and (4) that they are the only aunts and uncles she has that belong to her mother's family.[62] Any account of closure that leaves knowledge acquired via deductions as basic as Kim's unexplained will be less than fully satisfactory. For that reason we cannot settle for anything less than multi-premise closure.

The second improvement Hawthorne suggests is a refinement of the present version of the principle. Here is the one he offers for multi-premise closure:

---

[61] Cf. Williamson (2000), p. 117.

[62] One might try to respond on Hawthorne's behalf that in the present case all one needs is single premise closure: Kim knows the proposition that consists of the conjunction of (1) to (4) and deduces her conclusion from it. However, it would seem that whether or not Kim deduces her conclusion from the proposition that consists of the conjunction of (1) to (4) is a matter of psychological fact. For instance, suppose Kim reasons in the following way: I have three aunts that belong to my mother's family. Therefore, my mother has three sisters. I have two uncles that belong to my mother's family. Therefore, my mother has two brothers. I have no other aunts and uncles that belong to my mother's family. Therefore, my mother does not have any other siblings. If my mother has three sisters and two brothers and no other siblings, then she has more sisters than brothers. So, my mother has more sisters than brothers. Certainly, our intuitions that Kim knows that her mother has more sisters than brothers lingers even if we suppose that she figured it out in this way. However, if she does so figure it out, then single premise closure won't be strong enough to explain that she knows. After all, she needs the conclusions of her earlier deductions in order to make the step to the last conclusion. That is to say, however, that in order to be able to explain that Kim knows the conclusion of her deduction we need multi-premise closure.

*Multi-Premise Closure* (MPC). Necessarily, if S knows that $p_1$ … $p_n$, competently deduces $q$, and thereby comes to believe $q$, *while retaining knowledge of $p_1$ … $p_n$ throughout*, then S knows $q$.[63]

In addition to the present version of the closure principle, Hawthorne claims that in order to acquire closure-based knowledge one must retain one's knowledge of the premises throughout. It is not clear, however, that this way of weakening the closure principle does any good. Suppose I suspect that a certain proposition I don't yet know follows from other propositions I know. I want to find out. I take out my paper and pencil, introduce letters for the propositions known and for proposition I suspect to follow from them, specify the logical relations I know to hold between them in the obvious ways and manipulate them in accordance with the rules of my preferred logic. A long but competent deduction shows that the proposition I suspected to follow does indeed follow from the propositions I know. Intuitively, I now know the proposition I antecedently suspected to follow from my premises. But now suppose that while I was going through the deduction, an evil demon destroyed my memory of one or more of the premises as well as my memory of what letter I introduced for which proposition and the key I had written down at the outset that could remind me of all these things. It is as implausible, I submit, that in this case I do not know the conclusion of my deduction as it is implausible that I should lose my knowledge of a conclusion of a deduction I have once gone through when I forget one of the premises. However, I do not satisfy all the conditions required by MPC for deduction-based knowledge. Hence, MPC does not predict that I know the conclusion of my deduction in the present case. Hawthorne's version of the closure principle is unnecessarily weak. For that reason I suggest hanging on to the closure platitude as initially formulated—the one akin to Williamson's 'intuitive closure' principle.

---

[63] Hawthorne (2004), p. 33, *my italics*.

## 5  The Anti-Luck Platitude

Knowledge excludes luck. Of course, this platitude is, at least in its present form, too general to be ultimately plausible. Some refinement is in order. However, progress can be made rather quickly here. First, notice that by the belief and factivity platitudes knowledge is a subspecies of true belief. Since a true belief can be lucky in almost every conceivable fashion—or, in other words, since there is virtually no sense in which true belief excludes luck—it is reasonable to suppose that the anti-luck platitude states one of the essential differences between true belief and knowledge. That is to say that part of what distinguishes true belief from knowledge is that the latter excludes luck in a sense in which the former doesn't. Accordingly, we may expect the following to be a more refined version of the initial anti-luck platitude:

Anti-Luck Platitude*: Knowledge is luck-free true belief.

Alternatively, if one adopts the convention of calling luck insofar as it may afflict true belief *epistemic luck* one can also state the anti-luck platitude in the following way:

Anti-Luck Platitude**: Knowledge excludes epistemic luck.

For present purposes, I will treat the two refined versions of the anti-luck platitude as equivalent. Furthermore, I will henceforth use the term 'anti-luck platitude' to refer to the refined version as captured in Anti-Luck Platitude* and Anti-Luck Platitude** respectively. With this quick development of the anti-luck platitude in place, I will now try to adduce some considerations in support of it. Again, the strategy is to unearth facts about our use of the word 'know' in ordinary language and thought that point to an implicit commitment to the anti-luck platitude.

To begin with, consider cases of unreliably formed belief. For instance, take the case of someone who superstitiously believes that she can predict the future by looking

at a crystal ball. Intuitively, such a person does not come to know what she sees (or thinks she sees) in the ball even if the belief she acquires upon looking at the ball turns out to be true. Alternatively, people who form beliefs on the basis of wild guesses[64] do not know either—again, no matter whether their beliefs turn out to be true. Our intuition here is, of course, reflected in our ordinary use of the word 'know' in that we are unwilling to apply 'know' to guesses or beliefs based on superstition. At the same time, it seems that the subjects' beliefs in cases like these ones, even if true, are only luckily true. After all, the guesser might easily have guessed a falsehood and the superstitious believer might easily have 'seen' a falsehood in the crystal ball. The anti-luck platitude provides a simple and elegant explanation of our intuition that beliefs formed in unreliably ways—such as beliefs formed on the basis of guessing and superstition—do not qualify as knowledge: They are just too lucky to count as knowledge. In this way, cases of unreliably formed belief provide evidence for the anti-luck platitude.

As I have already pointed out in the last chapter, one of the strongest intuitions we have about knowledge is that one does not know, in advance, that a given ticket in a fair lottery won't win. Again, this intuition is reflected in a feature of our ordinary usage of the word 'know'—*viz.*, that we don't apply the word 'know' to people who believe of a given ticket in a fair lottery that it won't win. Now notice that a fair lottery draw is a paradigm case of an event that is influenced by luck. There is a certain number of possible outcomes and each possible outcome might become the actual outcome. Which one it will be is down to luck. Now, given that the outcome of a fair lottery draw is an event that is down to luck, if one forms a belief, in advance of the drawing, that a given ticket in such a lottery won't win, then whether or not one's belief turns out to be true

---

[64] For brevity's sake I will henceforth use the term 'guess' instead of 'belief formed on the basis of guessing'. It is important to keep in mind that the former is just shorthand for the latter. For it might be thought that a mere guess normally excludes belief (cf. the discussion of Sartwell in ch. II, section 1.1). So, one might think that the intuition that guesses do not qualify as knowledge is grounded in a failure of the belief condition on knowledge. If we keep in mind that the intended sense of 'guess' is 'belief formed on the basis of guessing' one won't be led astray by this kind of thought.

will also be an event that is down to luck. Thus, if one forms a belief, in advance of the drawing, that a given ticket in a fair lottery won't win, one's belief will not be luck-free—even if it turns out to be true. Or, in other words, epistemic luck will not be excluded from one's true belief. Thus, if knowledge excludes epistemic luck, then one's belief that a given ticket in a fair lottery won't win will not qualify as knowledge. The anti-luck platitude thus also provides a simple and elegant explanation of our intuition that we don't know, in advance, that given tickets in fair lotteries won't win. Lottery cases, then, provide us with further evidence for the anti-luck platitude.

Finally, another type of case from which support for the anti-luck platitude can be extracted comprises the various Gettier cases. Again, there is a strong intuition that subjects in Gettier cases do not know what they justifiably and truly believe and a corresponding unwillingness to apply the word 'know' to subjects in Gettier cases. So, let us now ask whether there is reason to believe that Gettiered subjects' ignorance is grounded in a violation of the anti-luck platitude. To begin with, it is noteworthy that there is a widespread perception that the reason why subjects in Gettier cases lack knowledge is that their belief, although justified and true, is too lucky to qualify as knowledge. Here, for instance, is Jonathan Dancy on what is going on in Gettier cases:

> [K]nowledge must somehow not depend on coincidence or luck. This was just the point of Gettier counterexamples; nothing in the tripartite definition excluded knowledge by luck.[65]

Since it is imprudent to rest the case simply on a widespread perception, however, it will be advisable to look for further arguments to buttress the perception. One such argument can be found in Linda Zagzebski's diagnosis of how Gettier cases work. Zagzebski argues that Gettier cases work by cancelling out a bit of bad luck by a bit of

---

[65] Dancy (1985), p. 134.

good luck.[66] The a bit of bad luck is such that, in the circumstances, the subject would normally form a justified but false belief. However, the bad luck is cancelled out by some good luck such that the subject's justified belief is true after all. It is the combination of these two features that drives our intuition in Gettier cases. By way of illustration, take the case of Henry in Barn Façade County. He has bad luck in that he is in an environment in which barn façades prevail. In his circumstances, Henry will normally look at a barn façade and form a false belief that he is looking at a barn. However, Henry's bad luck is counteracted by some good luck—*viz.*, that although barn façades prevail, there is one real barn in the field and that is the one Henry is looking at. According to Zagzebski, our intuition that Henry does not know is driven by these two features of the case. Zagzebski's diagnosis is further supported by the observation that one can construct new Gettier cases simply by combining these two features. If Zagzebski's diagnosis as to how Gettier cases work is correct, however, then, of course, it must only be expected that the crux with Gettier cases is grounded in the fact that luck is involved in such a way as to exclude knowledge.

Apart from these considerations, there are, furthermore, some striking parallels between lottery and Gettier cases. Recall that in lottery cases there are a certain number of possible outcomes each of which might become the actual one. Which one it will be is down to luck. Accordingly, when one forms a belief, in advance of the drawing, that a given ticket won't win, whether or not one's belief turns out to be true is down to luck. Now, consider the belief-forming situation of a subject in a Gettier case such as, for instance, Henry's. An appealing way to describe Henry's belief-forming situation is that there are a couple of (relevant) cases that might obtain: Henry might be looking at a barn façade or at a real barn. Since in his present environment both barn façades and real barns are present each of these cases might be the actual one. Which one it will be

---

[66] Cf. Zagzabski (1994), p. 66.

is down to luck. At the same time, Henry will form a belief that he is looking at a barn—whether he is in fact looking at a barn façade or whether he is looking at a real barn. Accordingly, whether or not he ends up with a true or a false belief is down to luck. Just as in lottery cases, the anti-luck platitude will provide a simple and elegant explanation of why Henry's belief that he is looking at a barn does not qualify as knowledge. Just as in lottery cases, then, there is evidence that our intuitions in Gettier cases point to an implicit commitment to the anti-luck platitude.[67]

The anti-luck platitude provides a simple and elegant explanation of our intuitions about knowledge and our ordinary language use of 'know' in the each of three types of case—that is, first, cases of unreliably formed belief, second, lottery cases, and, third, Gettier cases. By the same token, these types of case provide some reason to believe that our intuition that subjects in such cases lack knowledge and our corresponding unwillingness to apply the word 'know' to them is grounded in a violation of the anti-luck platitude.

In conclusion, this chapter has outlined the details of a first version of a minimalist approach to the theory of knowledge. The crucial idea here was to pin down the concept of knowledge—albeit for the time being only partially—by a set of platitudes about knowledge. Seven platitudes have been defended: The factivity, belief, good informant, informative speech act, assertion, closure and anti-luck platitudes. It has been argued that each platitude is supported by facts about our use of the word 'know' in ordinary language and thought and by further theoretical argument where

---

[67] The proposal that the anti-luck condition is the anti-Gettier condition is not on par with the problematic proposals the post-Gettier literature witnesses. To see why this is so, notice that most advocates of the problematic proposals would accept that the anti-luck condition is the anti-Gettier condition. They consider themselves to be filling in the details of this rather uninformative way of stating the anti-Gettier condition. Notice also that, for present purposes, we need not be worried by the fact that the anti-Gettier condition proposed is quite uninformative. After all, we are trying to spell out the details of a *minimalist* approach to the theory of knowledge.

applicable. In the next chapter I will attempt to develop a more substantial condition on

knowledge from the minimalist framework here defended.

# IV   From Minimalism to a Safety-Based Conception of Knowledge

Recall that the aim of the last chapter was to develop a minimalist framework for the theory of knowledge. The central idea of a minimalist approach to epistemology was to devise platitudes about knowledge in order to locate the concept of knowledge in the conceptual landscape, as it were. In the last chapter, seven such platitudes have been defended, *viz.* platitudes relating knowledge and truth, belief, good informants, informative speech acts, assertions, competent deduction, and luck. These platitudes constitute the first version of a minimalist framework for the theory of knowledge.

In the present chapter I will discuss one promising proposal for a predicate satisfying the platitudes of the minimalist framework outlined in the last chapter. The predicate I have in mind is the predicate '*safely believes that p*'. Recall, however, that I emphasised that since I do not claim to have uncovered all the platitudes about knowledge, the minimalist framework will allow us to pin down the concept of knowledge *only partially*. In terms of the landscape simile, we have located the concept of knowledge in the conceptual landscape only roughly—we have circled an area, as it were, within which the concept of knowledge is to be found. In consequence, we found that there is no ultimate guarantee that if a given predicate satisfies the platitudes of the minimalist framework, then it will also be a knowledge predicate. It would seem, however, that even if the minimalist framework does not locate the concept of knowledge precisely, any predicate that does satisfy the platitudes has a fair shot at signifying at the very least a necessary condition for knowledge. Thus, if the predicate 'safely believes that *p*' does satisfy the platitudes, it does have a fair shot at signifying a property that is at the very least necessary for knowledge that *p*.

In fact, however, the argument I will present in this chapter is somewhat more ambitious. Rather than merely showing that the predicate 'safely believes that *p*' satisfies the platitudes, I will try to make a case for the thesis that the property of *safe belief that p*—that is, the property signified by the predicate 'safely believes that *p*'—is in fact requisite for knowledge that *p*. In order to do so, I will provide a couple of arguments in support of this thesis. The first one, which is due to Pritchard, combines one of the platitudes of the minimalist framework—*viz.* the anti-luck platitude—with some independently plausible further assumptions to derive what we may call the 'safety condition' on knowledge. As opposed to that, the second argument shows that Craig's method of objectivisation of the concept of good informant can be used to make the same point. If the first argument is sound, it will, of course, also have been shown that the predicate 'safely believes that *p*', that is, the predicate signifying the property of safe belief that *p,* is consistent with the anti-luck platitude. All that needs to be done to complete the argument that safe belief that *p* is necessary for knowledge that *p* is to show that the predicate 'safely believes that *p*' does not conflict with any of the other platitudes of the minimalist framework. Let me begin, however, by shedding some light on the notion of safe belief that figures so centrally in the argument of the present chapter.

## 1  Safety: The Core Thesis

The thesis that knowledge requires safe belief has enjoyed a significant degree of attention in recent epistemology and has been at the heart of one of the most promising responses to scepticism—*viz.* the neo-Moorean response[1]. It has been defended by a

---

[1] The neo-Moorean response to scepticism takes up an idea from G.E. Moore (1925), (1939)—*viz.*, that one can come to know the denials of sceptical hypotheses by deducing them from known ordinary empirical propositions. Safety plays such a central part in the neo-Moorean response to scepticism

number of renowned philosophers such as Sosa, Williamson and Pritchard. One intuition that is behind all conceptions of knowledge that place a safety condition on knowledge is that in order to know one must have a true belief that is not in danger of—or, in other words, that is *safe* from—error. The notion of safety here is modal in nature. What I will argue is that the most prominent characterisations of safe belief that can be found in the literature can be reconstructed as maintaining that in order to believe safely, one must not only believe truly in the actual world but also avoid false belief across a range of nearby possible worlds.

Consider, to begin with, Pritchard's characterisation of safety. Pritchard glosses the notion of safe belief explicitly in the terminology of nearby possible worlds:

(P-Safety)     S's belief is safe *iff* in [all[2], CK] near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true.[3]

According to Pritchard's version of the safety principle, in order to believe a proposition safely one must believe it truly at all nearby worlds at which one forms one's belief in the same way as in the actual world. Or, equivalently, in order to believe a proposition safely one must avoid believing it falsely at all nearby worlds at which one forms one's belief in the same way as in the actual world.

As opposed to that, Sosa's characterisation of safety seems to be slightly different than Pritchard's. He offers the following safety principle:

---

because it allows neo-Mooreans to argue that our beliefs in the denials of sceptical hypotheses can enjoy some positive epistemic support: Our beliefs in the denials of sceptical hypotheses can be—and provided the world cooperates will be—safe. Most prominently Sosa (1999) and Pritchard (2005) motivate a neo-Moorean response to scepticism by appeal to a safety-based conception of knowledge.

[2] Pritchard's official version says "nearly all (if not all)" at this point. I am inclined to favour the stronger of the two alternatives, according to which one must believe truly at *all* nearby worlds with the relevant feature, since only if so construed there is a clear argument that safe belief is truth-entailing (cf. section 4.2). Since, arguably, in order to have the anti-Gettier import we may expect it to have (cf. chapter V, section 1) the safety condition on knowledge—knowledge entails safe belief—must be also factive (cf. chapter V, section 3.2), and since whether safety is construed in the stronger or the slightly weaker way is otherwise of little importance for present purposes, in order to keep things neat and simple I will construe Pritchard's version of the safety principle in the stronger fashion outlined above.

[3] Pritchard (forthcoming*a*), p. 6.

(S-Safety)    [A] belief by S is "safe" iff: … not easily would S believe that p without it being the case that p.[4]

However, given a standard possible worlds semantics of the notion of easy possibility, according to which A could easily be true iff A is true at some nearby possible world, Sosa's claim is tantamount to the claim that S's belief in a proposition is safe if and only if S avoids believing it falsely at nearby possible worlds.[5,6]

In this way, it turns out that Sosa and Pritchard's respective formulations of the safety principle are very similar. There is, however, at least one noticeable difference between Pritchard and Sosa's version. While, according to Sosa's version of the safety principle, in order to believe a proposition safely one must avoid believing it falsely at *all* nearby possible worlds, according to Pritchard's version, in order to believe a proposition safely one must avoid believing it falsely only across a restricted range of nearby possible worlds—*viz.* the ones at which one comes by one's belief *in the same way* as in the actual world. Although the difference between the two proposals may appear to be only slight, one may still wonder which of the two versions of the safety principle is preferable.

There is excellent reason to believe that Pritchard's version of the safety principle has an edge over Sosa's. The reason why this is so is that Sosa's version of the safety principle—or any version of the safety principle that does not place further restrictions on the range of nearby possible worlds across which error must be avoided—yields a safety condition on knowledge that is bound to be too strong. The following well-known case due to Robert Nozick brings the point home. Suppose granny is visited by her grandson and comes to believe by looking at him that he is well.

---

[4] Sosa (1999), p. 142.
[5] I take it that Sosa's 'not easily would S believe that *p* without it being the case that *p*' is equivalent to 'it could not easily be the case that S believes that *p* and *p* is false'.
[6] It is noteworthy that Sosa also provides a slightly different statement of the safety principle according to which S's belief that *p* is safe iff were S to believe that *p*, *p* would be true. I take it that safety is better construed by appeal to the notion of easy possibility than in terms of a counterfactual conditional—for one argument to that effect see Pritchard (2005), pp. 72-3.

Granny is good at telling these things by looking. Intuitively, when she forms a belief that her grandson is well by looking, her belief qualifies as knowledge. Now suppose that at some nearby possible worlds her grandson is ill. In order to save granny from distress, at (some of) these worlds her family tells her that her grandson is well but had something important to do and for that reason couldn't come and visit. Granny forms a false belief at these nearby possible worlds.[7] Any conception of the safety principle that construes safe belief as requiring avoidance of error at nearby possible worlds without further restrictions predicts that granny's true belief that her grandson is well, acquired in the actual world by looking, is unsafe. In consequence, any theory of knowledge that makes safe belief, so construed, necessary for knowledge predicts, counterintuitively, that granny does not know that her grandson is well when she looks at him and on that basis forms a true belief to that effect. Indexing to ways of belief-formation will remedy this defect. After all, at those nearby possible worlds at which granny forms a false belief she comes by her belief in a different way than in the actual world. She relies on testimony rather than on looking. So, there is reason to believe that Pritchard's version of the safety principle is preferable to Sosa's because the latter yields a safety condition on knowledge that is too strong.[8]

Unfortunately, however, there is reason to believe that even Pritchard's version of the safety principle yields a safety condition on knowledge that is too strong. Again, the problem is that there are cases in which, intuitively, the subject knows a proposition while her belief does not satisfy the safety condition derived from Pritchard's version of the safety principle. One such case is a variation of what, presumably, was the first Gettier case (due, somewhat surprisingly, to Bertrand Russell). Suppose I wake up in the morning, come down the stairs, have a look at the grandfather clock that is standing

---

[7] Cf. Nozick (1981), p. 179.
[8] It is noteworthy that, in Sosa (2002), Sosa provides a more refined version of his safety principle that will also avoid the problem posed by this case (cf. also p. 104, n. 16).

in the hallway, see that it reads 8.22 and on that basis form a belief that it's 8.22. My belief is well justified: I know the clock to be highly reliable, have no reason to believe that it is not working properly etc. Moreover, my belief is true. It is in fact 8.22. However, here comes the catch, the clock has stopped working exactly twelve hours earlier. Intuitively, my belief acquired by reading the clock does not qualify as knowledge. It is just too lucky to do so. After all, I could easily have come down the stairs a minute earlier or later in which case I would have acquired a false belief.[9]

Notice that in the present version of the case, my belief that it is 8.22 does not satisfy the safety condition derived from Pritchard's version of the safety principle. After all, given a standard possible worlds semantics of the relevant modal notions, saying that I could easily have come down the stairs a minute earlier or later is tantamount to saying that at a nearby possible world I come down a minute earlier or later. If at such a world I acquire a belief by reading the stopped clock, I will form a false belief. A theory that makes Pritchard-style safety requisite for knowledge will predict, correctly, that I do not know that it's 8.22.

But now consider the following variation of the case. Suppose my arch-nemesis has an interest that I form a belief that it's 8.22 when I come down the stairs. He is prepared to do whatever it may take to ensure that I acquire a belief that it's 8.22 when I come down the stairs. (We may suppose that he is prepared to set the clock to 8.22 if I chance to come down at another time.) However, my arch-nemesis is also lazy. He will act only if I do not come down the stairs at 8.22 of my own accord. Suppose, as it so happens, I do come down the stairs at 8.22. My arch-nemesis remains inactive. I form a belief that it's 8.22. It is 8.22. The grandfather clock is working reliably as always. Intuitively, I take it, I know that it's 8.22 upon reading the clock. After all, I look at a perfectly working clock, I have the ability to read the clock, exercise it and, as a result,

---

[9] Cf. Russell (1948), pp. 170-1.

form a true belief about the time. However, my belief that it's 8.22 does not satisfy Pritchard's safety principle. At all nearby possible worlds at which I come down a minute earlier or later my arch-nemesis steps on the scene and sets the clock to 8.22 anyway. At those worlds, I form a false belief that it's 8.22. At the same time, I form my belief in the same way as in the actual world—by reading the clock. If at all nearby possible worlds at which I come down a minute earlier or later I still form a belief that it's 8.22 in the same way, it is not the case that at all nearby possible worlds at which I form my belief in that way I avoid forming a false belief. In this variation of the case, then, I intuitively know that it's 8.22 even though my belief does not satisfy the safety condition derived from Pritchard's safety principle.[10]

This case suggests that the safety condition derived from Pritchard's safety principle is still too strong. If so, neither Sosa's nor Pritchard's version of the safety principle will yield a plausible safety condition on knowledge. However, there is a third version of the safety principle that can be found in the literature which is due to Williamson. Although Williamson uses the term 'reliable' and its cognates instead of 'safe' and its cognates, it is widely accepted that Williamson effectively states a safety condition on knowledge. Accordingly, I will describe Williamson's relevant condition on knowledge as a safety condition. Williamson characterises the relevant notion of safety in the following way:

> Reliability and unreliability, stability and instability, safety and danger, robustness and fragility are modal states. They concern what could easily have happened. They depend on what happens under slight variations in the initial conditions.[11]
>
> For present purposes [i.e. for the purposes of spelling out the notion of reliability that, according to Williamson, is necessary for knowledge, CK], we are interested in a notion of reliability on which, in given circumstances, something happens reliably if and only

---

[10] Similar cases can be found in Neta and Rohrbaugh (2004) as well as Comesaña (2005). It should be said, however, that I owe inspiration for this example (and the example in chapter V, section 3) to Harry Frankfurt who uses examples with a similar structure to argue against the view that moral responsibility requires ability to do otherwise (cf. Frankfurt (1969)).

[11] Williamson (2000), p. 123.

> if it is not in danger of not happening… In particular, one avoids false belief reliably in [a case, CK] α if and only if one avoids false belief in every case similar to α.[12]

The second part of the quotation states that in order to believe safely one must avoid false belief in similar cases. At the same time the first part indicates that cases are similar if they could easily have obtained. But now recall that, according to a standard possible worlds semantics of the notion of easy possibility, a statement of the form 'A could easily have obtained' is true if and only if A obtains at some nearby possible world. Putting these three things together we can reconstruct Williamson's safety principle in possible worlds terminology:

(W-Safety)     One believes safely if and only if one avoids false belief at nearby possible worlds.

Now one might think that in doing so we will have made little to no progress towards finding a version of the safety principle that avoids the problems that beset Sosa and Pritchard's versions of it. After all, the resulting safety condition on knowledge derived from it will face the same problem as the one derived from Sosa's version of the safety principle: If granny's grandson is ill at a nearby possible world while her family tells her that he's well granny will come by a false belief at some nearby possible world. All the while she may know by looking that her grandson is well in the actual world. However, rather than taking this to be an objection to a Williamson-style safety principle, we might think that we have not been careful enough in translating Williamson's word into possible worlds terminology. After all, Williamson would probably want to complain that a case in which granny is *told* about her grandson's state of health is in crucial respects *dissimilar* to a case in which she forms a belief about it *by looking*. Alternatively, the variation in the initial conditions between the two cases is more than *slight*. So, how can we improve our translation of Williamson's remarks into

---

[12] Williamson (2000), p. 124.

possible worlds terminology so as to do justice to Williamson's complaints? One way of improving our translation, it seems to me, is the following:

(W-Safety*)     One's belief that $p$ is safe if and only if one believes that $p$ in the actual world and one avoids false belief at nearby possible worlds at which the relevant initial conditions for the formation of the belief are the same as in the actual world.

It is plausible that holding the relevant initial conditions for the formation of the belief fixed will ensure that no more than slight variations in the initial conditions are admissible. At the same time, it is also plausible that some variations will be admissible. After all, it is only the *relevant* initial conditions, rather than *all* initial conditions, which are held fixed.[13]

Let us then ask whether the safety condition on knowledge derived from W-Safety* remedies the defects of the ones derived from Pritchard and Sosa's versions of the safety principle. To begin with there is reason to believe that it avoids the problem posed by the grandmother case. After all, it is plausible that it is part of the relevant initial conditions for the formation of granny's belief that she forms her belief by looking. Accordingly, granny's belief that her grandson is well formed *by looking* may satisfy W-Safety* even if at some nearby possible world her grandson is unwell while she continues to believe that he is well on the basis of *testimony* from her family. At those nearby worlds the relevant initial conditions for the formation of the belief have changed.

But how about the case in which I form a belief that it's 8.22 via reading a perfectly functioning clock while my arch-nemesis ensures that at nearby possible worlds I form a false belief that it's 8.22? Does my belief satisfy W-Safety*? One might think that it doesn't. After all, it seems only plausible that it is consistent with holding

---

[13] It is noteworthy that the notion of relevant initial conditions at issue here is notoriously vague. Even so, I agree with Pritchard ((forthcoming), p. 3) that we still have a pretty good intuitive grasp as to what has to be held fixed as part of the relevant initial conditions (cf. section 2.2).

the relevant initial conditions for the formation of the belief fixed that at some nearby possible worlds, I come down a minute earlier or later in which case I form a false belief that it's 8.22. If so, the safety condition derived from W-Safety* faces the same problem as the safety conditions derived from Sosa and Pritchard's versions of the safety principle. However, there may be some room for the defender of W-Safety* here. Again, it is Williamson himself who points to the way of solving the problem. As Williamson points out: "Safety and danger are highly contingent and *temporary* matters."[14] For instance, a runaway driver may be safe once he has taken the turn onto the motorway even though he might easily have taken the wrong turn in which case he would have been caught. If safety and danger are in general highly temporary matters, then so they will be in the case of belief-formation. We can accommodate this idea in the application of W-Safety* by allowing that the facts that need to be held fixed as part of the relevant initial conditions may change over time. If so, then it can be argued that my belief that it's 8.22 satisfies W-Safety* by maintaining that once I come down the stairs at 8.22 it becomes part of the relevant initial conditions that my arch-nemesis remains inactive. If so, since, additionally, I acquire my belief via a capacity to read the clock and the clock is working properly, there will be no nearby possible worlds at which I form a false belief. Thus my belief satisfies W-Safety*.[15]

---

[14] Williamson (2000), p. 124, *my italics*.

[15] One might be worried that if we make room for updating the relevant initial conditions in the way envisaged, then we will no longer get the right predictions in some other Gettier cases. Thus, consider, for instance, the original version of the clock case. If it becomes part of the relevant initial conditions that I come down at 8.22 when I do so and if it is also part of the relevant initial conditions that the clock has stopped at 8.22, then my belief that it's 8.22 will be safe. Alternatively, consider the case of Henry in Barn Façade County. If we allow for updating of the relevant initial conditions, it is no longer clear that when Henry looks at the only real barn in a field full of barn façades and thereupon forms a belief that he is looking at a barn, his belief will be unsafe. After all, if it becomes part of the relevant initial conditions that Henry looks out of the window at exactly the point in time at which he does, then it would seem that his belief turns out to be safe.

In response to this worry notice, first, that for now I am only interested in an argument suggesting that safety is a *necessary* condition for knowledge. For that reason, at this stage of the argument there would be no special problem if my belief in the original version of the clock-case or Henry's belief in the barn façade case turned out to be safe. However, second, the defender of W-Safety* may point out that there is a difference between the problematic variation of the clock-case on the one hand, and the original version of it and the barn façade case on the other. In the original version of the

Notice that while this move is available to the defender of W-Safety*, a parallel move is not open to the defender of P-Safety. In order to get the parallel move off the ground, the defender of P-Safety would have to show that the fact that the arch-nemesis remains inactive at the time of belief-formation enters into the description of the way in which the belief is formed. However, it would seem that there is no plausible way of individuating ways of belief-formation in which part of a way of belief-formation is that some other person refrains from doing a certain thing. In consequence, the defender of P-Safety cannot solve the problem posed by the present case in the same way as the defender of W-Safety*.[16]

If W-Safety* is an acceptable version of the safety principle, and if knowledge requires safe belief, then we can derive the following safety condition on knowledge:

(Safety)     One knows that *p* only if in the actual world one believes that *p* and one avoids false belief at nearby possible worlds at which the relevant initial conditions for the formation of the belief are the same as in the actual world.

---

clock-case the danger of error is generated by the fact that the clock I read is stopped. In the barn façade case the danger of error is generated by the fact that Henry is otherwise surrounded by barn façades. Notice that, at the time of belief-formation, the facts that generate the danger of error are still facts about the subjects' respective environments. That is, it is still the case that the clock I am reading is stopped and it is still the case that Henry is otherwise surrounded by barn façades. As opposed to that, in the problematic variation of the clock-case, the danger of error is generated by the fact that my arch-nemesis is in the background prepared to manipulate the clock. But now notice that, at the time of belief-formation, it is no longer the case that my arch-nemesis is prepared to manipulate the clock. That is to say, however, that the fact that earlier on generated the danger of error has, at the time of belief-formation, ceased to be. (If this formulation does not mesh with the reader's preferred theory of time, I enjoin the reader to make the obvious changes.) The defender of W-Safety* can use this difference between these cases in order to motivate the following restriction on updating of the relevant initial conditions: One may update the relevant initial conditions only if the fact that generated the danger of error in the first place has, at the time of belief-formation, ceased to be. To provide further support for this restriction consider a variation of the barn façade case in which a benevolent demon will replace all barn façades in the field alongside which Henry is driving by real barns just before Henry looks out of the window. In this variation of the case, at the time of belief-formation, the fact that generated the danger of error has ceased to be. Updating may be allowed. Notice, however, that it seems that in this variation of the case, updating gives the right predictions. After all, it would seem that when looking at a field full of barns, Henry's belief that he is facing a barn does qualify as knowledge. At the same time, if we update the relevant initial conditions, we get the result that Henry's belief that he is facing a barn is safe.

[16] Notice that W-Safety* closely resembles a later version of safety due to Sosa (2002) according to which there may be cases of *dependently* safe belief. According to Sosa, one's belief is dependently safe iff it is safe provided that a certain condition obtains (cf. Sosa (2002), pp. 268, 275-6). If one interprets that condition as the conjunction of the facts that belong to the relevant initial conditions, one will also get W-Safety*. It is noteworthy that Sosa defends the additional thesis that in order to know one must be *guided* by the facts dependent on which one's belief is safe (cf. Sosa (2002), pp. 275-6). For present purposes I will leave it open whether we have to construe knowledge such that the additional thesis holds as well.

## 2 The Argument from the Anti-Luck Platitude

In this section I will rehearse an argument, originally due to Pritchard, that the safety condition on knowledge is entailed by a plausible development of the anti-luck platitude and an independently plausible modal conception of luck. In what follows, I will first outline the development of the anti-luck platitude, then move on to the credentials of and the main motivations for Pritchard's modal conception of luck and, finally, show how the argument for the safety condition can be run on the basis of this conception of luck and the developed anti-luck platitude.

### 2.1 The Anti-Luck Platitude Developed

### 2.1.1 Varieties of Epistemic Luck

In the last chapter we saw that there is reason to believe that the anti-luck platitude holds—that is, that knowledge excludes epistemic luck. While the considerations adduced in the last chapter provide the relevant support for the anti-luck platitude, there is reason to believe that more can be done in order to bring out in what way exactly knowledge excludes epistemic luck. Recall that 'epistemic luck' was defined as luck that may afflict true belief. Although cases of unreliably formed belief, lottery cases, and Gettier cases are all cases in which luck afflicts true belief in such a way that, intuitively, the subjects lack knowledge, as we will see in a moment, there are also other ways in which luck may afflict true belief, ways in which one's true belief may be lucky, whilst, intuitively, one still counts as knowing what one truly believes. There are, then, types of epistemic luck that are compatible with knowledge. Or, to put it in Pritchard's words, there are "benign varieties of epistemic luck".[17] Thus, in order to

---

[17] Cf. Pritchard (2005), p. 133.

bring out the precise sense in which knowledge excludes luck, it will be instructive to distinguish the type or types of epistemic luck that are inconsistent with knowledge—henceforth also 'malignant types of epistemic luck'—from benign varieties of epistemic luck.

Although there has been some discussion of benign versus malignant varieties of epistemic luck in the literature dating back to the late sixties (most notably in Unger (1968), Statman (1991), and Engel (1992)), the treatment of the issue has been rather fragmentary.[18] The first more detailed taxonomy of benign versus malignant varieties of epistemic luck can be found in Harper (1996). The most comprehensive treatment of epistemic luck has recently been developed in Pritchard (2004*b*) and (2005). In order to save time and space, I will follow Pritchard's account of the varieties of epistemic luck rather than picking up the pieces from the earlier literature.

Pritchard distinguishes five different ways in which a subject's true belief can be afflicted by luck that are of interest for present purposes[19], that is, five varieties of epistemic luck. Here they are:

*Content epistemic luck*
It is lucky that the proposition is true.[20]

---

[18] One might object that Unger's treatment of varieties of epistemic luck is relatively complete. After all, Harper (1996, p. 279) claims nothing more than having extracted the varieties of epistemic luck he countenances from Unger's writings. However, it is not clear that Unger explicitly countenances all the relevant varieties of luck (cf. Pritchard (2005), p. 136). Moreover, Unger's discussion is actually stated in terms of the notion of accident rather than in terms of the notion of luck.

[19] In fact, he distinguishes six varieties of epistemic luck. The sixth type of epistemic luck, which he labels 'reflective epistemic luck', is characterised in the following way:

*Reflective epistemic luck*
Given only what the agent is able to know by reflection alone, it is a matter of luck that her belief is true. (Pritchard (2005), p. 175.)

Pritchard claims that while reflective epistemic luck is not strictly speaking incompatible with knowledge it is still epistemically problematic and for that reason stands in need of diagnosis (cf. Pritchard (2005), chs. 6-9). Since most of his discussion of reflective epistemic luck is not directly relevant for present purposes, whilst providing a proper outline of the issues associated with it would be rather space-consuming, I omit discussion of this type of luck, rest content with Pritchard's claim that reflective luck is compatible with knowledge, and refer the reader to Pritchard (2005) for further discussion.

[20] Pritchard (2005), p. 134.

*Capacity epistemic luck*
It is lucky that the agent is capable of knowledge.[21]

*Evidential epistemic luck*
It is lucky that the agent acquires the evidence that she has in favour of her belief.[22]

*Doxastic epistemic luck*
It is lucky that the agent believes the proposition.[23]

*Veritic epistemic luck*
It is a matter of luck that the agent's belief is true.[24]

I will outline Pritchard's argument (which is of the same form as Harper's, Unger's and Engel's arguments[25]) that the first four types of luck—content, capacity, evidential and doxastic luck—are benign types of epistemic luck, while the fifth type of luck—veritic epistemic luck—is incompatible with knowledge.

## 2.1.2   Benign Varieties of Epistemic Luck

Pritchard's et al. project of bringing out the exact way in which knowledge excludes luck starts out with an argument that content, capacity, evidential and doxastic epistemic luck are compatible with knowledge. The argument consists of a series of cases. In each case, intuitively, the protagonist knows a proposition, *P*, while, at the same time, her true belief that *p* is afflicted by a different type of these types of epistemic luck. In this way, the cases suggest that knowledge is compatible with these types of epistemic luck. Here, then, are the cases:

*Bumping into a friend (illustrating that knowledge is compatible with content epistemic luck)*
Suppose I have an old friend whom I have not seen in a long time. Suppose my friend has moved abroad a few years ago and rarely ever comes back home. However, just today he is stopping over in the city I have recently moved to. While he is on the way from the airport to his hotel, I bump into him in the streets. Realising that this is what has happened I form a

---

[21] Pritchard (2005), p. 134.
[22] Ibid., p. 136.
[23] Ibid. p. 138.
[24] Ibid. p. 146. Harper also distinguishes these five types of epistemic luck (Harper (1996), p. 279).
[25] Notice that Engel only discusses veritic and evidential luck.

true belief that I have bumped into my friend. Intuitively, my belief qualifies as knowledge. However, my true belief is also afflicted by content epistemic luck. After all, my bumping into my friend is a lucky event if anything is. Accordingly, it also is lucky that the proposition that I have bumped into my friend is true. So, knowledge is compatible with content luck.[26]

*Disease of the eye (illustrating that knowledge is compatible with capacity epistemic luck)*
Suppose I suffer from a rare disease of the eye: I have a tendency not to be able to see whenever I am exposed to bright light. Suppose I step out onto my friend's balcony. It's a very bright day, so bright indeed that my disease would normally cause me to be temporarily blind when I am outside. Luckily, however, my disease does not prevent me from seeing on this particular occasion. I look around, see that the cherry trees are blossoming and form a true belief to that effect. Intuitively, my true belief qualifies as knowledge. However, my true belief is also afflicted by capacity epistemic luck. After all, given my disease, it is a matter of luck that I am able to see in the bright sunlight to which I am exposed on my friend's balcony. So, knowledge is compatible with capacity epistemic luck.[27]

*The fired employee (illustrating that knowledge is compatible with evidential epistemic luck)*
Suppose Hynes walks past his employer's office and overhears his employer saying that he (Hynes) is about to be fired. Intuitively, when he comes to truly believe on the basis of what his employer says that he will soon be fired, his belief qualifies as knowledge. At the same time, it may be a matter of luck that he comes across the evidence for his belief: It may be a matter of luck that he walks past his employer's office. He is not normally around this part of the building—only today he had to deliver a report to someone over here. If it is a matter of luck that Hynes walks past his employer's office, then it is also a matter of luck that he comes by the evidence for his belief that he is about to be fired. If so, his belief is afflicted by evidential luck. So, knowledge is compatible with evidential epistemic luck.[28]

---

[26] Unger ((1968), p. 159) and Pritchard ((2005), p. 134) both use the example of a car accident to illustrate the compatibility of knowledge and content luck. Even though I do not want to dispute that the example works, I am not certain that it is the best example available. After all, car accidents are often caused by irresponsible—drunk, reckless etc.—driving. And it is not at all clear that it in such cases it is a matter of luck that the accident occurs. (Notice that this is so even by the lights of Pritchard's account of luck (Pritchard (2005), p. 128). Given that someone is drunk it may not be the case that the accident does not occur at most nearby worlds. In that case, by Pritchard's account of luck, the event of the accident fails to satisfy a necessary condition for being a lucky event.) Accordingly, the event of a car accident might give rise to misunderstandings here. The event of bumping into somebody seems to be less prone to do so.
[27] A similar example can be found in Pritchard (2005), p. 135-6.
[28] The first one to give the example is again Unger ((1968), p. 159). However, as Pritchard points out, Unger does not identify evidential luck as an independent, benign type of epistemic luck (cf. Pritchard (2005), pp. 135-6). Rather, he appears to be subsuming it under capacity luck. The first one to be explicit about evidential luck is Mylan Engel ((1992), p. 67) who takes evidential luck to be a benign variety of epistemic luck and contrasts it with veritic luck which he takes to be malignant. However, Engel only contrasts evidential and veritic luck and seems to be unaware of Unger's treatment of the other benign varieties of luck. Both Harper ((1996), pp. 159-60) and Pritchard ((2005), pp. 134-41) connect both discussions. Statman (1991) approaches the issue of epistemic luck from a slightly different angle. He construes varieties of epistemic luck in a parallel fashion to Thomas Nagel's varieties of moral luck (cf. Nagel (1979*a*)). In this way he develops in an independent fashion analogues of what Pritchard calls 'capacity epistemic luck' and 'veritic epistemic luck' (roughly Statman's 'constitutive luck' and 'resultant luck') ((1991), pp. 149-51). However, Statman's taxonomy does not fully parallel the above one. In particular, what Statman calls 'circumstantial' epistemic luck ((1991), p. 147), is a mixture of capacity luck and veritic luck. It is not clear that the category as construed by Statman really captures an independent type of epistemic luck.

*The sceptical detective (illustrating that knowledge is compatible with doxastic epistemic luck)*

Suppose Detective Green is trying to solve the mystery of the murder of Hynes's employer. He comes to believe, truly, that it was Hynes on the basis of forensic evidence that establishes beyond reasonable doubt that it was Hynes. Intuitively, Green knows that it was Hynes. But now suppose Detective Green is a sceptical person. In order to form a belief that a suspect is a murderer, Green normally requires much more evidence than would otherwise be needed in order to know. For instance, we may suppose that, normally, Green would not form a belief that a suspect did the deed unless he had evidence that the suspect had a sound motive for committing the crime. While, in the present case, Green does not and would not have this evidence (all the relevant evidence has been destroyed and Hynes himself is unwilling to impart the information), the striking way in which the investigation has developed leads Green to form a belief that Hynes did the deed on the basis of the forensic evidence only. Now we need only suppose that it was a matter of luck that the investigation developed in the striking way it did. If so, then it is a matter of luck that Green formed a belief that it was Hynes. Green's belief is doxastically lucky. So, knowledge is compatible with doxastic luck.

What these cases suggest is that content-, capacity-, evidential- and doxastic epistemic luck are all benign varieties of epistemic luck—they are all compatible with knowledge.[29]

---

[29] It is noteworthy that while I agree with Pritchard that the abovementioned types of epistemic luck are benign, there is one difference between Pritchard and myself. Pritchard maintains that evidential luck and doxastic luck are, at least insofar as they are benign types of epistemic luck, roughly equivalent. That is to say that there are no *uncontroversial* cases—that is, according to Pritchard, cases without further controversial epistemic assumptions—in which a subject knows a proposition, whilst also being evidentially but not doxastically lucky, and no uncontroversial cases in which a subject knows a proposition, whilst also being doxastically but not evidentially lucky. (Cf. Pritchard (2005), p. 140. In the earlier article, Pritchard is even more sanguine about the equivalence of the two types of luck (cf. Pritchard (2004*b*), pp. 202-4).) However, the case of the sceptical detective provides us with at least *prima facie* evidence that Pritchard's conclusion here is mistaken. For, while Green is doxastically lucky—it is a matter of luck that the investigation developed in the way it did and Green would not have formed the relevant belief if it had not so developed—he is not evidentially lucky. After all, it may well be that even if the investigation had taken a different turn, Green, being the conscientious detective he is (and Hynes being the amateurish murderer he is), would have come across the relevant forensic evidence. If so, however, Green is not evidentially lucky. At the same time, on the face of it, no contentious epistemic assumptions have been invoked. The assumptions that are needed to get the case off the ground is that the investigation develops in a special way and that this makes a psychological difference to Green (in that it makes him form a belief in circumstances in which he would normally not form a belief). Green might be a non-standard believer—normal believers would already form a belief when they have sufficient forensic evidence—but such believers may well exist. Moreover, it would seem that when non-standard believers like Green form beliefs in the way that standard believers would form them, they may still acquire knowledge—only provided, of course, that their belief is true and that they base their belief on suitable evidence. But basing a true belief on suitable evidence Green does. So, we have a *prima facie* uncontroversial case in which a subject knows a proposition, whilst also being doxastically but not evidentially lucky.

The case of Detective Green thus provides us with *prima facie* evidence that doxastic luck and evidential luck are not equivalent, not even in the rough sense intended by Pritchard. It is a further question, however, whether they are fully independent kinds of epistemic luck. The case of Detective Green gives us only the first half of the argument that doxastic and evidential epistemic luck are fully independent kinds of luck. After all, every (uncontroversial) case of evidential luck might still be a case of doxastic luck. To complete the argument that doxastic and evidential epistemic luck come apart, we need an uncontroversial case in which a subject knows a proposition, whilst also being evidentially but

## 2.1.3 Veritic Epistemic Luck

In this section I will look at evidence suggesting that veritic epistemic luck is involved in all the types of case that motivated the anti-luck platitude. If so, then there is reason to believe that veritic epistemic luck is the type of luck at issue in the anti-luck platitude and that veritic epistemic luck is incompatible with knowledge. Let us go back to these cases, then. Recall that the first type of case comprised cases of unreliably formed belief. By way of example, we looked at a case in which a subject acquires a belief via guessing and another one in which a subject acquires a belief on the basis of what she sees (or thinks she sees) in a crystal ball. We found that, intuitively, beliefs acquired in these ways, even if true, do not qualify as knowledge. And the diagnosis we issued was that the subjects' beliefs are too lucky to qualify as knowledge. Now let's ask whether the subjects' beliefs are afflicted by veritic luck; or, in other words, whether it is a matter of luck that their beliefs are true. It is plausible that the answer to this question is 'yes'. After all, the problem in these cases is that there is nothing about the way of belief-formation involved that ensures or even makes it more likely than not that a belief formed in such a way is *true*. Subjects who form beliefs in unreliable ways, if

---

not doxastically lucky. However, it seems that such a case is not too hard to come by. Suppose that instead of Green, the detective who is working on the case of the murder of Hynes's employer is Detective Red. Red has accumulated quite a bit of evidence suggesting that it was Hynes who committed the murder. While Red's evidence makes it overwhelmingly reasonable to believe that it was Hynes, it isn't strong enough to turn a belief based on it into knowledge. Now, suppose that by luck Red stumbles across the bit of evidence that seals the case against Hynes. Intuitively, Red now knows that Hynes committed the crime. At the same time, Red is evidentially lucky. He was lucky to stumble across the missing bit of evidence. However, Red is not doxastically lucky. Being the reflective investigator he is, if Red hadn't stumbled across the missing evidence, he would have been aware of the fact that, given the evidence he has, it is overwhelmingly reasonable to believe that it was Hynes. On that basis he would have believed that it was Hynes anyway. Again, it seems to me that there aren't any contentious epistemic assumptions in play here. For instance, Detective Red is not to be blamed for his belief-forming practises and he is not ignoring any defeaters either. In this way we have reason to believe that we have an uncontroversial case in which, intuitively, a subject knows a proposition, whilst also being evidentially but not doxastically lucky. The argument that evidential and doxastic luck come apart can thus be completed.

they hit upon the truth, will be lucky to have done so. So, beliefs formed in unreliable ways are veritically lucky.[30]

Moreover, there is reason to believe that beliefs that given tickets in fair lotteries won't win are bound to be veritically lucky. Recall that we said that (fair) lottery draws are the paradigmatic cases of lucky events in that each ticket might win and which one it will be is down to luck. Given that the outcome of a fair lottery is down to luck, if, in advance of the drawing, one forms a belief that a given ticket in such a lottery won't win, then whether or not one's belief *turns out to be true* will be down to luck. That is to say, however, that one's belief, even if true, will be afflicted by veritic luck. There is, therefore, also reason to believe that beliefs that given tickets in fair lotteries won't win are veritically lucky.

Finally, some of the considerations that suggested that the problem with Gettier cases is that Gettiered subjects are too lucky to qualify as knowing can be used to argue that the beliefs of Gettiered subjects are afflicted by veritic luck. First, recall that, according to Zagzebski's diagnosis, Gettier cases are cases in which an element of bad luck is cancelled out by an element of good luck. Recall, furthermore, that we found it plausible that the element of bad luck is such that, in the Gettiered subject's circumstances, she would normally form a *false* belief, while the element of good is such that she ends up with a *true* belief after all. Notice that if the mixture of good and bad luck that drives our intuitions in Gettier cases works in this way, then it is also plausible that the Gettiered subject ends up with a belief that is only luckily *true*. That is to say, however, that the beliefs of Gettiered subjects are veritically lucky.[31]

Second, the parallels I have pointed out between lottery and Gettier cases can be used to make a case that the argument that lottery cases are cases of veritically lucky

---

[30] Pritchard's also adduces cases of unreliably formed belief in which the subject is veritically lucky (cf. Pritchard (2005), pp. 146-7).
[31] Cf. also Pritchard (2005), pp. 149-50.

belief extends to Gettier cases. To see how this works, consider, again, the case of Henry in Barn Façade County. Henry's case parallels lottery cases in that it can be described as a case with the following features: In Henry's belief-forming situation there are a couple of (relevant) cases that might obtain: Henry might be looking at a barn façade or at a real barn. Which one it will be is down to luck. Since, at the same time, Henry will form a belief that he is looking at a barn in either case that means that whether or not Henry's belief *turns out to be true* will be down to luck. That means, however, that Henry's belief is afflicted by veritic epistemic luck.

There is reason to believe, then, that in all three types of case that motivated the anti-luck platitude, the subjects' beliefs are afflicted by veritic luck. Given that this is so, there is also reason to believe that veritic epistemic luck is the type of luck at issue in the anti-luck platitude and that veritic epistemic luck is incompatible with knowledge. By the same token, the following principle appears to be a plausible development of the anti-luck platitude:

*Veritic Luck Principle*   Knowledge excludes veritic epistemic luck.[32]

## 2.2   Pritchard's modal conception of luck

As Pritchard has argued, there is strong reason to believe that the safety condition can be derived from the veritic luck principle. At the heart of Pritchard's argument is a

---

[32] Now one might wonder why I have presented the veritic luck principle as a *development* of the anti-luck platitude rather than as just another version of the anti-luck platitude. After all, it seems just as platitudinous to say that knowledge is true belief that is not merely luckily true as it is platitudinous to say that knowledge is luck-free true belief. True belief that is not merely luckily true, however, just is true belief that is not afflicted by veritic epistemic luck. If all this is correct, then the veritic luck principle ought to be just as platitudinous as the anti-luck platitude and ought to be viewed as another version rather than as a development of the anti-luck platitude. As I will argue in due course (chapter VI, section 3), however, appearances are misleading. There is another principle that competes with the veritic luck principle over cashing out the type of luck at issue in the anti-luck platitude. For that reason I regard the veritic luck principle as a development rather than just another version of the anti-luck principle.

modal conception of luck. Roughly, the notion of luck is captured in the following biconditional:

(LE)    An event is lucky *iff* it obtains in the actual world but does not obtain in a wide class of nearby possible worlds in which the relevant initial conditions for that event are the same as in the actual world.[33]

Pritchard offers two reasons for thinking that LE captures the notion of luck only roughly. First, it omits a dimension of the notion of luck—*viz.*, that an event is lucky only if it is in some way significant to someone. Second, the notion of relevant initial conditions that figures in it is notoriously vague. However, since the dimension of significance is, for present purposes, of little importance[34] and since we normally have a good intuitive grasp of what the relevant initial conditions of an event are, Pritchard argues that we can still expect LE to be illuminating in the project of uncovering the relation between luck and safety.[35,36]

Pritchard points out that there are a variety of facts that provide independent support for this modal conception of luck. In what follows, I will highlight what seem to me to be the two most important of these facts. First, the modal conception of luck appears to succeed in matching our intuitive verdicts of lucky and non-lucky events across a wide range of cases. For present purposes, it will be sufficient to look at the following few cases. As Pritchard correctly points out, a paradigm case of a lucky event is a lottery win. For suppose I buy a ticket in a fair lottery. Suppose my chance of winning is one in a million. I win. What could be a better illustration of a lucky event than that? Importantly, the event of me winning the lottery does satisfy LE. Although I

---

[33] Pritchard (forthcoming*a*), p. 3.

[34] We may with Pritchard assume that the significance condition is met. As Pritchard points out: "[S]ince [in the epistemic cases we are concerned with, CK] we are talking here about events which the agent has already formed a belief about, it is clear that these events must have had some impact on the agent concerned." (Pritchard (2005), p. 178, n. 2).

[35] Cf. Pritchard (forthcoming*a*), p. 3.

[36] I take it that in this case just as in the corresponding case of W-Safety* it is plausible that holding the relevant initial conditions fixed will admit of some variations in the initial conditions but only slight ones. It is important to keep that in mind in order to avoid potential equivocations in later arguments.

win in the actual world, there also is wide range of nearby possible worlds at which the relevant initial conditions for that event are held fixed—that is, at which I still own a ticket, my chance of winning is still one in a million, the lottery is still fair etc.—while I don't win the lottery. For that reason, the event of me winning the lottery satisfies LE. Another paradigm case of a lucky event Pritchard discusses is the event of someone who chances to dig a hole at the beach, for instance, and discovers a hidden treasure. Again, LE will predict that the event of the discovery is lucky. After all, there will be a wide range of nearby possible worlds at which the relevant initial conditions for that event are held fixed—I still dig a hole at the beach etc.—while I dig at a slightly different location and do not discover the treasure. On the other hand, a paradigm case of an event that is not lucky is the event of the sun rising in the morning. Again, LE succeeds in matching our intuitive verdicts. There is no nearby possible world—never mind a wide range of such worlds—at which the relevant initial conditions for that event are held fixed—in that the earth still revolves around the sun at the same speed as it does in the actual world, the are no comets approaching, we aren't at the verge of nuclear disaster etc.—at which the sun does not rise in the morning. If there is no wide range of nearby such worlds, LE predicts that the event of the sun rising in the morning is not lucky.

Moreover, as Pritchard argues in a joint paper with Matthew Smith, his modal conception of luck is continuous with and indeed serves to explain some of the results of recent scientific research on luck. Most importantly in this regard, scientific research suggests that counterfactual thinking plays a central role in attributions of luck. For instance, Johnson (1986) set up an experiment in which subjects were presented with descriptions of events with either a major positive outcome, a major negative outcome, a major positive outcome that almost did not occur, or a major negative outcome that almost did not occur and a control situation in which events with no major positive or

negative outcome were described. The subjects were asked to rate the events, among other things, in terms of luckiness. The result of the experiment was that subjects judged protagonists in events with major positive or negative outcomes that almost did not occur as luckier (unluckier) than the other ones. In this way the study suggests that counterfactual thinking, that is, thinking about what almost does or does not occur influences our attributions of luck.[37]

In another experiment by Teigen (1995) subjects were given descriptions of lucky and unlucky events (described as such by independent subjects) and asked how attractive they rated the events, whether they thought that something else might easily have happened instead and how attractive that other event would have been. Most importantly, the experiment showed that subjects who were given descriptions of events that were described as lucky or unlucky could very easily imagine something else to have happened instead. As opposed to that, subjects who were given descriptions of events that were not described as lucky or unlucky could not so easily imagine something else to have happened instead. Again the study suggests that counterfactual thinking is crucial to attributions of luck.[38]

A third type of study that highlights the role of counterfactual thinking in attributions of luck looks at the influence comparison processes have on such attributions. Experiments by Heider (1988) and Janoff-Bulman (1992) show that subjects who have had bad luck may judge themselves lucky because they can easily imagine an even worse outcome of the situation. For instance, someone who breaks an arm in a car accident may deem himself lucky because he was not killed instead.[39]

It can easily be seen that all of these studies provide support for Pritchard's modal conception of luck. After all, the counterfactual thinking the studies show to be

---

[37] Cf. Pritchard and Smith (2004), pp. 10-11.
[38] Cf. Ibid., p. 11.
[39] Cf. Ibid., p. 10.

central to attributions of luck concerns what might easily have happened. As we have already seen, however, given a standard possible worlds semantics of the notion of easy possibility, that is tantamount to saying that the relevant counterfactual thinking concerns what happens at nearby worlds. In this way, the studies provide reason to believe that Pritchard's modal conception of luck is on the right track.[40]

## 2.3   The Veritic Luck Principle and Safety

Given LE, we are also in a position to characterise the idea of a non-lucky event and the idea of an event that is not at all lucky. An event is non-lucky if and only if it occurs in the actual world and there is no wide range of nearby possible worlds at which the relevant initial conditions for that event are held fixed while it does not occur. An event is not at all lucky if and only if it occurs in the actual world and there is no range (however small) of nearby possible worlds at which the relevant initial conditions for that event are held fixed while it does not occur.

Recall that in section 2 reason was provided to believe that the type of luck at issue in the anti-luck platitude is veritic luck and that therefore the veritic luck principle is a plausible development of the anti-luck platitude. Now it is plausible that a true belief from which veritic epistemic luck is eliminated is a true belief that is not *at all* afflicted by veritic luck.[41] Recall, furthermore, that a belief is veritically lucky if and only if it is a matter of luck that the belief is *true*. If so, and if an event is not at all lucky if and only if fails to occur at no range (however small) of nearby possible worlds at which the relevant initial conditions for that event are held fixed, the idea of a true

---

[40] It is noteworthy that Pritchard further substantiates his modal conception of luck by arguing that it can do justice to intuitions about varying degrees of luck (cf. Pritchard (2005), p. 130 and (forthcoming), pp. 19-20) and that it can explain what is intuitive about alternative but ultimately unsatisfactory accounts of luck such as accounts that try to analyse (or elucidate) luck in terms of control, accidentality or chance.

[41] If this is not independently plausible, it can easily be seen that the veritic luck principle must be interpreted in this way if it is to explain our intuition that we do not know that given tickets in fair lotteries won't win (cf. Unger (1968), Pritchard (2005) and (forthcoming)).

belief from which veritic epistemic luck is eliminated, is naturally characterised along the following lines:

*Not-At-All-Lucky True Belief (NLTB)*    Veritic luck is eliminated from one's true belief if and only if one believes truly in the actual world and there is no range (however small) of nearby possible worlds in which the relevant initial conditions for the formation of the belief are the same as in the actual world while one believes falsely.[42]

Provided that, according to the veritic luck principle, knowledge requires that veritic luck be eliminated from one's true belief, we can now derive the following necessary condition for knowledge.

*Anti-Luck Condition (ALC)*    One knows that $p$ only if one truly believes that $p$ in the actual world and there is no range (however small) of nearby possible worlds in which the relevant initial conditions for the formation of the belief are the same as in the actual world while one believes falsely.[43]

In other words, ALC states that in order to know, one must have a true belief and avoid false belief at nearby possible worlds at which the relevant initial conditions for the formation of the belief are held fixed. Now notice that ALC and Safety are very similar indeed. In fact the only difference between the two is that while ALC requires true belief in the actual world, Safety only requires belief. Since true belief entails belief, ALC entails Safety. Given that this is so, we have the first argument for the safety condition on knowledge. In a nutshell, the argument for Safety takes the following shape:

    (1)  The veritic luck principle holds.
    (2)  Pritchard's modal conception of luck holds.
    (3)  By Pritchard's modal conception of luck, NLTB captures what it is for veritic epistemic luck to be eliminated from one's belief. (NLTB is the relevant instance of Pritchard's modal conception of luck.)

---

[42] Cf. Pritchard (forthcoming*a*), p. 6. Notice that Pritchard characterises the notion of a non-lucky belief here rather than the notion of a belief that is not at all lucky.
[43] Cf. Ibid., p. 7.

(4)  NLTB and the veritic luck principle entail ALC.
(5)  ALC entails Safety.
(6)  Hence, Safety holds.


## 3   A Craig-Style Argument for the Safety Condition

Apart from considerations concerning luck, there is yet another argument for the safety condition on knowledge. This argument exploits some of the considerations Craig adduces to support the claim that the concept of knowledge is the objectivised version of the concept of good informant. Recall in order to show this much, Craig looks at the application conditions of the concept of good informant and argues that they approximate the application conditions of the concept of knowledge as ordinarily understood and that objectivisation will deal with remaining deviations. Now recall that we observed earlier on that if Craig is right, if the concept of knowledge is the objectivised version of the concept of good informant, we may expect that the concept of knowledge retains the application conditions of the concept of good informant that remain after the objectivisation of the latter concept. What I will argue in this section is that there is one application condition of the concept of good informant such that (at least part of) what remains after filtering out all the subjective features pertaining to it is safe belief. If the argument is successful, then there is reason to believe that safe belief is an application condition of the concept of knowledge.

Recall that in order to get at the application conditions of the concept of good informant, Craig looks at a state of nature scenario and asks what properties an inquirer in such a scenario will want his informant to have. Accordingly, we can get at the application conditions of the concept of knowledge by filtering out any subjective features that pertain to the properties that an inquirer in such a scenario will want his informant to have.

Craig stated three properties that an inquirer in such a scenario will want his informant to have: To begin with, an inquirer who wants to acquire a belief about $P$ wants his informant to (a) truly and (b) believe that $p$. We have already seen that Craig maintains that the property of true belief will survive objectivisation and that we therefore have reason to believe that the application conditions of the concept of knowledge will include a truth and a belief condition. However, recall that Craig claims that an inquirer will also want his informant (c) to possess a "detectable property … which correlates well with being right about p".[44] Craig hastens to add that as regards the precise nature of the relevant property, or, in other words, as regards the answer to the question: What is X? in "S has a detectable property that correlates well with being right about $p$ iff X" there is no univocal answer. Rather, the relevant property is gerrymandered: "There could be almost as many different answers [to the question as to what the precise nature of the property at issue is, CK] as there are types of thing that the inquirer might want to know about."[45] Although it is plausible that Craig is right in claiming that there is no univocal answer to the question as to what the nature of this property of the good informant is, it may well be the case that once all the subjective features of this property of the good informant are filtered out, some property with a precisely identifiable nature remains. Again, it is my contention that at least part of what remains after filtering out the subjective features of the good informant's detectable property that correlates well with being right is the property of safe belief.

Let us turn to the detectable property that correlates well with being right that, according to Craig, an inquirer will want his prospective informant to have. As Craig himself notices one clearly subjective feature of this property lies in the requirement of detectability. Accordingly, objectivisation will filter out the requirement of detectability. So objectivisation leaves us with a property (detectable or not) that

---

[44] Craig (1990), p. 18.
[45] Ibid., p. 27.

correlates well with being right.[46] But now recall that, according to Craig, the property of correlating well with being right has a certain modal feature. A prospective informant who has this property must avoid false belief at nearby and open possibilities. An open possibility is characterised as a possible world that the inquirer does not know not to be the actual world. Notice that there is also a subjective element in this modal feature of the property of correlating well with being right: Whether or not a certain possibility is an open possibility for a given subject depends on what that subject knows about the actual world. However, different subjects may know different things about the actual world. In consequence, whether or not a given possibility is an open possibility may vary from one subject to another. Objectivisation will filter out this subjective element of the modal feature of the property at issue. What remains of this modal feature after objectivisation is the property of avoiding error at nearby possible worlds. Since the set of nearby worlds at which the relevant initial conditions for the formation of the informant's belief are held fixed is a subset of the set of nearby worlds, however, that also means that safe belief is at least part of what remains after the objectivisation of the modal feature of the third property of good informants.[47] So if we continue Craig's

---

[46] Cf. Craig (1990), p. 90.

[47] It is noteworthy that Craig would do well to place further restrictions on the range of nearby possible worlds at which the good informant must avoid false belief. To be more precise, he had better place restrictions on the range of nearby possible worlds at which the good informant must avoid false belief such that the informant needs to avoid false belief only at nearby possible worlds at which the relevant initial conditions for the formation of his belief are held fixed. If this is right, then we may expect safe belief to be exactly what remains after objectivisation of the modal feature of the property of correlating well with being right. Notice, however, that it will still be the case that the argument at best establishes that safe belief is part of what remains after the subjective features of the property at issue have been filtered out. After all, it may be that the modal feature is but one of a range of features of this property. And objectivisation of the other features may lead to other properties.

      The reason why Craig would do well to include the further restriction just mentioned is that, otherwise, his construal of the third property of good informants is bound to place too strong a demand on good informants. For suppose in the initial state of nature scenario there is no such restriction. Now recall that the inquirer in the state of nature scenario will want his informant to avoid false belief at nearby and *open* possibilities. Since inquirers often do not know in which way their informants have formed their beliefs, for instance, that means that among the nearby and open possibilities may be worlds at which the informant has formed her belief in a different way. This, however, will result in too strong a demand on good informants. For instance, in Nozick's grandmother case, by the lights of the present proposal, granny will not be a good informant for any inquirer who does not know that she has formed her belief by looking. After all, for any such inquirer it is an open possibility that granny formed her belief in another way, for instance, by testimony. However, at (some of) the nearby worlds at which granny forms her

project of filtering out subjective elements of the detectable property that correlates well with being right that, according to Craig, the inquirer in the state of nature scenario wants his informant to have, at least part of what we end up with is the property of safe belief. Since we have seen above that we may expect that filtering out any subjective features that pertain to this property will lead us to an application condition of the concept of knowledge, that means that the property of safe belief will figure in the application conditions of this concept: One cannot know a proposition unless one believes it safely. In this way we have a second argument for the safety condition on knowledge.

Now, even if one finds this argument appealing, one might wonder what sorts of considerations motivate the idea that the third property of good informants has the modal feature Craig maintains it has—*viz.*, that good informants avoid error at nearby possible worlds that are also open possibilities for the inquirer. Notice that this idea had better be well-motivated. For if it turned out to be ill-motivated, the argument for the safety condition on knowledge, which exploits this very modal feature of the third property of good informants, is bound to lose a great deal of any appeal it might initially have had.

One way in which one might attempt to motivate the modal feature of the third property of good informants is the following. Inquirers want to come to know the proposition they inquire about. If so, then since knowledge excludes veritic epistemic luck and since Pritchard's modal conception of luck holds, inquirers will want prospective informants to avoid false belief at nearby possible worlds. Unfortunately,

---

belief via testimony, she forms a false belief. Accordingly, for any inquirer who does not know that granny formed her belief by looking, granny will not avoid false belief at all open and nearby possibilities. In consequence, by the lights of the present proposal, granny is not a good informant for such inquirers. However, that does not seem to be right. Given that granny has formed her belief in the way and in the circumstances she did, it seems that she is a good informant with respect to whether or not her grandson is well. I, for one, would happily turn to granny to find out how her grandson is doing. So, it would seem that Craig had better place a further restriction on the nearby worlds at which a good informant has to avoid false belief—*viz.* a restriction such that they need only be worlds at which the relevant initial conditions for the formation of his belief are held fixed.

however, Craig cannot avail himself of this sort of explanation. After all, Craig wants to get from a state of nature scenario in which an inquirer wants to obtain a true belief from an informant to the concept of good informant and from that concept to the concept of knowledge. If so, then, on pain of circularity, he cannot avail himself of considerations regarding knowledge to motivate certain properties of the good informant.[48]

But how, if not by appeal to the idea that inquirers want to come to know, that knowledge excludes veritic luck and that Pritchard's modal conception of luck holds can Craig motivate the idea that inquirers want prospective informants to have a property with the relevant modal feature? Craig answers this question by arguing that inquirers have every reason to care about what happens at nearby and open possibilities. Here is the relevant passage:

> Imagine someone about to go out, and wishing to stay dry, but not knowing whether it will be raining or not, so that he faces one 'possible world' in which it will, and one in which it won't. Can he say 'I only want to keep dry in the actual world; I'm not bothered about whether I would have kept dry in whichever of those worlds turns out to be merely possible'? He can *say* it, and in a sense it is true. But if he proposes to do something about it he will either have to guess which possibility will be realised or take such action as will work in either case, even though that means planning for at least one eventuality which will turn out to have been merely possible. And the same applies to our inquirer: He will either have to guess which of the possible worlds he is actually in, or he must adopt a strategy which works in many merely possible worlds as well as the actual one.[49]

What Craig's remarks suggest is that when it comes to being successful in one's projects, what happens at nearby and open possibilities will be of importance to one. In the case he provides, an agent who wishes to stay dry and does not know whether it is going to rain (so that worlds at which it is going to rain are open possibilities for him)

---

[48] Notice, furthermore, that even though there is nothing about the methodology *I* am exploring that would render this argument inherently problematic, adopting it will still not advance my cause here. After all, we have already seen that the safety condition can be derived from the veritic luck principle and Pritchard's modal conception of luck. So, any argument that rests on both the veritic luck principle and Pritchard's modal conception of luck will not succeed in making an independent case for the safety condition.

[49] Craig (1990), p. 20.

will do well to plan for the eventuality of rain, by taking an umbrella, for instance. Now even if one agrees with this, one might wonder why the analogy is supposed to work, that is, why inquirers will also be interested in avoiding false belief at nearby and open possibilities. In response it may be pointed out that the relation between the agent whose interest brings with it an interest in what happens at nearby and open possibilities and the inquirer is more intimate than Craig's remarks in the above quotation may suggest. To see why this is so, notice that in order to be successful in one's projects one needs true beliefs. For instance, Craig's agent will be successful in his project of staying dry by taking an umbrella to protect him from any rain that may fall while he is outside only if his belief that an umbrella will protect him from rain is true. Now if one has reason to care about what happens at nearby and open possibilities when it comes to being successful in one's projects, then, since in order to be successful in one's projects one needs true beliefs, one has also reason to care about continuing to believe truly at nearby and open possibilities. Accordingly, if one were to acquire a belief from an informant, one also has reason to care that one's informant believes truly not only in the actual world but also at nearby and open possibilities.

Now, one might object to this line of thought by pointing out that while it shows that, for *some* of our beliefs, we have reason to care whether they continue to be true at nearby and open possibilities, it certainly does not show that this is so for *all* of our beliefs. After all, it is plausible that not all of our beliefs—not even all of our beliefs that we acquire via testimony from other people—will become relevant to some project of ours. So, even if the above claim about the relation between successful projects and caring about what happens at nearby and open possibilities is true, and even if the relation between successful projects and true belief holds in the way indicated, we will not have done enough to motivate the claim that *any* good informant must avoid false belief at nearby and open possibilities. In order to do so, it would have to be shown that

any inquirer would care about what happens to the truth-value of the proposition she comes to believe at nearby and open possibilities. But since the above argument establishes at most that inquirers who want to put their belief to use in one of their projects have reason to care about what happens to the truth-value of the propositions they come to believe at nearby and open possibilities, it would now seem that there is a gap in the argument.

However, it might seem that things are even worse. For there is reason to believe that it is not only the case that the intended development of Craig's argument from analogy won't do the job we wanted it to do, but also that even the undeveloped argument from analogy does not go through. For one might think that, for the argument from analogy to work, it must be the case that *any* agent has reason to care about what happens at nearby and open possibilities. If it turned out that *only some* agents have reason to care about what happens at nearby and open possibilities, then even if the intended analogy between agents and inquirers did hold, all the analogy could establish would be that there are *some* inquirers who have reason to care about what happens at nearby and open possibilities. That, in turn, will only be enough to motivate the claim that *some* good informants must avoid false belief at nearby and open possibilities. Again, we would not have done enough to motivate the claim that *any* good informant will have to do so. Again, we would have failed to properly motivate the modal feature of the third property of good informants. The problem is, of course, that there is excellent reason to believe that not all—or equivalently, only some—agents have reason to care about what happens at nearby and open possibilities. To see why this is so, reconsider the very example Craig provides. A person who wishes to stay dry and does not know whether it is going to rain is well advised to prepare for the worst—even if that means preparing for one world that turns out to be merely possible. This example is supposed to illustrate that agents have reason to care about what happens at nearby and

open possibilities. For the example to work, however, it is crucial that the agent *wishes to stay dry*. For suppose that the agent just does not care whether or not she gets wet. It is far from clear that in that case the agent is well advised to care about worlds that will remain merely possible, for instance, by taking an umbrella. Indeed, if the agent does not care whether she will get wet, it seems that she is better off not taking an umbrella. After all, she will have one thing less to carry around with her. What this variation of Craig's case suggests is, of course, that not all agents have reason to care about what happens at nearby and open possibilities. So, even Craig's undeveloped argument from analogy fails to motivate the modal feature of the third property of good informants. As it stands, the claim that the third property of good informants has the relevant modal feature lacks satisfactory motivation. In consequence, the argument for the safety condition on knowledge, which rests on this claim, threatens not to work.

Now, one might be inclined to simply reply that the situation of the person who does or does not want to get wet is different from the situation of the inquirer. While wishing to avoid getting wet is optional for the former person, wishing to avoid false belief is not optional for the inquirer. The thought here is that to be an inquirer is to be a person who aims at believing the truth and—more importantly for present purposes— also aims at avoiding false belief. Any person who gives up either aim ceases to be an inquirer. So, while the argument depends on the assumption that the inquirer wishes to avoid false belief, the assumption can be discharged. The problem we are now facing is, of course, that even if the assumption can be discharged, we are still missing a crucial bit to complete the argument as it was initially intended. For it still seems that, first, since the assumption cannot be discharged in cases of agents we have no basis to appeal to *analogy* between cases of agents and cases of inquirers, and, second, we will also be unable to run the argument that exploits the intimate link between true belief and successful projects of agents. The argument is still at best incomplete.

But maybe we can now run the argument without relying on Craig's original analogy or on the development I have suggested. We might just go back to the state of nature scenario in which we have an inquirer who wants to get a belief from a prospective informant and ask ourselves what features such an inquirer would want his informant to have. We now notice that an inquirer is by his very nature someone who aims at avoiding false belief. So, he will want his prospective informant to help him secure his aims. At the same time, however, the inquirer's knowledge of the actual world is incomplete. In consequence, there are many nearby possible worlds that, for all the inquirer knows, might be the actual world. In order to properly *secure* his aim of avoiding false belief, the inquirer must avoid false belief not only in the actual world but also at nearby and open possibilities. Accordingly, given that he wants his prospective informant to help him secure his aim of avoiding false belief, he will want his prospective informant to avoid error not only in the actual world but also at nearby possible worlds which, for all he, the inquirer, knows, might be actual. In this way, one might think, the modal feature of the third property of good informants can be motivated.

Again, however, one may have some worries about this argument. While one might be willing to concede that inquirers by nature aim at avoiding false belief, how important avoidance of error is may vary dramatically from one inquirer to another and from one situation in which a given inquirer may find himself to another. Consider, for instance, a situation in which the cost of error is relatively low, while the gain of being right is very high. Suppose one has to choose between two prospective informants. One informant is ready to hand but he is only moderately likely to be right. He avoids false belief at most nearby possible worlds that, for all one knows, might be actual but not at all of them. The other informant is certain to be right. He avoids false belief at every nearby possible world that, for all one knows, might be actual. However, it is extremely

hard to get hold of him. (Perhaps one would have to travel to the other end of the world to get the information from him.) It is far from clear that the only good informant in such a situation is the one who avoids error across the entire range of nearby possible worlds one does not know not to be actual. It may well be that the prospective informant whom one can easily consult is still a good enough informant given one's present intents and purposes. If so, however, then even if inquirers by their very nature seek to avoid false belief, it need not always be the case that the good informant avoids false belief at every nearby possible world the inquirer does not know not to be the actual world. The problem is, of course, that even if inquirers by nature seek to avoid false belief, they do not always have reason to *secure* or *guarantee* that they avoid false belief. The situation may sometimes be such that doing less than *securing* avoidance of error will be enough. If good informants need not always avoid false belief at all nearby possible worlds that, for all the inquirer knows, might be actual, that means that the relation between good informants and having the property of avoiding false belief at nearby possible worlds that, for all the inquirer knows, might be actual that we aimed to establish on Craig's behalf breaks down again.

It might be responded that the objector does not describe the situation in the right way. The right way of describing it is the following: The only good informant is the one who avoids false belief at all nearby worlds the inquirer does not know not to be the actual world. However, in the present case, it is reasonable for the inquirer to compromise quality for availability. If so, it is open to us to say that the available informant at issue is, strictly speaking, still not a good one.

The problem with this response is, of course, that, on Craig's premises, we get to the properties of the good informant from a state of nature scenario, a scenario, that is, in which an inquirer wants to acquire a belief from someone else and by asking what features the inquirer would want his prospective informant to have. It just not clear at

all, however, that an inquirer in that situation will always want his informant to avoid false belief at all nearby possible worlds he does not know not to be the actual world. In the situation described, for instance, being right at less than all of these nearby possible worlds might do just as well. Or so the objector might argue.

I do not want to decide the argument between Craig and the objector here. Instead, I will try to show that even if the objector is right and good informants do not always have to avoid false belief at all nearby possible worlds the inquirer does not know not to be actual, it can be argued that the argument that I am interested in—*viz.*, the one showing that knowledge requires safety—can be recovered. To begin with, we start with a concession to the objector and revise the modal feature of the property that, by Craig's lights, every good informant must possess. Rather than construing the modal feature such that good informants must avoid error across the entire range of nearby and open possibilities, we now hold that good informants must avoid false belief across as wide a range of nearby possible worlds the inquirer does not know not to be the actual world as the purposes of the inquirer require.[50,51]

---

[50] Notice that this way of construing the property will also do justice to the above problematic case. The prospective informant who is readily available to one is still a good informant in one's situation because he avoids error at a range of nearby and open possibilities that is wide enough for one's purposes.

[51] One might now wonder why we should think that the third property of good informants has a modal feature in the first place. By way of response, reconsider first the last argument we suggested on Craig's behalf to motivate the modal feature of the third property of good informants: We go back to the state of nature scenario in which we have an inquirer who wants to get a belief from a prospective informant and ask ourselves what properties such an inquirer would want his informant to have. We then notice that an inquirer is by his very nature someone who aims at avoiding false belief. At the same time, his knowledge of the actual world is incomplete such that there are many possible worlds that, for all he knows, might be the actual world. The inquirer wants his informant to help him secure his aim. We can motivate the idea that the third property of good informants as presently construed has a modal feature by arguing that it provides a precise and plausible account of what the inquirer wants his informant to be like *in case avoidance of error is of utmost importance to the inquirer*. In that case the inquirer will want the informant to *guarantee* avoidance of error. The informant will guarantee avoidance of error if he avoids false belief at nearby possible worlds that, for all the inquirer knows, might be actual. We can then make room for the suggestion that avoiding error need not always be of utmost importance to the inquirer by allowing that the inquirer may sometimes be content with an informant who avoids error at less than the entire range of nearby and open possibilities. Indeed, we can argue that different degrees in the importance to the inquirer of avoiding error determine different ranges of nearby and open possibilities across which a good informant must avoid error. In this way we can motivate the modal feature of the third property of good informants as presently construed.

Next, we go on filter out the subjective features of this property of good informants. To begin with, we observe that across exactly how wide a range of nearby and open possibilities an inquirer's purposes require his prospective informant to avoid false belief is a subjective matter. For a given proposition, the purposes of one inquirer may require a good informant to avoid false belief across the entire range of nearby and open possibilities, while, relative to the purposes of another inquirer, someone may be a good informant even though he avoids false belief only across a moderately wide range of nearby and open possibilities. Accordingly, this feature will be filtered out by objectivisation. It is plausible, however, objectivisation here will lead to a *tightening* of the demands on good informants.[52] To see why this is so, notice first, that we want the objectivised version of the concept of good informant to apply not just to individual inquirers and informants but also to linguistic communities with a practice of passing on information. But now notice, second, that members of such communities often do not know across how wide a range of nearby possible worlds the purposes of a given inquirer require a prospective informant to avoid false belief. For instance, suppose that Miller has hired Smith to find an informant about *P*. Suppose, Smith comes to me and asks me whether I can supply the relevant information. I can't but I know that Jones can and so I recommend Jones to Smith as an informant about *P*. In this case, neither Smith nor I may have any inkling across how wide a range of nearby possible worlds the prospective informant (Jones) has to avoid false belief in order to suit the purposes of Miller. Miller's purposes could require avoidance of error across as wide a range of nearby and open possibilities as anyone's purposes ever will. In such a case, however, we will be well advised to prepare for the worst, as it were, and make sure that the prospective informant avoids error at a wide enough range of nearby possible worlds to suit the purposes of the inquirer, in the present case Miller, *even if Miller's purposes*

[52] Cf. also Craig's argument that the demands on exactly how well a good informant's property must correlate with being right will be driven up by objectivisation ((1990), p. 91).

*require avoidance of error across as wide a range of nearby and open possibilities as anyone's purposes ever will.* It is plausible, however, that the widest such range of nearby and open possibilities is the entire range of nearby and open possibilities. If so, since we want the objectivised version of the concept of good informant to apply to cases like this one as well, there is reason to believe that objectivisation will lead to a tightening of the demands on good informants. What remains of the modal feature of the third property of good informants after objectivisation, then, is the property of avoiding false belief across the entire range of nearby possible worlds that, for all the inquirer knows, might be actual. We can now run the argument for the safety condition as rehearsed at the beginning of this section. Thus, even if the objector is right and Craig's claim that good informants must avoid false belief at all nearby possible worlds that, for all the inquirer knows, might be actual, is false, the argument that objectivisation of the third property of good informants leads to the property of safe belief can be recovered. In this way, Craig-style reasoning can still be put to use to argue that knowledge requires safety.

## 4   Safety is Consistent with the Other Platitudes

We have now looked at two arguments suggesting that safe belief is a necessary condition for knowledge. Since the safety condition can be derived from the anti-luck platitude—in conjunction with some further independently plausible assumptions—the safety condition will also be consistent with the anti-luck platitude. In order to complete the case for safety, however, it will be necessary to show that the safety condition is consistent with the other platitudes of the minimalist framework. This is what I will turn to now. I will begin with an argument suggesting that the safety condition on knowledge is consistent with the good informant platitude.

## 4.1 The Good Informant Platitude

According to the good informant platitude, knowers are typically—that is, except when they are unwilling to part their information, have lost their credibility etc.—good informants. Accordingly, the good informant platitude will be inconsistent with the safety condition just in case either of the following two entailments holds: First, if one satisfies the safety condition, then one is not a good informant (except when one unwilling to part one's information etc.). Or, second, if one is a good informant, then one does not satisfy the safety condition. It can easily be seen that neither entailment holds. Suppose Henry drives through Normal Barn Country, sees a barn and comes to believe that there is a barn in the field. Henry's belief is true. In since in Normal Barn County Henry would not mistake a barn for anything that is not a barn, Henry's belief will also be safe. Thus, Henry satisfies the safety condition. At the same time—assuming, as we may, that Henry is willing to part his information, is credible etc.—he is certainly also a good informant about whether he is facing a barn. If I, for one, were lying in the back, trying to sleep but were overcome by the desire to know whether there is a barn in the field, I would not hesitate to ask Henry. Thus the first entailment does not hold. As regards the second entailment, suppose a schoolboy cannot remember whether World War II ended in 1945 or 1944. Certainly, his history teacher—whom we may suppose to be willing to part the information, be credible etc.—will be a good informant for our schoolboy. At the same time, the history teacher not only knows that World War II ended in 1945, he also entertains his belief on the basis of reliable sources and memory neither of which would have led him to a false belief. Thus, the history teacher is not only a good informant but also satisfies the safety condition. The second entailment does not hold either. The safety condition and the good informant platitude are thus consistent.

## 4.2 The Informative Speech Act and Assertion Platitudes

One way of arguing that the safety condition on knowledge is consistent with the informative speech act platitude—according to which, remember, in performing an informative speech act that *p* one represents oneself as knowing that *p*—exploits the relation between the interactions between inquirers and informants Craig describes and the informative speech act platitude. Recall that in the last chapter we found that one argument for the informative speech act platitude was that if the interactions between inquirers and informants Craig describes were to be enshrined in a practice of passing on information, then the features the inquirer will want the informant to have will be turned into demands on the informant. Recall that we also found that since any subjective features the inquirer may want his informant to have will be of little importance to the general practice, the features that will be enshrined in this practice will be the ones that remain after filtering out these subjective features. Since, as we have just seen, one feature that remains after filtering out these subjective features is safe belief, we have reason to believe that one of the demands the practice places on informants is that they believe safely. If so, however, we may expect the informative speech act platitude and the safety condition to be consistent.

I will try to provide further support for the idea that the informative speech act platitude and the assertion platitude—which does not (at least not in any obvious manner) receive support from the above argument—are consistent with the safety condition by showing that there is reason to believe that some situations that would show the inconsistency of the safety condition on the one hand and the informative speech act and assertion platitudes on the other, do in fact not obtain. The situations I have in mind are situations in which one felicitously asserts/performs an informative speech act that *p* whilst, at the same time, not safely believing that *p*. Provided that any

such informative speech act/assertion will be felicitous only if one also knows its content, any such situation would show that the informative speech act/assertion platitude and the safety condition are inconsistent. Of course, we must be conscious of the fact that an assertion/informative speech act can sometimes be felicitous even if one does not know its content. For instance, an assertion/informative speech act may be felicitous if one reasonably believes that one knows its content or if the norm for assertion/informative speech acts is overridden by other norms.[53] So, in order to assess whether there are cases in which one felicitously performs an assertion/informative speech act without safely believing its content, we must be careful to construe the case in such a way that the subject does not also reasonably believe that she knows its content and that the norm governing assertion/informative speech acts is not overridden by other norms.

So are there situations in which one felicitously performs an informative speech act/makes an assertion whilst not safely believing its content? In order to answer this question, let us first reflect on the sorts of situation in which one believes a given proposition unsafely. It seems that the following four kinds of situation capture the widest range of such situations: first, one believes falsely; second, one believes unreliably; third, one's belief is Gettiered; and, fourth, one believes a lottery proposition.[54] Is it the case that one can felicitously perform an informative speech act/make an assertion expressing one's belief in either of those situations? It seems to me that the answer is 'no'. First, if one believes that $p$ falsely, then one's informative speech act/assertion that $p$ will be infelicitous. For instance, if I believe that Fermat's

---

[53] Cf. Williamson (2000), pp. 256-7.
[54] So far I have not argued that Gettiered subject's lack safe beliefs or that beliefs in lottery propositions are unsafe. I will provide some arguments for these claims later on. For now, notice that for present purposes it is not required that safety is the anti-Gettier condition or that beliefs in lottery propositions are always unsafe. All that I have said so far is that (some) Gettiered beliefs and (some) beliefs in lottery propositions are members of set of unsafe beliefs. That much seems plausible. Even if it turned out to be false, however, the argument would remain unaffected. If Gettiered beliefs and beliefs in lottery propositions were never unsafe it would only mean that they could be left out of the list of the main types of unsafe beliefs. The argument would only be somewhat shorter and simpler.

last theorem has not been proven just because it had not been proven by 1950[55] and I tell someone/assert what I believe, the informative speech act I perform/the assertion I make is not felicitous. Similarly, if I believe that Fermat's last theorem has been proven but believe it on the basis of having read the tea leaves[56] and then go on to tell someone/assert what I believe, my informative speech act will not be felicitous either. Moreover, we have already seen that informative speech acts/assertions to the effect that a given ticket in a fair lottery won't win are not felicitous.[57] So, it is not hard to find evidence that, in situations in which the content of an informative speech act/assertion is false or a lottery statement as well as in situations in which the belief represented in the informative speech act/assertion is unreliably formed, the informative speech act/assertion is also infelicitous. As opposed to these situations, situations in which the subject's belief is Gettiered are harder to assess because Gettiered subjects often have a reasonable belief that they know the proposition they in fact do not know. So, in order to justly judge whether a Gettiered subject's informative speech act/assertion is felicitous in the relevant sense we will have to make sure that we construe the case in such a way that the subject does not have a reasonable belief that she knows the proposition at issue. The following case seems to me to be a decent shot at such a case. Suppose a friend of mine and I are waiting for Jane to come and pick us up with her car. Suppose I see a black Citroen DS21 and since I know that Jane owns a black DS21, I also know that there aren't very many DS21s around, that she usually comes from that direction and that she is usually very punctual, I come to believe that it's Jane's car that I'm seeing and that Jane will be here in a few moments. My belief is justified and true. Unbeknownst to me, however, Jane has lent her black DS21 to a friend and is coming in

---

[55] I take it that if I base my belief on these grounds then I do not reasonably believe that I know what I believe.

[56] Again, someone who bases his belief in this way will not reasonably believe that he knows.

[57] We know that we do not know that given tickets in fair lotteries won't win. So we do not reasonably believe that we *do* know that they won't win.

her blue Alfa Giulia instead. However, the car I am seeing is still Jane's, although driven by her friend and it is still true that Jane (in her blue Giulia) is about to pick us up. Now, suppose I tell my friend/assert that I have just seen Jane's black DS21 and that Jane will be here to pick us up in a few moments. My friend, who is a slightly more cautious person when it comes to saying things, points out to me that Jane's is by far not the only black DS21 in town. Suppose I take his point, but continue to believe that it was Jane's black DS21 and that she will be here to pick us up soon. Now I take it that since I took my friend's point I no longer reasonably believe that I know that I saw Jane's car and that she will soon be here to pick us up. An accepted but undefeated challenge should establish this much. At the same time, the challenge does not undercut my justification. The number of black DS21s in town is still rather small and thus the probability that there should be one at exactly this time and at exactly this end of town is still low. So, I am still Gettiered. I still have a justified true belief that does not qualify as knowledge. Moreover, I do not withdraw my informative speech act/assertion. On the contrary, when a relative of mine passes by and asks us what we are doing, I tell him/assert the same as I told my friend/asserted earlier on. Now, it would seem that my informative speech act/assertion in this situation is not felicitous. As an indication of this, notice that when Jane turns up in her blue Giulia my friend may justifiably blame me for telling my relative/asserting what I did. He may advise me to be more cautious in what I tell people/what I assert and so on. So, there is also evidence that if a Gettiered subject performs an informative speech act/assertion—and does not reasonably believe that he knows the content of his speech act—her speech act will be infelicitous.

In this way, there is reason to believe that informative speech acts/assertions in which the belief represented is unsafe will not be felicitous—at least not unless exceptional circumstances prevail. So, the most imminent threat to the suggestion that

the informative speech act or the assertion platitude is consistent with the safety condition on knowledge can be defused.

## 4.3   The Factivity and Belief Platitudes

Let me then turn to the factivity- and belief platitudes. It is not at all hard to see that the condition on knowledge that safety-based views propound is consistent with the belief-platitude. After all, the condition at issue has it that in order to know a proposition, *P*, one has to have a safe belief that *p*. Thus, it follows from the safety condition on knowledge that one knows that *p* only if one also believes that *p*. The belief platitude is thus entailed by the safety condition on knowledge. Thus, the former is consistent with the latter.

Similarly, the reason why the safety condition on knowledge is consistent with the factivity platitude is that safe belief is factive. If so, then the safety condition on knowledge entails the factivity platitude, which, in turn, means that the two are consistent. Here comes the argument. To begin with, recall that one's belief that *p* is safe if and only if one believes that *p* at the actual world and one avoids false belief at nearby possible worlds at which the relevant initial conditions for the formation of the belief are held fixed. If one avoids false belief at nearby possible worlds at which the relevant initial conditions for the formation of the belief are held fixed, then, since actuality is close to itself, and since in actuality the relevant initial conditions for the formation of the belief are the same as in actuality, in actuality one avoids false belief. It follows that any safe belief must also be true. Safe belief is thus factive. So, if safe belief is required for knowledge, then knowledge is factive. Thus, the safety condition on knowledge is consistent with the factivity platitude.[58]

---

[58] Thanks to Patrick Greenough for pointing this out to me.

It is noteworthy that this argument also has another interesting consequence. Since safe belief is factive, if knowledge requires safe belief, it follows that knowledge requires safe true belief. Thus the safety condition implies:

(Safety*)        One knows that *p* only if one truly believes that *p* and one avoids false belief at all nearby possible worlds at which the relevant initial conditions for the formation of one's belief are held fixed.

But now notice that Safety* is equivalent to ALC—recall that, according to ALC, knowledge also requires true belief in actuality and avoidance of false belief at nearby worlds at which the relevant initial conditions for the formation of the belief are held fixed. Given that this is so, we can argue from the safety condition and Pritchard's modal conception of luck to the veritic luck principle. The argument proceeds along the following lines:

(1) Safety holds.
(2) Safe belief is factive.
(3) Hence, Safety* holds.
(4) Safety* is equivalent to ALC.
(5) Pritchard's modal conception of luck holds.
(6) By Pritchard's modal conception of luck, NLTB captures what it is for veritic epistemic luck to be eliminated from one's belief. (NLTB is the relevant instance of Pritchard's modal conception of luck.)
(7) ALC and NLTB entail that the veritic luck principle.
(8) Hence, the veritic luck principle holds.

Since we have already seen that, given Pritchard's modal conception of luck, the veritic luck principle entails the safety condition, it follows that, given Pritchard's modal conception of luck, the veritic luck principle and the safety condition are mutually derivable. Given, additionally, that veritic luck is the type of luck at issue in the anti-luck platitude, we may expect that the safety condition fully captures the sense in which knowledge excludes luck at issue in the anti-luck platitude.

## 4.4   The Closure Platitude

In order to show that the safety condition is consistent with the closure platitude, I will argue that any belief formed on the basis of competent deduction from known premises is also safe. Or, in other words, whenever one forms a belief in a proposition, $Q$, on the basis of competent deduction of $Q$ from a set of known premises $P_1 \ldots P_n$ then one's belief that $Q$ is also safe. If so, then there can be no threat to the closure platitude from the safety condition on knowledge. The two will be consistent.

To begin with, to say that one's belief in the conclusion, $Q$, of one's competent deduction from known premises, $P_1 \ldots P_n$, is safe is to say that there is no nearby world at which the relevant initial conditions for the formation of one's belief are held fixed, at which one believes $Q$ falsely. I take it that in the case of a belief formed on the basis of competent deduction from known premises, it will be part of the relevant initial conditions for the formation of the belief (a) that the premises are known and (b) that the subject competently deduces the conclusion from the premises. Now since knowledge is factive, at every possible world at which one knows the premises, $P_1 \ldots P_n$, of one's deduction, the premises are also true. If so, however, then, at every possible world at which one competently deduces $Q$ from $P_1 \ldots P_n$, where $P_1 \ldots P_n$ are known, $Q$ must also be true. After all, one cannot have competently deduced a proposition, $Q$, from a set of propositions, $P_1 \ldots P_n$, unless $P_1 \ldots P_n$ entail $Q$. Since the set of nearby possible worlds is a subset of the set of possible worlds, at all nearby worlds at which one competently deduces $Q$ from $P_1 \ldots P_n$, where $P_1 \ldots P_n$ are known, $Q$ will be true. Since (a) that the premises are known and (b) that the subject competently deduces the conclusion from the premises are part of the relevant initial conditions for the formation of the belief in $Q$ that need to remain fixed, that means that any subject who competently deduces and thereupon comes to believe $Q$ from a set of known premises

$P_1 \ldots P_n$ will avoid false belief at nearby possible worlds at which the relevant initial conditions for the formation of her belief are held fixed. That is to say, however, that such a subject's belief in $Q$ will be safe. The property of safety is transmitted to any belief based on competent deduction from known premises. So, the safety condition is consistent with the deduction platitude.

In summary, then, we have looked at two arguments suggesting that safe belief is a necessary condition for knowledge. The first argument had it that the safety condition captures the sense in which knowledge excludes luck, while the second argument exploited the fact that safe belief is (at least part of) what remains after objectivisation of Craig's third property of good informants. Moreover, it has been argued that the safety condition is consistent with the platitudes of the minimalist framework. In this way, a case has been built for a conception of knowledge according to which knowledge requires safe belief. In the next chapter, I will look at what we may expect from such a safety-based conception of knowledge and whether it can meet the expectations.

# V  A Closer Scrutiny of Minimalist Safety

Now that we have seen how the minimalist framework developed in chapter III can be used to argue for a safety-based conception of knowledge, in this chapter I will discuss what we may expect from a safety-based conception of knowledge thus developed (or, for short, 'minimalist safety') and whether it can meet these expectations.

## 1  What We May Expect from Minimalist Safety

Let us begin by describing some of the things that we may expect the safety condition on knowledge as developed from the minimalist framework to do for us. To begin with, notice that, according to the present proposal, minimalist safety must be understood as capturing the sense in which knowledge excludes luck. After all, recall that in the argument that connected the anti-luck platitude with the safety condition we first provided reason to believe that the type of luck at issue in the anti-luck platitude is veritic epistemic luck; second, derived the safety condition from the corresponding veritic luck principle and the independently plausible modal conception of luck, LE; and, third, saw that the veritic luck principle can also be derived from the safety condition and LE. Given that all this is so, we may expect minimalist safety to capture the sense in which knowledge excludes luck and we may also expect minimalist safety to do the job the anti-luck platitude does. First and foremost, that means that we may expect minimalist safety to explain those facts about our use of the word 'know' in ordinary language and thought that motivate the anti-luck platitude.

Now, recall that it was argued in chapter III, section 6 that there are at least three facts about our ordinary use of the word 'know' that motivate the anti-luck platitude. First, there are cases of unreliably formed belief in which we do not accord knowledge

to the subject—even if these beliefs turn out to be true. For instance, someone who looks at a crystal ball does not come to know what she 'sees' in the ball. Second, we do not count beliefs based on probabilistic evidence that given tickets in fair lotteries won't win as knowledge. For instance, my belief that my friend's numbers won't win the state lottery does not qualify as knowledge. Third, we do not apply the word 'knows' to subjects in Gettier cases. In Barn Façade County, Henry's belief that he is looking at a barn, even if true, does not qualify as knowledge. In all three types of case the diagnosis was that the subject is too lucky to be credited with knowledge. In this way, these cases motivated the anti-luck platitude.

If we may expect minimalist safety to capture the sense in which knowledge excludes luck, then we may also expect it to explain why subjects who hold unreliably formed beliefs, subjects who have beliefs that given tickets in fair lotteries won't win, and Gettiered subjects lack knowledge. That is to say that we may expect that the beliefs of subjects in these types of case turn out to be unsafe. One important caveat: We may expect minimalist safety to explain that subjects in these types of case lack knowledge only to the extent that the minimalist framework has identified the sources of our intuitions correctly. Or, in other words, if there are certain cases of the abovementioned types in which the subject's ignorance is not grounded in the belief's being afflicted by luck, but rather has some other ground, then we need not expect minimalist safety to explain the subject's ignorance.

## 2  Worries Arising for Minimalist Safety

### 2.1  Gettier Cases and the Analytical Problem

Notice that if the above considerations are true, then we must now also expect the safety condition to be the anti-Gettier condition. Now one might be worried that if the

arguments that yielded the safety condition commit us to this much, we must somehow have taken the wrong turn. After all, the persistence of the problem Gettier cases pose for the traditional approach to answering the question as to what knowledge is—*viz.* by providing an analysis in terms of necessary and sufficient conditions—contributed significantly to motivating the project of looking into alternative approaches to answering this very question. If we now claim to have found the solution to the Gettier problem do we not end up with a traditional analysis of knowledge after all? Or, at the very least, do we not lose any advantage our non-traditional approach to the analysis of knowledge may have had vis-à-vis traditional approaches? And does attempting to solve the Gettier problem not undermine the motivations for the non-traditional project?

Let us turn to the first question first. There is excellent reason to believe that we do not end up with a traditional analysis of knowledge. For, recall that we found that Gettier's refutation of the received view not only spurred the debate over the anti-Gettier condition but also over the status of the internalist justification condition at issue in the received view. In order to be in a position to claim to have identified a sufficient condition for knowledge, we would first have to resolve the internalism/externalism debate. However, nothing in the present discussion contributes towards resolving this debate. Moreover, since we have seen that the debate between internalists and externalists may be irresolvable, it may even still be the case that we cannot identify a set of individually necessary and jointly sufficient conditions for knowledge.

In response to the second question—whether we must expect that if we attempt to solve the Gettier problem we will end up losing whatever advantage we may have had vis-à-vis traditional approaches to the analytical problem—I would like to make the following remarks. To begin with, recall how the Gettier literature has evolved: It started out with Gettier's counterexamples to the received view. Epistemologists then tried to solve the problem posed by Gettier's counterexamples by patching up or

modifying the received view in such a way as to deal with the counterexamples. Once the patched-up or modified analysis was in place, however, it did not take long until Gettier-style counterexamples to the patched-up or modified analysis were discovered. This led to further patching-up or modification. The resulting proposed analysis was, of course, designed to deal with the then state-of-the-art Gettier cases. However, soon after the new proposed analysis appeared to have found its foothold in the theory of knowledge, it fell prey to a new Gettier-style counterexample and so on.[1] But now notice that, as opposed to the various proposals in the Gettier literature, the minimalist approach roots out the source of our intuitions in Gettier cases—*viz.* an implicit commitment to the anti-luck platitude. It is then argued that the type of luck at issue in the anti-luck platitude is veritic luck and that given an independently plausible further assumption—*viz.* Pritchard's modal conception of luck—the corresponding veritic luck principle and the safety condition are mutually derivable. It is only because and— remember the caveat of section 1—only to the extent that the minimalist framework has provided a convincing diagnosis of the sources of the Gettier problem that defenders of minimalist safety will claim to have identified the anti-Gettier condition. There is, then, an important difference between the various proposals in the Gettier literature and minimalist safety: while the former provide a set of conditions that is able to deal with some state-of-the-art Gettier cases and then hope for the best, the latter may be said to solve the problem only because and only insofar as it can also be said to tackle the problem at its roots. In this way, there is reason to believe that minimalist safety will successfully deal with the Gettier problem in a way in which there isn't such reason to believe that the various proposals in the Gettier literature will do so. So there is reason to believe that defenders of minimalist safety can retain (or regain) an edge over traditional approaches.

---

[1] Cf. Kirkham (1984), pp. 501-2.

As regards the third question—whether the fact that we now attempt to solve the Gettier problem undermines any motivation we may have had to adopt the non-traditional approach in the first place—notice, first, how the non-traditional project was motivated: It was the unfortunate series of proposed analyses of knowledge and Gettier cases on the one hand and the continuing debate between internalists and externalists about knowledge on the other that did the job. True, part of what motivated the non-traditional approach was that it promised to be illuminating without having to address the Gettier problem. However, it should be unproblematic to use the persistent failure of a certain type of approach to a given problem to motivate an alternative type of approach which does not need to address the problem that led to the failure of the problematic type, whilst, once there is special reason to believe that the alternative approach will be more successful than its competitors in addressing the problem that led to the failure of the problematic type of approach, ending up addressing the problem after all without thereby undermining the motivation for adopting the alternative approach in the first place. Since, as we have just seen, there is special reason to believe that minimalist safety will solve the Gettier problem in a way in which there is no such reason to believe that traditional approaches to the analytical problem will, there is also reason to believe that minimalist safety's claim to solve the Gettier problem does not undermine the motivations for adopting a minimalist approach to the analytical problem in the first place.

## 2.2 Craig, Floridi, and Zagzebski on the Unsolvability of the Gettier Problem

At this point, it will be instructive to note that the safety condition can also avoid probably the most serious worry any theory of knowledge that professes to solve the

Gettier problem faces. There is a family of arguments that provide very strong—if not conclusive—reason to believe that a certain type of analysis of knowledge will never solve the Gettier problem. One version of the argument can be found in Craig (1990) and Zagzebski (1994), (1996). An even stronger version of the argument has recently been provided by Floridi (2004), who argues not only that the relevant type of analysis will never solve the Gettier problem but even that the Gettier problem is logically unsolvable by analyses of this type. The relevant type of analysis is of the following form:

(TB+X)   S knows that $p$ if and only if   (a) $P$ is true,
(b) S believes that $p$
(c) S's belief that $p$ has property X (where having property X does not entail that the relevant belief that $p$ is true)

Let us look at Craig's version of the argument. Craig maintains that for (c) to be a plausible candidate for the missing link between true belief and knowledge it must always be the case that if S's belief that $p$ has X then it is highly probable that $P$ is true. As Craig notices the last statement is of the form: "E entails: It is highly probable that C."[2] However, unless E entails C statements of this form are bound to be false. For in that case there will always be a condition D such that when conjoined with E it is not highly probable that C is true. Applied to knowledge, that means that unless the property X that is supposed to turn true belief into knowledge is factive, there is always a further condition such that if conjoined with X it is no longer highly probable that the proposition believed is true. There is reason to believe, then, that unless X is factive, X cannot be a plausible candidate for the missing link between true belief and knowledge.

---

[2] Craig (1990), p. 50.

Or, in other words, there is reason to believe that no non-factive condition can be a plausible candidate for the missing link between true belief and knowledge.[3]

Zagzebski's argument improves on Craig's version of it in that it does not even rest on the assumption that any plausible candidate for the missing link between true belief and knowledge must render it highly probable that the subject's belief is true. Here is Zagzebski's version:

> It really does not matter how the particular element of knowledge in addition to true belief is analysed. As long as there is a small degree of independence between this other element and the truth, we can construct Gettier cases by using the following procedure: start with a case of false belief [that has property X, CK] … The falsity of the belief will not be due to any systematically describable element in the situation, for if it were, such a feature could be used in the analysis of the components of knowledge other than true belief, and the truth would be entailed by the other components of knowledge, contrary to the hypothesis. The falsity of the belief is therefore due to some element of luck. Now emend the case by adding another element of luck, only this time an element which makes the belief true after all. The second element must be independent of the element [of property X, CK] so that [property X, CK] is unchanged. The situation might be described as one element of luck counteracting another. We now have a case in which the belief [has property X, CK], the belief is true, but it is not knowledge. The conclusion is that as long as the concept of knowledge closely connects [the component given by property X, CK] and the truth component, but permits *some* degree of independence between them, true belief [that also has property X, CK] will never be sufficient for knowledge.[4]

Finally, Floridi argues that in order to solve the Gettier problem one has to solve the problem of ensuring the successful coordination of the components of knowledge—that is, on the one hand, truth, and, on the other, property X. Floridi goes on to show that, where X is construed as non-factive, the problem of the coordination of the components of knowledge is equivalent to the "coordinated attack" problem. Since it can be

---

[3] Cf. Craig (1990), pp. 50-52. It is noteworthy that Craig also maintains that if X is factive, then X is most likely to be too strong to be necessary for knowledge and that therefore the prospects for a traditional analysis of knowledge are very dim indeed. However, Craig acknowledges that it cannot be concluded that all attempts at finding the anti-Gettier condition will not work. Craig is aware that there are some candidate epistemic conditions that entail the truth of the proposition believed—such as the sensitivity condition (as defended by Dretske (1970), (1971) and Nozick (1981)), causal conditions (as defended by Goldman (1967)), and what Craig calls the "no-false-lemma" condition (as defended by Harman (1973) and Lehrer (1965)). Craig deals with them individually to complete his argument (pp. 72, 77-81). It is somewhat ironic that that Craig's own approach to the theory of knowledge, if further pursued, yields a factive condition on knowledge that we now have seen to be a promising candidate for dealing with Gettier cases.

[4] Zagzebski (1994), p. 69.

demonstrated in epistemic logic that the "coordinated attack problem" is unsolvable, it follows that the Gettier problem is logically unsolvable, too.[5]

It can easily be seen that the safety condition on knowledge avoids this sort of problem. After all, we have already seen that safe belief is factive.[6] So, any worry that may arise for conceptions of knowledge that attempt to deal with Gettier cases by placing a non-factive condition on knowledge will not arise for conceptions of knowledge that attempt to solve the Gettier problem by placing a safety condition on knowledge. Again, the safety-based conception of knowledge as developed from the minimalist framework seems to have an edge over a significant number of proposed solutions to the Gettier problem that can be found in the Gettier literature.


## 3 Minimalist Safety Safe Home?

We have seen that there is reason to believe that minimalist safety spells out the sense in which knowledge excludes luck, and that, given that this is so, we may expect it to explain our ignorance in cases of belief formed in unreliable ways, lottery cases and Gettier cases—caveat: provided luck really is the source of ignorance in these cases. Moreover, as regards the notorious Gettier cases, there is reason to believe that minimalist safety has an advantage over the various proposals in the Gettier literature in that by being derived from a plausible development of the anti-luck platitude it tackles the Gettier problem at its roots rather than at the surface. In addition, since safe belief is factive, it does not fall prey to a very strong argument against the prospects of a certain type of condition—that is, any non-factive condition—to be the anti-Gettier condition. Does that mean that we have brought minimalist safety safely home?

---

[5] Cf. Floridi (2004).
[6] In chapter IV, section 4.3.

Unfortunately, and perhaps somewhat surprisingly, the answer to this question is, I submit, 'no'. As I am about to argue, there are cases in which a subject safely believes a proposition whilst being Gettiered. The existence of such cases naturally exerts some pressure on minimalist safety. After all, on the face of it, in order to show that minimalist safety does not fail to meet the expectations it now has to be shown that luck is not the ground for the Gettiered subject's ignorance. It will come as no surprise that I will try to defend minimalist safety from the challenge such Gettier cases pose. However, instead of taking the challenge head-on, I will outline a way to sidestep it. To be more precise, I will argue that reflection on the case (and on similar non-epistemic cases) gives us strong reason to believe that Pritchard's modal conception of luck, at least in its present form, fails. I will also propose a slightly weaker version of the modal conception of luck and will argue that once Pritchard's conception of luck is replaced by the slightly weaker version, it can still be argued—along the lines outlined in the last chapter—that safe belief is requisite for knowledge. At the same time, the weaker version will not licence the argument from the safety condition to the veritic luck principle. Given that this is so, we need no longer expect the safety condition to fully capture the sense in which knowledge excludes luck. Accordingly, we need no longer expect it to explain our intuitions in all Gettier cases in which case the pressure exerted upon minimalist safety will evaporate. To begin with, let us now turn to the problematic Gettier case in which the subject also has a safe belief.

The relevant Gettier case is yet another variation of Russell's grandfather-clock case. Recall that in the last chapter I used a variation of the case to argue against a certain way of stating the safety condition on knowledge—*viz.* one according to which one must avoid false belief at nearby possible worlds at which one forms one's belief in the same way as in actuality. In that variation of the case, I come down the stairs, read the clock and thereby acquire a true belief that it's 8.22. Since I have the ability to read

the clock and since the clock works just fine, intuitively, my belief counts as knowledge. However, my powerful arch-nemesis who for some reason wants me to have a belief that it's 8.22 when I come down the stairs would have set the clock to 8.22 at nearby-worlds at which I do not come down the stairs at 8.22. Since there is a wide range of such nearby worlds while, at the same time, at those worlds I continue to form a belief that it's 8.22 in the same way as in the actual world, the relevant version of the safety condition turned out to be too strong.

Now consider the following variation of the case: Again, my arch-nemesis is in the background. This time he has an interest that I form a *true* belief that it's 8.22 when I come down the stairs. Again, he is prepared to do whatever it may take in order to ensure that I acquire a true belief that it's 8.22 when I come down the stairs. (My arch-nemesis's life might be a bit more difficult here. Since we are concerned with a conceptual claim, however, my arch-nemesis may have means available to attain his goal that one could imagine only in one's wildest dreams. For instance, my arch-nemesis may be an evil demon who could make me go downstairs at precisely the right point in time.) However, my arch-nemesis is still lazy. He will act only if I do not come down the stairs at 8.22 of my own accord. Suppose, as it so happens, I do come down the stairs at 8.22. My arch-nemesis remains inactive. I form a belief that it's 8.22. It is 8.22. However, unbeknownst to me (and to my arch-nemesis, suppose), the clock has stopped working exactly twelve hours ago. Intuitively, I take it, I do not know that it's 8.22. After all, I cannot come to know what time it is if the only information I have comes from a stopped clock.[7] I am Gettiered.

---

[7] It is, of course, not in general impossible to come to know what time it is by looking at a stopped clock. If one knows that one's arch-nemesis will ensure that one comes to form a true belief about the time upon reading the clock, then upon reading the clock, one can come to know what time it is. However, here one has additional information. It remains true that if the only relevant information one has is information from a stopped clock, intuitively, one's information is not good enough to give one knowledge what time it is.

At the same time, however, there is reason to believe that my belief is safe. After all, at all nearby possible worlds at which the relevant initial conditions for the formation of my belief are held fixed—that is, at which my arch-nemesis is still prepared to ensure that I come down the stairs at 8.22, at which he still has the power to do so etc.—I will come down at 8.22 and thus form a true belief that it's 8.22 (by looking at the stopped clock). Notice that the move that got the defender of safety off the hook in the first variation of the case won't work here. Even if at the time of belief-formation it has become part of the relevant initial conditions that my arch-nemesis remains inactive, it is still true that at all nearby possible worlds at which the relevant initial conditions for the formation of my belief are held fixed, it is 8.22. After all, we may assume that, for all nearby possible worlds, my arch-nemesis remains inactive if and only if I come down the stairs at 8.22 of my own accord. So, the set of nearby worlds at which my arch-nemesis remains inactive is the set of worlds at which I come down the stairs at 8.22 of my own accord. So, even if it has become part of the relevant initial conditions for the formation of my belief that my arch-nemesis remains inactive, at all nearby worlds at which these relevant initial conditions are held fixed, it is still the case that it is 8.22 when I come down the stairs. Since it is also part of these relevant initial conditions that the clock has stopped at 8.22, it follows that I will still form a true belief that it is 8.22 by reading a stopped clock. So, my belief is safe. Since I cannot come to know what time it is by reading a stopped clock, however, my belief is still Gettiered. Thus, I have a Gettiered belief that is also safe. Thus, safety cannot be the anti-Gettier condition.

Now, one might object that even if it is true that at all nearby possible worlds at which the relevant initial conditions are held fixed I come down the stairs at 8.22, it need not be the case that at all nearby possible worlds I avoid forming a false belief. After all, one might want to point out, the clock might easily have stopped a minute

earlier or later in which case I would have acquired a false belief. Notice, however, that the objection can be dealt with rather easily. We may construe the case in such a way that if the clock stops at all, it is bound to stop at 8.22. Maybe there is a weakness in the construction such that there is an outside chance that the clock stops. At the same time, the weakness is such that if it stops at all, it stops at 8.22. So construed, there are no nearby possible worlds at which the relevant initial conditions for the formation of my belief are held fixed at which the clock stops at any other time than 8.22. Still I do not come to know that it's 8.22 when reading the clock. I cannot acquire knowledge of the time if I base my belief only on information from a stopped clock. My belief is still Gettiered.[8]


## 4  Pritchard's Modal Conception of Luck Reviewed

The existence of a case in which a subject has a safe but Gettiered belief exerts pressure on minimalist safety. On the face of it, in order to show that minimalist safety does not fall short of what may be expected from it we need to show that in the problematic Gettier case my ignorance is not grounded in luckiness of my belief. I have already anticipated that I am not going to take the challenge posed by the Gettier case head-on. Rather, as I have also indicated, I will argue that reflection on the case (and on relevantly similar non-epistemic cases) provides us with strong reason to believe that Pritchard's modal conception of luck, at least in its present form, fails and that, in consequence, we need no longer expect minimalist safety to be the anti-Gettier condition.

---

[8] I provide another Gettier case of this type and a recipe how to construct such cases in the next chapter (section 3.2).

## 4.1   A Diagnosis of the Grounds of Ignorance in the Problematic Gettier Case

I will start off my argument by issuing what seems to me to be a plausible diagnosis of the problem. To get the diagnosis off the ground let us have another look at the consideration that suggested that my belief that it is 8.22, even if safe, does not qualify as knowledge: I cannot come to know what time it is by looking at a stopped clock. An attractive explanation of why my belief does not qualify as knowledge appeals to the notion of a cognitive ability. According to this explanation, the reason why I lack knowledge is that the truth of my belief about the time was not *ensured through* my relevant cognitive abilities—in this case, presumably, an ability to read the clock. Consider, by way of analogy, a basketball player—let's call him Larry—who has the ability to score from within a range of three metres. When he shoots from a position well outwith his range and scores, his success will not be *ensured through* his ability. Analogously, my cognitive ability to read the clock cannot ensure the truth of a given belief about the time if the clock I am reading is stopped. In this way, according to the current explanation, my ignorance is grounded in the fact that the truth of my belief has not been ensured through the exercise of my relevant cognitive abilities.

## 4.2   Interlude: Minimalist Safety off the Hook?

One might be inclined to believe that the defender of minimalist safety is already off the hook. After all, according to the above diagnosis, my ignorance is grounded in the fact that the truth of my belief was not ensured through the exercise of my relevant cognitive abilities. Now, one might think that if this is so, then my ignorance is not grounded in the luckiness of my belief. If it isn't, however, then, by the caveat of section 1, we need not expect minimalist safety to predict ignorance in this Gettier case.

Unfortunately, however, things are not that straightforward. For suppose the following entailment holds:

*Ability Luck Entailment (ALE)*   If one succeeds but not through the exercise of one's relevant abilities, then, relative to one's performance, it is a matter of luck that one has succeeded.

Since, according to the above diagnosis, I do not hit upon the truth through the exercise of my relevant cognitive abilities, if ALE holds, relative to my cognitive performance, it will also be a matter of luck that I have acquired a true belief. Given that this is so, however, it is still possible and indeed quite plausible that my ignorance is grounded in the fact that my belief is too lucky to count as knowledge.

Notice that, for now, I do not wish to endorse ALE.[9] The only point I want to make is that, insofar as ALE is not entirely implausible, it is not clear that the defender of minimalist safety is off the hook merely in virtue of the proposed diagnosis of my ignorance. Since in what follows I will argue that there is another way to get the defender of minimalist safety off the hook, there will be no need to settle this issue just now.

## 4.3   Pritchard's Modal Conception of Luck in Trouble

After this short interlude, let me now return to my argument. The next step in the argument focuses on the relation between cases in which one succeeds but not though the exercise of one's relevant abilities on the one hand and luck on the other hand. There is reason to believe that there is a wide range of cases in which one succeeds but not through the exercise of one's relevant abilities, while, at the same time, *relative to one's relevant abilities*, it is a matter of luck that one succeeds. By way of illustration,

---

[9] However, I will defend a slight variant of ALE in the next chapter (section 2).

consider once more the case of Larry who scores reliably only from within three metres. In virtually all situations, when he takes a shot from a position well outwith his range he will miss. Now, suppose he shoots from a distance of seven metres and scores. Certainly, *relative to his shooting abilities*, it is a matter of luck that he has scored. And similarly in the epistemic case: It is plausible that in the problematic Gettier case in which I hit upon the truth but not through the exercise of my relevant cognitive abilities—that is, my ability to read the clock—*relative to my cognitive abilities*, it is also a matter of luck that I have hit upon the truth.

With these observations about the relation between cases of success but not through the exercise of one's relevant abilities on the one hand and luck on the other in play, it is now time to move on to the crucial point—*viz.* that there is reason to believe that Pritchard's modal conception of luck fails. In order to show that this is so, I will argue that, contrary to Pritchard's modal conception of luck, the type of luck at issue in these cases—that is, luck relative to one's relevant abilities—can obtain even if the event at issue does not exhibit the modal instability that, according to Pritchard's modal conception of luck, it must exhibit. In this regard, let us first go back to the case of Larry who shoots from well outwith his range and scores. Relative to his shooting abilities, he is lucky to have scored. Or, in other words, the event of his scoring is lucky relative to his shooting abilities. If the modal conception of luck held, that would mean that the event of his scoring does not occur at a wide range of nearby possible worlds at which the relevant initial conditions for that event—he still takes a shot from the same distance, he still has the same shooting abilities etc.—are held fixed. But now notice that even if this is *typically* the case, one can easily augment the case in such a way that Larry will succeed at such nearby possible worlds, too. Perhaps a suitable piece of metal in the ball and a strong magnet in the hoop secure that he scores not only at the actual world but also at nearby possible worlds, or, if that is not good enough, let a benevolent

demon do the job. Even so, relative to his shooting abilities, it is a matter of luck that he scores. Similarly in the epistemic case: Relative to my cognitive abilities, it is a matter of luck that I hit upon the truth. At the same time, my powerful arch-nemesis ensures that I avoid false belief at nearby possible worlds. What these considerations suggest is, of course, that if there is such a thing as luck relative to one's relevant abilities, as there is excellent reason to believe there is, then it is not always required that a lucky event fails to obtain at a range of nearby possible worlds. If so, the modal conception of luck, at least in its present form, fails.

## 4.4  A Slightly Weaker Modal Conception of Luck

But now remember just how plausible we found the modal conception of luck to be: On the one hand, it allowed us to derive the highly plausible safety condition on knowledge. On the other hand, it allowed us to explain our intuitions concerning luck in a wide range of cases. And, what's more, it was also continuous with and indeed served to explain some results of scientific research on luck. In particular, recall that there is excellent scientific evidence for the claim that counterfactual thinking is of central importance to attributions of luck. If we discard the modal conception of luck on the basis of a handful of armchair counterexamples, are we not a bit too quick in dismissing an apparently valuable tool for our theory of knowledge and for our understanding of the foundations of some solid scientific research? And as regards the latter in particular, ought we not rather welcome attempts of philosophers to devise theories that are continuous with scientific research in the way that Pritchard's conception of luck is? And if such a theory runs into problems that some philosopher has cooked up from the armchair, ought we not reject or at least reconsider some of the considerations that led to the problems rather than give up the theory?

I won't even attempt to answer these questions. Fortunately, however, I don't have to. For even if we take the present considerations to indicate that the modal conception in its present form is false, there is a slightly weaker version of it that can avoid the problems Pritchard's version faces whilst retaining its virtues. To be more precise, there is a weaker modal conception of luck such that, first, it still holds out the promise of explaining the results of the relevant scientific research, while, second, the safety condition can still be derived from it and the veritic luck principle, and, third, we can avoid the problem posed by the Gettier case and the case of Larry the basketball player. On top of all this, if we replace Pritchard's modal conception of luck by the weaker one, the argument from the safety condition to the veritic luck principle no longer goes through. If so, we need not expect minimalist safety to fully capture the sense in which knowledge excludes luck. Hence, whatever pressure may have been exerted on minimalist safety by the problematic Gettier case will evaporate.

Let's look at how all this can be done. To begin with, recall that the modal conception of luck as proposed by Pritchard takes the following form:

(LE)    An event is lucky *iff* it obtains in the actual world but does not obtain in a wide class of nearby possible worlds in which the relevant initial conditions for that event are the same as in the actual world.[10]

LE does well on explaining the results of the relevant scientific research and in conjunction with the veritic luck principle it also serves to derive the safety condition on knowledge. At the same time, since LE has it that if an event is lucky, then it must exhibit a certain modal instability, certain cases in which a person's success is lucky relative to his abilities pose a problem for LE.

My suggestion now is to weaken Pritchard's version of the modal conception of luck. The weakened version I have in mind takes the following shape:

---

[10] Pritchard (forthcoming*a*), p. 3.

(LE*)   An event is lucky *if* it obtains in the actual world but does not obtain in a wide class of nearby possible worlds in which the relevant initial conditions for that event are the same as in the actual world.

LE* is weaker than LE in that failure to obtain at a wide range of nearby possible worlds is a sufficient but not a necessary condition for an event's being lucky. But how does LE* avoid the vices of LE whilst retaining its virtues? To begin with, notice that since, by the lights of LE*, there is still a close conceptual connection between certain modal statements and statements that a certain event is lucky, LE* still holds out the promise of explaining a variety of phenomena associated with luck that LE did explain such as, for instance, why scientific research suggests that counterfactual thinking is of importance to attributions of luck. Moreover, there is reason to believe that the argument from the veritic luck principle to the safety condition on knowledge will go through even if LE is replaced by the weaker LE*. To see why this is so, notice that the contraposition of LE* is

(CLE*) If an event is not lucky then it does not only obtain in the actual world but also there is no wide class of nearby possible worlds at which the relevant initial conditions for that event are the same as in the actual world and at which it does not obtain.

Accordingly, the conditional for an event that is not at all lucky, from which luck is excluded will take the following shape:

(NLE*) If luck is excluded from an event then it does not only obtain in the actual world but also there is no class, however small, of nearby possible worlds at which the relevant initial conditions for that event are the same as in the actual world and at which it does not obtain.

It can easily be seen that the safety condition can be derived from the veritic luck principle and NLE*. So, our argument for the safety condition on knowledge can be recovered. In this way, LE* can retain the virtues of LE.

At the same time, there is reason to believe that in replacing LE by LE* we can avoid the problems that LE encountered. To begin with, since modal instability is not a necessary condition for luck, it is consistent with LE* that an event exhibits modal stability—that is, that it occurs not only at the actual world but also at nearby possible worlds at which the relevant initial conditions for that event are held fixed—whilst being lucky. In consequence, cases of success that is lucky relative to an agent's relevant abilities in which the agent nevertheless succeeds at all nearby possible worlds at which the relevant initial conditions are held fixed do not pose a problem for LE*. Or, to put the point in a slightly different way, it is consistent with LE* that there be other ways in which an event may be lucky, ways, that is, that do not involve modal instability of the lucky event. One way in which the event of a success may be lucky whilst not involving modal instability may be that the success is not ensured through the exercise of one's relevant abilities. In this way, LE* can avoid the problem cases of luck relative to one's relevant abilities posed for LE.

Finally, if we replace LE by LE* the pressure the problematic Gettier case exerted on minimalist safety evaporates. To see why this is so, recall that the pressure was on minimalist safety in the first place because, first, we may expect that minimalist safety captures the sense in which knowledge excludes luck and, second, it is plausible that the reason why Gettiered subjects lack knowledge is that their true beliefs are afflicted by luck. Recall, furthermore, that the reason why we may expect that the safety condition captures the sense in which knowledge excludes luck is, first, that there is reason to believe that the veritic luck principle is a plausible development of the anti-luck platitude, that veritic luck is the type of luck at issue in the anti-luck platitude, and, second, that, given LE, the veritic luck principle and the safety condition are mutually derivable. Now, notice that if we replace LE by LE*, the veritic luck principle and the safety condition are no longer mutually derivable. To be more precise, while we can

derive the safety condition from LE* (or rather NLE*) and the veritic luck principle, we will not get the veritic luck principle from the safety condition and LE* (or NLE* for that matter). In consequence, there is no longer reason to expect that minimalist safety fully captures the sense in which knowledge excludes luck. If there isn't, however, then we need no longer expect minimalist safety to deal with all Gettier cases. The fact that there are cases of Gettiered safe belief no longer exerts pressure on minimalist safety.

In conclusion, then, it has been argued that an appealing way to explain why in the problematic Gettier case my belief does not qualify as knowledge is that I have not ensured the truth of my belief through the exercise of my relevant cognitive abilities. There is a wide range of cases—both epistemic and non-epistemic—in which one is successful but not through the exercise of any of one's relevant abilities, while, at the same time, one's success is lucky relative to one's relevant abilities. Since, in some of these cases, one may be successful at all nearby possible worlds, we have reason to believe that the modal conception of luck as proposed by Pritchard fails. However, we can recover the close connection between lucky events and their failure to occur at nearby possible worlds by weakening Pritchard's modal conception of luck. Instead of holding that an event is lucky *if and only if* it fails to occur at a certain range of nearby possible worlds, we maintain that it is lucky *if* it does so. In this way, we may still hope to be able to explain the relevant scientific results concerning the importance of counterfactual thinking to attributions of luck, we can hold onto the argument from the veritic luck principle and the (weakened) modal conception of luck to the safety condition, whilst avoiding commitment to the claim that the resulting minimalist safety condition must be able to explain our intuitions in all Gettier cases.

# VI  Knowledge, Achievements, and Abilities

In the last chapter I provided a case in which a subject has a safe belief whilst being Gettiered. It was argued that a plausible diagnosis of the subject's ignorance is that the truth of his belief is not ensured through the exercise of his relevant cognitive abilities. The question I would like to address now is how this diagnosis is to be fitted into the minimalist framework. A straightforward way of fitting the diagnosis into the minimalist framework is by adopting the following additional platitude:

> The Ability Platitude: If one knows that $p$, then the truth of one's belief that $p$ must have been ensured through the exercise of one's relevant cognitive abilities.[1,2]

Now one might wonder whether the ability platitude really is intuitive and lightweight enough to count as a genuine platitude of the kind suitable for the minimalist framework. While some have claimed that we have an ability intuition about knowledge[3]—in which case some platitude relating knowledge and abilities may be suitable for the minimalist framework—one might be worried that the above ability platitude looks much more like a full-blown substantive condition on knowledge. While I am inclined to accept even the above ability platitude (henceforth, again, simply 'ability platitude') as a genuine platitude about knowledge, in order to allay any worries a critic may have, I will start out with an argument that derives a slight variant of the condition on knowledge stated in the ability platitude from another platitude and some

---

[1] A similar condition for knowledge has been defended by Sosa (e.g. in Sosa (1988*a*), (1988*b*) and (1991*a*)), Greco (e.g. in Greco (2003), (2007), (forthcoming)) and Riggs (2002) and (forthcoming). A closely related view has been developed in Zagzebski (1996) and (1999).
[2] Sometimes the condition on knowledge at issue in the ability platitude is stated as follows: One knows that $p$ only if one believes the truth *because of* the exercise of one's relevant cognitive abilities. For present purposes, I consider this alternative a notational variant of the condition at issue in the ability platitude. (Greco, for instance, appears to do so as well: In (forthcoming) he alternates between the locutions 'through' and 'because' in the statement of his version of the relevant condition.) What both conditions must be contrasted with is success *and*—that is, but not through/because of—the exercise of a set of relevant abilities (cf. Greco (forthcoming), p. 14).
[3] Cf. Pritchard (2007), p. 15.

independently plausible principle. This argument will be useful for those who are unwilling to count the ability platitude as a genuine platitude about knowledge in that the condition derived explains the lack of knowledge in the problematic Gettier case and thus allows fitting last chapter's diagnosis into the minimalist framework without adopting the ability platitude. At the same time, however, it will also be useful for those who are inclined to accept the ability platitude in that, as I will argue in due course, the condition so derived can be viewed as sharpening the ability platitude in a way that will allow its defenders to solve one of the most pertinent problems that it may otherwise encounter.

# 1 The Achievement Platitude and the Achievement Ability Principle

It seems to me that the following platitude will eventually allow us to derive a condition very similar to the one at issue in the ability platitude:

The Achievement Platitude: Knowledge is a cognitive achievement.

The achievement platitude enjoys a great degree of intuitive plausibility. The intuitive plausibility can be backed up by the observation that paradigm types of knowledge—such as perceptual knowledge, introspective knowledge, inferential knowledge and memory-based knowledge—are quite clearly also cognitive achievements.[4] In this way the achievement platitude appears to be just the kind of principle that is suitable for the minimalist framework. I will not provide any further direct support for the achievement platitude. However, the arguments I will adduce in this chapter will suggest that adopting the achievement platitude will give us significant theoretical benefits. In this

---

[4] I will discuss testimonial knowledge, which might also be regarded as a paradigm case of knowledge, in due course (section 5.2).

way I hope that these arguments will provide if not direct support then at least further indirect support for the achievement platitude.

Now, how do we get from the achievement platitude to the condition stated in the ability platitude? Well, one way of getting there is by arguing that achievement requires success through the exercise of ability. In this vein, I will now provide some support for the following principle:

*Achievement Ability Principle (AAP)*    x is an achievement only if x is a success obtained through the exercise of a set of relevant abilities.[5]

The following considerations provide at least *prima facie* evidence for AAP. To begin with, it is plausible that only successes qualify as achievements. Thus, for instance, even a professional juggler will secure the type of achievement involved in juggling only if he juggles successfully. If he tries and fails he will not have secured this kind of achievement. Similarly, a police dog will secure the achievement involved in locating the drugs only if he successfully locates the drugs. If he tries and fails he will not have secured this achievement. In this way there is *prima facie* evidence that achievement requires success.

Notice next that only agents can secure achievements. By way of evidence, it is plausible that only humans and non-human animals with a certain degree of sophistication can secure achievements. Inanimate objects and certain animate objects with insufficient sophistication are not in the ballpark of achievements.[6] Given that this is correct, there is reason to believe that achievement requires success by an agent.

---

[5] A similar thesis has been defended in Riggs (2002) and (forthcoming). Riggs claims that one secures an achievement only if one's abilities and intentions are sufficiently creditworthy for bringing about the success ((forthcoming), p. 26). Furthermore, he makes it clear that this will be the case only if the success is causally due to one's abilities ((forthcoming), p. 11).

[6] It is hard to see, for instance, how an amoeba could secure an achievement. At the same time, it is equally hard to see how an amoeba could be regarded as an agent.

Moreover, there is also reason to believe that achievement requires success by an agent through the exercise of a set of abilities. Thus compare Daniel Barenboim playing the *Moonlight Sonata* with a person who is in a frenzy in which he has no control over his body but which causes his fingers to randomly hit the keys of the piano he is sitting at. As it so happens he hits the keys in the sequence of the notes of the *Moonlight Sonata*. Alternatively, compare Hemmingway writing *For Whom the Bell Tolls* with a monkey at a typewriter who chances to hit the keys in such a way as to produce the sequence of characters of *For Whom the Bell Tolls.* It is plausible that Barenboim and Hemmingway secure the relevant musical and literary achievements while the person in frenzy and the monkey don't. At the same time, Barenboim and Hemmingway succeed in their respective performances through the exercise of their musical and literary abilities, while in the cases of the person in frenzy and the monkey no ability appears to be involved.

Finally, it would seem that achievements require success not just through the exercise of any old set of abilities but rather through the exercise of a set of *relevant* abilities. Thus compare our basketball player, Larry, who takes a shot from within his range and scores through the exercise of his shooting abilities with another player, Harry, who is to take a shot but is somehow tricked into believing that he is in defence. In fact he believes that he must now block an opponent's shot. In order to do so, he jumps up and stretches out his arm. Since—at least in Harry's case, we may suppose—the blocking movement is the same as the shooting movement, he ends up throwing the ball and, lo and behold, it goes in. It would seem that the Harry's athletic success is secured through the exercise of a set of abilities. After all, he not only has shooting abilities but also blocking abilities and it is the exercise of his blocking abilities that ensures his success. However, while Larry has certainly secured the achievement involved in scoring, it would be inadequate to credit Harry with such an achievement.

And the reason why it would be inadequate to do so is that he has exercised the wrong abilities.[7] These considerations, then, provide *prima facie* support for AAP.

Notice that AAP contains the notorious term 'relevant'. While I have no fully-fledged account of what makes abilities relevant and will therefore have to rely on intuition to a certain extent, I take it that it is plausible that one possesses a set of relevant abilities only if one possesses a set of abilities that allows one to succeed reliably in one's circumstances. Now one might think that this is hardly an interesting claim. After all, it is widely agreed that there is a reliability condition even on ability possession *simpliciter*, as it were. For instance, both Sosa (1991*a*) and Greco (2007*a*) have argued for such a reliability condition.[8] Roughly, Sosa and Greco's idea is that one possesses an ability relative to a set of circumstances only if one succeeds reliably in these circumstances.[9] While the present proposal agrees with Sosa and Greco's proposal that there is a relation between ability and reliability, there is a significant difference between the two. While Sosa and Greco relativise ability possession *simpliciter* to circumstances, the present proposal only relativises possession of a set of *relevant* abilities—that is, relevant to securing a given achievement—to circumstances. It is thus consistent with my proposal but not with Sosa and Greco's that one may possess an ability in a set of circumstances in which it does not reliably lead one to success. There

---

[7] One might wonder whether it is really necessary to wheel in the notion of a set of relevant abilities to explain why Harry fails to secure an achievement. It might be thought that Riggs's intention condition on achievement (cf. p. 162, n. 5) will do the job. After all, when Harry comes to believe that he must block his opponent's shot, he no longer intends to score. So, he will fail Riggs's intention condition on achievement. However, it would seem that we can easily amend the case so that Riggs's intention condition won't do the job. For instance, we may suppose that Harry intends to score all along but, say, nervousness causes him to exercise his blocking abilities rather than his shooting abilities. If he does and scores, then he still won't have secured the relevant achievement. At the same time, he satisfies Riggs's intention condition.

[8] Notice, however, that while Greco maintains that there is a reliability condition on possession of *any* kind of ability ((2007*a*), p. 7), Sosa maintains only that there is a reliability condition on possession of *cognitive abilities* or intellectual virtues ((1991*a*), p. 286). Since the problematic examples I am about to provide are non-cognitive in nature, they are a cause of concern in the first instance for Greco's account of abilities. However, I think that there are parallel cases which suggest that we had better not place a reliability condition of the kind envisaged by Sosa and Greco on possession of cognitive abilities *simpliciter* either. I will adduce one such case in due course.

[9] Greco and Sosa interpret the relevant notion of reliability in different ways: While Sosa spells out the notion of reliability in probabilistic terms ((1991*a*), p. 286), Greco gives it a modal spin ((2007*a*), p. 7).

is excellent reason to believe that this is an advantage of my view over Sosa and Greco's. When our basketball player, Larry, shoots in high winds, for instance, he will in all likelihood be in circumstances in which he is highly unreliable in securing success. At the same time, it is plausible that when taking a shot in high winds Larry may retain and exercise the abilities that enable him to secure the achievement involved in scoring in low winds. To see why this is plausible, notice, first, that he must certainly exercise *some* abilities even when shooting in high winds. Otherwise it will be hard to make sense of the idea that the ball so frequently heads at least roughly towards the hoop. Now suppose, furthermore, that the circumstances are such that there are frequently gusts of strong wind around the hoop. Intuitively, Larry will secure achievements when his shots are not tampered with by wind but not when they are so tampered with. A plausible explanation of why this is so is that when the wind is not blowing, he is in circumstances in which he is reliable, while, when the wind is blowing, he is in circumstances in which he is unreliable. On Sosa and Greco's proposal it is not clear which ability Larry has exercised until the ball comes within the region of the hoop. If there is no gust of wind, he will have exercised one set of abilities. If, however, there is a gust of wind, he will have exercised a different set of abilities. That, however, seems quite implausible. For that reason we had better not construe ability possession *simpliciter* as relative to circumstances in the way envisaged by Sosa and Greco.[10]

Fortunately, the present proposal can avoid this problem. The only claim that is made is that possession of a set of abilities relevant for securing a certain achievement is relative to circumstances in the sense that such a set of abilities must lead to reliable

---

[10] This is not to say, however, that once one has an ability one has it in all circumstances. Thus, it would seem that when losing his shooting arm, for instance, Larry may also lose his ability to score. The only claim I wish to make here is that there are *some* circumstances in which some of one's abilities do not lead one to reliable success, while, at the same time, one retains and exercises them.

success in the circumstances in which it is exercised.[11] Now, this claim is much more plausible. After all, when Larry's shots are tampered with by wind, for instance, he will certainly not secure an achievement even on occasions on which he happens to be successful.[12,13]

An obvious consequence of the present conception of achievements is that whether or not one is in a position to secure a given achievement may depend heavily on (external) circumstances. To see why this is so, notice, to begin with, that whether or not the abilities one possesses and exercises will lead one to reliable success may vary dramatically from one set of circumstances to another. For instance, while we may expect an average basketball player—such as Larry—to have a set of abilities that will allow him to score reliably (from within a certain range) when playing indoors (or outdoors in low winds), we need not expect him to have a set of abilities that will allow him to score reliably (even from within the same range) when shooting in high winds. Of course, some players may even possess sets of abilities that allow them to score reliably in high winds. Only those players will be in a position to secure achievements in such circumstances. Moreover, some circumstances may render achievement (humanly) impossible. It may be humanly impossible to acquire a set of abilities that will lead one to score reliably on the moon. Alternatively, it would certainly be humanly impossible to acquire a set of abilities that will lead one to score reliably when

---

[11] Notice that the reliability condition on sets of relevant abilities is at best a necessary condition. For that reason in a given case, we may still have to rely on intuition as to what makes for a set of relevant abilities.

[12] One might wonder what the present proposal makes of the claim that there is a reliability condition on ability possession *simpliciter*. A plausible suggestion is that in order to possess a given ability there must be *some* set of circumstances in which one succeeds reliably. If one is not reliably successful in any set of circumstances, then one does not possess the ability at issue. This move will make sense of the idea that ability possession *simpliciter* requires reliability while also allowing that one can possess and exercise abilities in circumstances in which one does not secure success reliably.

[13] A straightforward consequence of the present conception of abilities is that abilities are fallible. If one can possess and exercise an ability in circumstances in which one is highly unreliable, then one must also be able to possess and exercise an ability and fail.

blindfolded while, at the same time, shooting at a target that is moving randomly (and otherwise imperceptibly).

With AAP in play, we can now use the achievement platitude to derive the following condition on knowledge: One knows only if one secures cognitive success through the exercise of a set of relevant cognitive abilities. Since it is plausible that the relevant kind of cognitive success at issue here is true belief, we get:

*The Ability Condition*   One knows that *p* only if the truth of one's belief that *p* has been ensured through the exercise of a set of relevant cognitive abilities.[14]

Notice that the ability condition will also do the job the ability platitude outlined at the outset of this chapter promised to do: It will allow us to fit the diagnosis of the problematic Gettier case discussed in the last chapter into the minimalist framework. After all, the ability condition claims that knowledge requires exactly what, according to the diagnosis of the problematic Gettier case, I lack when I form a belief about the time by looking at a stopped clock. I therefore fail to satisfy the ability condition. Since, provided that AAP holds[15], the ability condition is entailed by the achievement platitude, it follows that I also fail to satisfy the achievement platitude. In this way, those who are unwilling to accept the ability platitude as a genuine platitude about

---

[14] Notice that the ability condition is very similar to the ability platitude. Indeed, the only noticeable difference between the two is that in the ability condition the notion of a set of relevant abilities is given a more precise interpretation in the following respects: First, the notion of a set of relevant abilities is to be understood in the same way as the notion of a set of relevant abilities at issue in other achievements; second, having such set of relevant abilities requires having a set of abilities that, in the circumstances, lead one to reliable success; and, third, whether or not one possesses such a set of relevant abilities may depend heavily on circumstances. As I will argue in due course, interpreting the notion of relevant ability at issue in the ability platitude in this way will allow the defender of the ability platitude to avoid one of the most serious problems he may otherwise encounter. For that reason the defender of the ability platitude will do well to accept the present interpretation of the notion of relevant ability as a development of the platitude. Moreover, if the ability condition identifies the sense of relevant abilities at issue in the ability platitude, those who are inclined to accept the ability platitude now have some reason to believe that the achievement platitude entails the ability platitude.

[15] For stylistic reasons I will henceforth sometimes leave this proviso implicit. Thus, when I say things like: "The achievement platitude entails the ability condition." I mean: "Provided that AAP holds, the achievement platitude entails the ability condition."

knowledge will now also be able to fit the diagnosis of the problematic Gettier case provided in the last chapter into the minimalist framework.

If this line of thought is sound, it follows that there are at least some Gettier cases in which the ability condition and the achievement platitude can explain why the Gettiered subject lacks knowledge. But now recall that it was argued in chapter III that there is excellent reason to believe that the crux with Gettier cases is luck and that we may therefore expect the anti-luck platitude to explain why Gettiered subjects lack knowledge. Given that this is so, however, we may wonder what exactly the relation between the achievement platitude and the anti-luck platitude is. Are the two platitudes independent of one another or is there some relation between them? If they are independent, do we need both of them to explain why subjects in Gettier cases lack knowledge? And what about the other types of case that motivated the anti-luck platitude? On the other hand, if there is a relation between the two, what exactly is it?

In what follows I will argue that there is reason to believe that the achievement platitude entails the anti-luck platitude. In order to do so, I will show, first, that the ability condition entails an anti-luck principle; second, that this anti-luck principle can be viewed as a development of the anti-luck platitude; and, third, that this anti-luck principle outperforms the veritic luck principle—which, recall, could also be viewed as a development of the anti-luck platitude—in explanatory power: it explains our intuitions in all cases in which the veritic luck principle may hope to do so but not vice versa. Given that all this is so, there will be reason to believe that the anti-luck principle which is entailed by the ability condition captures the sense of luck at issue in the anti-luck platitude. Since, additionally the achievement platitude entails the ability condition, there is reason to believe that the cognitive achievement involved in knowing excludes the type of luck at issue in the anti-luck platitude. That is to say, however, that there is reason to believe that the achievement platitude entails the anti-luck platitude.

## 2  The Ability Condition and Luck

To begin with, let us take a closer look at the relation between successes obtained through the exercise of a set of abilities relevant for achievement and luck. In this section I will provide evidence that such successes are luck-excluding in the following way:

*Relevant Abilities Luck Principle (RALP)*   If one succeeds through the exercise of a set of abilities relevant for an athletic, cognitive etc. achievement, then relative to one's athletic, cognitive etc. performance, it is not a matter of luck that one has succeeded.

How can RALP be supported? It looks as though the following considerations provide *prima facie* support for RALP. To begin with, RALP appears to be supported by a number of paradigm cases of success through the exercise of a set abilities relevant for the cognitive, athletic etc. achievement at issue. Suppose, first, that Larry scores through the exercise of a set of abilities that, in his circumstances, is relevant for the athletic achievement involved in scoring. It is plausible that, relative to his athletic performance, it is not a matter of luck that he has scored. Similarly, when Barenboim succeeds in playing the *Moonlight Sonata* or Jamie Oliver succeeds in preparing a Fillet Wellington through the exercise of a set of piano or cooking abilities that are relevant, in their respective circumstances, for the musical and cooking achievements at issue, relative to their respective performances as a piano player and chef, it is not a matter of luck that they have succeeded.

By way of contrast, let us now look at the ways in which it can be a matter of luck, relative to one's relevant performance, that one succeeds. It is plausible that there are three ways in which, relative to one's relevant performance, it may be a matter of luck that one succeeds. To begin with, one may perform whilst *not possessing* a set of relevant abilities and succeed. This will be the case, for instance, when Larry shoots

from outwith his range and scores. In this case he does not possess a set of abilities rendered relevant by his circumstances for the athletic achievement involved in scoring. Second, relative to one's performance, it may be a matter of luck that one succeeds if, whilst possessing a set of relevant abilities, one exercises the *wrong* set of abilities in one's performance. For instance, our other basketball player, Harry, will perform with the wrong set of abilities when he scores through exercising his blocking abilities rather than his shooting abilities. The third way in which, relative to one's performance, it may be a matter of luck that one succeeds is the following: One performs *badly* and succeeds. Thus suppose that a third basketball player, Gary, takes a shot in circumstances relative to which he possesses a set of relevant abilities. However, on this particular occasion he performs badly—his shot is off target, let us suppose. At the same time, circumstances conspire so that he is successful anyway—perhaps a gust of wind brings his bad shot back on target.[16] It seems that these three ways are the ways in which, relative to one's performance, it may be a matter of luck that one has succeeded. At the same time, it is plausible that in neither case one manages to ensure success through the exercise of a set of relevant abilities. Thus Larry, who does not even possess a set of abilities relevant for the achievement at issue, obviously cannot succeed through the exercise of such a set of abilities. Similarly, Harry, who performs with the wrong set of abilities, won't succeed through the exercise of a set of abilities relevant for the achievement at issue. After all, to exercise a wrong set of abilities is to exercise a set of abilities *irrelevant* for this achievement. Finally, Gary, who performs badly, won't succeed *through* ability. He needs circumstances to conspire in such a way as to lead him to success anyway.

---

[16] The example can easily be amended to show that even sets of relevant abilities are fallible. Just suppose that there is no gust of wind that brings the shot back on target. Since, on the present conception of abilities, abilities and sets of relevant abilities are fallible, the conception of knowledge that is based on this conception of abilities can adequately be labelled 'fallibilist'.

There is, then, some *prima facie* evidence that RALP holds; that securing success through the exercise of a set of relevant abilities entails that, relative to one's relevant performance, it is not a matter of luck that one has succeeded. Applied to the epistemic case, that means that if the truth of one's belief is ensured through the exercise of a set of relevant cognitive abilities, then, relative to one's cognitive performance, it is not a matter of luck that one has formed a true belief. In conjunction with the ability condition we can then derive the following principle:

*Knowledge Luck Principle (KLP)* One knows that *p* only if, relative to one's cognitive performance, it is not a matter of luck that one has formed a true belief.

According to KLP, knowledge precludes luck relative to one's cognitive performance. Since the achievement platitude entails the ability condition that means that the achievement platitude entails KLP. But now notice that, according to KLP, knowledge requires that luck be excluded from one's true belief. Or, to be more precise, KLP specifies a type of luck that must not afflict one's true belief if one's true belief is to qualify as knowledge. Given that this is so, however, KLP is in the ballpark of providing a sharpened version of the anti-luck platitude. As such, it competes with the veritic luck principle, which, it was argued, is also a development of the anti-luck platitude.[17] And, in fact, KLP and the veritic luck principle are very similar to one another. Indeed the only difference between the two is that KLP relativises luck to cognitive performance in a way in which the veritic luck principle does not. Yet the question remains which of the two principles is the more plausible development of the anti-luck platitude.

In the next section, I will argue that the scales tip in favour of KLP. To be more precise, I will argue that KLP explains all the phenomena the veritic luck principle

---

[17] Cf. chapter IV, section 2.1.

explains but not vice versa. Moreover, I will argue that it is plausible that the problem with the phenomenon that the veritic luck principle does not explain is that the subject's belief is too lucky to qualify as knowledge. If my argument is sound, there is reason to believe that KLP is the more plausible development of the anti-luck platitude. By the same token, we also have some reason to believe that the type of luck at issue in the anti-luck platitude is the type of luck at issue in the consequent of KLP and, since the achievement platitude entails KLP, that the achievement platitude entails the anti-luck platitude.

## 3   The Veritic Luck Principle v KLP

## 3.1   KLP Explains the Phenomena that Motivated the Veritic Luck Principle

Let us first turn to the question as to whether KLP succeeds in explaining all the phenomena that the veritic luck principle promises to explain. I will go through the types of case in the by now familiar order: first, cases of unreliably formed belief, second, lottery cases and, third, Gettier cases.

### 3.1.1   Unreliably Formed Belief

There is excellent reason to believe that in cases of unreliably formed belief KLP will explain our intuitions with equal success as the veritic luck principle. After all, if one forms a belief in an unreliable way, then there is nothing about the cognitive abilities one exercises that makes it more likely than not that one hits upon the truth. That is to say, there is nothing about one's cognitive performance that makes it more likely than not that one hits upon the truth. Given that this is so, however, if one hits upon the truth

despite having formed one's belief unreliably, relative to one's cognitive performance, it will be a matter of luck that one has done so. That means, however, that unreliable believers do not satisfy KLP. There is reason to believe, then, that KLP successfully explains one of the three phenomena that the veritic luck principle promises to explain. Let us move on quickly and ask whether KLP succeeds in explaining our intuitions in lottery cases with equal success as the veritic luck principle.

### 3.1.2 Lottery Cases

So, do people who believe that given tickets in fair lotteries won't win satisfy KLP? It seems to me that the answer to this question is 'no'. After all, it would seem that cognitive agents like us are constituted such that whenever they form beliefs that given tickets in fair lotteries won't win—in circumstances in which the lottery is in fact fair and will in fact be held etc.—whether or not their beliefs turn out to be true is at the end of the day down to luck. Given that this is so, however, then, no matter how well their cognitive performance, at the end of the day it will always be down to luck that such a belief turns out to be true. If so, then whenever cognitive agents like us form beliefs that given tickets in fair lotteries won't win, they will not satisfy KLP. KLP will be able to provide a successful explanation as to why our beliefs in lottery propositions do not qualify as knowledge as well.[18]

---

[18] It is noteworthy that Greco provides a different analysis of lottery cases. In Greco (2003), he explains why people who entertain beliefs in lottery propositions lack knowledge in the following way: He argues for a version of contextualism, a view, that is, according to which sentences of the form 'S knows that *p*' can express different propositions in different conversational contexts. Very roughly, Greco argues that lottery cases always create a conversational context in which sentences of the form "S knows that ticket x won't win the lottery" are false. I am making this point here only to put it to one side. For even if Greco's version of contextualism (or any other version of contextualism for that matter) does explain why we don't know that given tickets in fair lotteries won't win, it would still be interesting to see whether our intuitions can be explained without appeal to contextualism but in terms of the relevant subjects' failure to satisfy KLP as presently construed, for instance.

### 3.1.3 Gettier Cases

In order to see whether KLP can successfully explain our intuitions in Gettier cases, recall Zagzebski's appealing diagnosis of how Gettier cases work—*viz.* by introducing a bit of bad luck and cancelling it out by a bit of good luck. Recall that in chapter III we used Zagzebski's diagnosis to argue that the anti-luck platitude can explain the subject's ignorance in Gettier cases in the following way: We observed that the bit of bad luck is such that, in the subject's circumstances, it would normally lead her to form a false belief and that it is only through the bit of good luck that she ends up with a true belief after all. Accordingly, the Gettiered subject will end up with a belief that is only luckily true. It is not hard to see that this argument can be recovered by the defender of KLP. After all, it would seem that it is just as plausible to say that the bit of bad luck ensures that, in the subject's circumstances, the cognitive abilities she exercises would normally lead her to a false belief and that it is only through the bit of good luck that she ends up with a true belief after all. If so, the Gettiered subject will end up with a belief that, relative to her cognitive performance, will only be luckily true. It follows that Gettiered subjects fail to satisfy KLP. Given that this is so, then there is excellent reason to believe that KLP explains all the phenomena that the veritic luck principle promises to explain.

## 3.2   KLP Does Better than the Veritic Luck Principle

While I do believe that KLP explains all the phenomena that the veritic luck principle may hope to explain, I think that there is reason to believe that there are some cases in which KLP does a better explanatory job than the veritic luck principle. To be more precise, as I am about to argue, there is a type of case in which, intuitively, a subject does not know a proposition, whilst satisfying the veritic luck principle but not KLP.

Since, intuitively, the reason for his ignorance is that he is too lucky to qualify as knowing, we then have reason to believe that the type of luck at issue in the anti-luck platitude is the one that figures in KLP (rather than veritic epistemic luck). By the same token we will have reason to believe that the achievement platitude—since it entails KLP so that the cognitive achievement involved in knowing excludes the type of luck we have reason to believe is at issue in the anti-luck platitude—entails the anti-luck platitude.

The case I have in mind is a variation of a Gettier case originally due to Roderick Chisholm: I am driving through the countryside, see a sheep-shaped object in the field and on that basis come to believe that there is a sheep in the field. While there happens to be a sheep in the field, the object I am looking at is in fact a non-sheep that looks like a sheep. While my belief is both justified and true, intuitively, it does not qualify as knowledge. I am Gettiered.[19] Notice that in the original version of the case, my belief fails to satisfy both the veritic luck principle and KLP. My belief fails to satisfy KLP because my otherwise reliable cognitive ability to identify sheep fails me in the present case. In forming a belief on the basis of looking at a non-sheep I perform badly with my cognitive ability. As we have seen above, if one succeeds despite bad performance, then, relative to one's performance, it will be a matter of luck that one has succeeded. It follows that, relative to my cognitive performance, it is a matter of luck that I have hit upon the truth about there being a sheep in the field. My belief thus fails to satisfy KLP. My belief also fails to satisfy the veritic luck principle because it is merely a matter of luck that my belief turns out to be true. After all, it is a matter of luck that a sheep happens to be somewhere in the field.

But now consider the following variation of the present Gettier case: While I still form my belief that there is a sheep in the field on the basis of looking at a sheep-shaped

[19] Cf. Chisholm (1977), p. 105.

object that in fact isn't a sheep, it is not as if there merely *happened* to be a sheep in the field. Rather the environment I am driving through is strongly populated by sheep. Sheep are in all the fields of the country. I take it that our intuition that my belief, while justified and true, does not qualify as knowledge carries over to this slight variation of the original case. Moreover, it seems to me that an intuitively appealing explanation of why I lack knowledge is—just as in the other Gettier cases—that I am too lucky to count as knowing.

Now, let's ask whether my belief satisfies the veritic luck principle or KLP or both. First, it would seem as though my belief does not satisfy KLP. After all, I am still forming my belief on the basis of looking at a sheep-shaped object that in fact is not a sheep. In this way I am still successful despite performing badly. Thus, relative to my cognitive performance, it is still a matter of luck that I hit upon the truth. At the same time, however, it would seem that my belief satisfies the veritic luck principle. After all, provided that the fields of the country I am driving through are all strongly populated with sheep, it is not at all a matter of luck that there is a sheep in the field (or in any other field in the country). In consequence, it is hard to see how it could still be a matter of luck that the belief I form is true—at the very least if, in my circumstances, it is not at all a matter of luck that I form a belief that there is a sheep in the field.[20] And we may suppose that it is not at all a matter of luck that I form such a belief: In the present circumstances, I would take a sheep-shaped object to be a sheep and since there are so many sheep around I was bound to see a sheep-shaped object when looking out of the window. The underlying point can be nicely brought out by considering the structural features of the situation in terms of the relation between the different types of luck introduced in chapter IV, section 2: If it is not at all a matter of luck that *p* is true and it

---

[20] In fact, all that is needed here is that I form a belief that there is a sheep in the field and that it is not at all a matter of luck that I do not form a belief in a proposition inconsistent with or logically stronger than the proposition that there is a sheep in the field instead. If so then the belief I form will not be afflicted by veritic epistemic luck.

is not at all a matter of luck that one believes that *p*, then it is not at all a matter of luck that the belief one forms is true either.[21] Or, in other words, if it is not at all a matter of *content* luck that *p* is true and it is not at all a matter of *doxastic* luck that one believes that *p*, then one's belief is not afflicted by *veritic* luck either. In the present case the fact that the fields of the country I am driving through are so strongly populated with sheep ensures that it is not at all a matter of content luck that there is a sheep in the field. At the same time, the fact that, in the present circumstances, I would take sheep-shaped objects to be sheep and the fact that, given the strong sheep population, I was bound to see a sheep-shaped object when looking out of the window ensure that, in the present circumstances, it is at all not matter of doxastic luck that I believe that there is a sheep in the field. In consequence, the belief I form—*viz.*, that there is a sheep in the field—is not afflicted by veritic epistemic luck.[22]

In the present variation of the sheep-case, and in similar cases in which I perform badly and hit upon the truth only despite my bad performance, I intuitively lack knowledge. At the same time, I may satisfy the veritic luck principle—that is, provided that it is not at all a matter of content luck that the proposition I believe is true and that it is not at all a matter of doxastic luck that I believe the proposition. Since I perform badly, however, I will not satisfy KLP. Thus, KLP will be able to explain our intuitions

---

[21] Or, again, alternatively, if it is not at all a matter of luck that *p* is true and if one believes that *p* and it is not at all a matter of luck that one does not believe something inconsistent with *P* or logically stronger than *P* instead of *P*, then it is not at all a matter of luck that the belief one forms is true.

[22] It is noteworthy that the present case is yet another case in which a Gettiered subject believes safely. After all, the fact that the fields of the country I am driving through are so strongly populated with sheep ensures that at all nearby possible worlds it remains the case that there is a sheep in the field. Moreover, the facts that, in my present circumstances, I would take sheep-shaped objects to be sheep and that I was bound to see a sheep when looking out of the window ensure that at nearby possible worlds continue to I see a sheep-shaped object and form a belief that there is a sheep in the field. Given that both facts enter into the relevant initial conditions for the formation of my belief, the belief I form is safe. At the same time, intuitively, my belief does not qualify as knowledge. I am still Gettiered. Thus safe belief cannot be the anti-Gettier condition.

Moreover, the present considerations provide us with a recipe to generate Gettiered safe beliefs. All one needs to do is to set up a case in which one performs badly with one's cognitive abilities such that one's belief is too lucky to qualify as knowledge, while, at the same time, it is not at all a matter of content luck that *p* is true or of doxastic luck that one believes that *p* (such that *P* continues to be true and one continues to believe *P* at nearby possible worlds). In such a case one will have a safe belief which, intuitively, does not qualify as knowledge.

in such cases, while the veritic luck principle fails to be able to do so. Moreover, at least in some such cases (as in the above variation of the sheep-case) it seems correct to say that I am too lucky to count as knowing.[23] Given that this is so, and given that KLP explains our intuitions in all the other types of case in which the veritic luck principle promises to do so, there is reason to believe that the type of luck at issue in the anti-luck platitude is the type of luck at issue in KLP rather than veritic epistemic luck. By the same token, there is reason to believe that there is a relation between the achievement platitude and the anti-luck platitude—*viz.*, that the former entails the latter.

## 4  The Ability Condition

### 4.1  The Case for the Ability Condition

In the previous sections, I have argued that the achievement platitude allows us not only to fit the diagnosis of last chapter's problematic Gettier case into the minimalist framework, but also that there is reason to believe that the achievement platitude entails the anti-luck platitude.[24] In this way, I believe that adopting the achievement platitude

---

[23] Given that this is so, something must have been wrong with the argument that if Zagzebski's diagnosis of Gettier cases holds, then Gettiered beliefs will be veritically lucky. Recall that, according to Zagzebski's diagnosis, Gettier cases work by introducing a bit of bad luck and cancelling it out by a bit of good luck. We then observed that, in the circumstances, the bit of bad luck would normally lead one to form a false belief and that it is only through the bit of good luck that one ends up with a true belief after all. It seems to me that the problem for those who wish to say that the problem with Gettier cases is veritic luck is that they cannot make sense of the idea that bad luck would *normally* lead one to form a false belief. What the present argument shows is that one can construct a Gettier case in which the subject would *normally* form a false belief—I take it that it is plausible that someone who forms a belief about the presence of sheep by looking at a non-sheep would normally form a false belief—while in the particular situation, she does not and *would not* form a false belief. Since whether or not a belief is afflicted by veritic luck depends only on the particular situation, the present Gettier case creates a problem for those who wish to defend the claim that the problem with Gettier cases is veritic luck.

As opposed to that, the type of luck at issue in KLP can make sense of the idea that bad luck would *normally* lead one to form a false belief in the sense required. After all, it is plausible that exercising the cognitive abilities in the type of situation one finds oneself in would normally lead one to a false belief. Exercising one's sheep-spotting abilities when looking at a non-sheep will normally lead one to a false belief. In this way, there is reason to believe that the defender of KLP can hold onto the parallel argument from Zagzebski's diagnosis of Gettier cases to the idea that KLP will successfully explain ignorance in such cases.

[24] Recall that I also said that there is reason to believe that the ability condition entailed by the achievement platitude provides a plausible development of the ability platitude and that, if the ability

brings with it further theoretical benefits. But now notice that much of the argument rested on the claim—defended at the outset of this chapter—that the achievement platitude entails an ability condition on knowledge. Let us now take a closer look at this condition. To begin with, as I will argue in due course, those who accept the ability platitude will do well to accept the ability condition as a development of the ability platitude because the ability condition will allow them to solve one of the most pressing problems they may otherwise encounter. For those who do not accept the ability platitude the fact that the ability condition explains our intuitions in the problematic Gettier case provided in the last chapter ought to provide reason to accept it. Moreover, the fact that the ability condition is entailed by the achievement platitude and the achievement ability principle (AAP) gives some motivation for the ability condition.[25] Most importantly, however, as we have seen, the ability condition entails the knowledge luck principle (KLP). And, as we have also seen, there is reason to believe that KLP captures the type of luck at issue in the anti-luck platitude. After all, it not only predicts ignorance in all the types of case that motivated the anti-luck platitude, but it also outperforms the competing veritic luck principle in explanatory power. Since the ability condition entails KLP, however, that means that we also have reason to believe that neither subjects who have formed their beliefs in unreliable ways, nor Gettiered subjects nor subjects who entertain beliefs in lottery propositions will satisfy the ability condition. That, however, is probably as strong support for the ability condition as we may hope to find. In this way, then, the ability condition is a very strong candidate for a

---

condition does identify the sense of relevant ability in the ability platitude, the achievement platitude entails the ability platitude. Since the ability condition entails an anti-luck principle that we have reason to believe identifies the type of luck at issue in the anti-luck platitude, if the ability condition does identify the sense of relevant ability in the ability platitude, then we have reason to believe that the three platitudes are related in the following way: The achievement platitude entails the ability platitude which, in turn, entails the anti-luck platitude.

[25] Notice that while I may be bootstrapping here, I think that my argument is not circular. After all, the achievement platitude has some independent (intuitive) plausibility which can be used to motivate the ability condition, while, at the same time, the theoretical benefits that the ability condition provides give *further* support for the achievement platitude.

necessary condition for knowledge. But now recall that the minimalist approach dictates that in order to qualify as a condition for knowledge a candidate must also be consistent with all the platitudes of the minimalist framework. In order to complete the case for the ability condition I will now adduce some considerations suggesting that the ability condition is consistent with the platitudes defended here.

## 4.2 The Ability Condition is Consistent with the Platitudes

### 4.2.1 The Factivity and Belief Platitudes

It can easily be seen that the ability condition is consistent with the factivity and belief platitudes. According to the ability condition, one knows that *p* only if the *truth* of one's *belief* that *p* has been ensured through the exercise of a set of relevant cognitive abilities. In order to satisfy the ability condition, then, one must have a true belief. The ability condition entails and is therefore consistent with the belief and factivity platitudes.

### 4.2.2 The Good Informant Platitude

According to the good informant platitude, knowers are typically—that is, except when they are unwilling to part their information, have lost their credibility etc.—good informants. Accordingly, the good informant platitude will be inconsistent with the ability condition just in case either of the following two entailments holds: First, if one satisfies the ability condition, then one is not a good informant (except when one unwilling to part one's information etc.). Or, second, if one is a good informant, then one does not satisfy the ability condition. It can easily be seen that neither entailment holds. Suppose Henry drives through Normal Barn Country, sees a barn and comes to

believe that there is a barn in the field. Henry's belief is true. Moreover, the truth has been secured through the exercise of a set of relevant cognitive abilities. Thus, Henry satisfies the ability condition. At the same time—assuming, as we may, that Henry is willing to part his information, is credible etc.—he is certainly also a good informant about whether there is a barn in the field. Again, if I were lying in the back, trying to sleep but were overcome by the desire to know whether there is a barn in the field, I would not hesitate to ask Henry. Thus the first entailment does not hold. As regards the second entailment, suppose a contemporary of Gödel's cannot remember whether Gödel proved the completeness or the incompleteness of arithmetic. Since we may assume that Gödel is willing to part the information, is credible etc., certainly, Gödel himself will be a good informant for our man. At the same time, Gödel not only knows that arithmetic is incomplete, he has also hit upon this fact through the exercise of a set of relevant mathematical abilities. Thus, Gödel is not only a good informant but also satisfies the ability condition. The second entailment fails as well. The ability condition and the good informant platitude are thus consistent.

### 4.2.3 The Informative Speech Act and Assertion Platitudes

In order to provide reason to believe that the ability condition is consistent with the assertion and informative speech act platitudes, I will—just as in the parallel argument for the consistency of the safety condition with these platitudes[26]—argue that some situations that would show the two to be inconsistent do in fact not obtain. The relevant situations are situations in which one felicitously asserts/performs an informative speech act that $p$, whilst, at the same time, one has not ensured the truth of one's belief that $p$ through the exercise of a set of relevant cognitive abilities. In this regard, let's

---

[26] Cf. chapter IV, section 4.1.

now reflect on the sorts of situations in which one does not satisfy the ability condition. It would seem that the following three types of situation capture the widest range of such situations: first, one believes unreliably; second, one believes a lottery proposition; and, third, one's belief is Gettiered. As we have seen in the parallel discussion of the safety condition, however, it is plausible that in neither type of situation one's assertion/informative speech act will be felicitous—at least not unless exceptional circumstances obtain.[27] The most imminent threat to the suggestion that the assertion and informative speech act platitudes on the one hand and the ability condition on the other are consistent can thus be defused.

### 4.2.4 The Closure Platitude

Recall that the closure platitude has it that one knows the conclusions of one's competently executed deductions provided that one knows the premises and comes to believe the conclusions on the basis of one's deductions. To begin with, it is plausible that to competently deduce one proposition from a set of others is to exercise a cognitive ability. At the same time, if one also knows the premises of one's deduction and one's deduction is competent, then one is bound to hit upon the truth with respect to the conclusion. After all, knowledge is factive and one will have competently deduced one proposition from a set of others only if the members of latter jointly entail the former. That is to say, however, that if one satisfies the closure platitude, the ability one exercises will ensure that one hits upon the truth. Moreover, since the closure platitude requires that one form a belief in the conclusion on the basis of the competent deduction, if one satisfies the closure platitude, one will also have formed a belief the truth of which is ensured through the exercise of a cognitive ability. Given that this is

---

[27] Cf. chapter IV, section 4.2.

so, the only way in which one can satisfy the closure platitude and fail to satisfy the ability condition is by having exercised an irrelevant cognitive ability. However, it is massively implausible that this could be so. After all, the closure platitude specifies a specific cognitive ability that must be exercised in acquiring closure-based knowledge. A plausible way of understanding what is going on here is that the closure platitude identifies the cognitive ability *relevant* for closure-based knowledge. If so, one cannot satisfy the closure platitude and have formed a belief the truth of which is ensured through the exercise of irrelevant cognitive abilities. In consequence, anyone who satisfies the closure platitude will also satisfy the ability condition. Thus the two are consistent.

### 4.2.5   The Achievement, Ability and Anti-Luck Platitudes

To begin with, since the achievement platitude entails the ability condition the latter is obviously consistent with the former. Since the ability condition can be viewed as a sharpened version of the ability platitude the two are also consistent. Finally, in order to see that the ability condition is consistent with the anti-luck platitude recall that the ability condition entails KLP and that we have reason to believe that the type of luck at issue in the anti-luck platitude is the type of luck at issue in KLP. Given that this is so, there is also reason to believe that the ability condition is consistent with the anti-luck platitude.

In conclusion, it has been argued that the ability condition can be derived from the achievement platitude and the achievement ability principle. We have also seen that there is reason to believe that the ability condition entails KLP and that it will therefore explain our intuitions in the types of case that motivated the anti-luck platitude—*viz.* cases of unreliably formed belief, lottery cases and Gettier cases. Finally, I have

provided reason to believe that the ability condition is consistent with the platitudes of the minimalist framework. In this way, a case has built for a conception of knowledge according to which knowledge requires that the truth of one's belief be ensured through the exercise of a set of relevant cognitive abilities. Before closing, I will consider a couple of objections to the arguments and theses defended in this chapter.

## 5 Objections

In this section I will discuss one objection to the ability condition on knowledge—or, to be more precise, to what I have argued the ability condition can do for us—and another one to the achievement platitude. The first objection, which is due to Pritchard, is an argument that the ability condition does not successfully explain ignorance in all Gettier cases, while the second objection, which can be developed form an argument due to Jennifer Lackey, calls into doubt the claim that knowledge requires achievement.

### 5.1 Pritchard's Objection: The Ability Condition Does Not Explain Ignorance in All Gettier Cases

In order to convince us that the ability condition will not explain ignorance in all Gettier cases, Pritchard invites us to consider the following case: Archie is a professional archer. He goes to the shooting range, picks a target, and takes a shot. In taking a shot Archie exercises his arching abilities. Now suppose that Archie hits the target right in the middle. Intuitively, Archie has succeeded through the exercise of his arching abilities. The intuition that Archie succeeds through the exercise of his arching abilities lingers even if we suppose, in addition, that, unbeknownst to Archie, he is shooting at

the only target at the shooting range that isn't fitted with a hidden forcefield that would repel any shot fired at it.[28]

Now, Pritchard points out that the case of Archie is formally analogous to a Gettier case we have already encountered on various occasions—*viz.* the one of Henry in Barn Façade County. Recall that, in that case, Henry drives through the countryside, looks at the only barn in a field otherwise full of barn façades so cleverly constructed that they cannot be distinguished from real barns by people driving along the road and forms a belief that he is facing a barn. Henry's case is analogous to Archie's in that the barn façades play the role of the sabotaged targets and the fact that Henry is looking at the only real barn parallels the fact that Archie is shooting at the only non-sabotaged target. Accordingly, if we acknowledge, as is plausible, that Archie succeeds in hitting the target through the arching abilities he exercises, then we must also agree that Henry exercises his barn-spotting abilities and that he hits upon the truth through the exercise of these abilities. It would seem, however, that that is to agree that Henry satisfies the ability condition. At the same time, intuitively, Henry's justified and true belief that he is facing a barn does not qualify as knowledge. Henry is Gettiered. Given that this is so, it would also seem that the ability condition fails to explain Henry's ignorance and, in consequence, that there are some Gettier cases in which the ability condition will be unable to explain the Gettiered subject's ignorance.[29]

It is in dealing with this particular objection that the ability condition as developed from the achievement platitude and AAP does better than the ability platitude as stated at the outset of the chapter. After all, notice what Pritchard's argument makes plausible is that Henry exercises his barn-spotting abilities. Moreover, since, at first glance at least, in Henry's circumstances, his barn-spotting abilities are relevant to his coming to know that he is facing a barn, it may be hard to see how Henry could fail to

---

[28] Cf. Pritchard (forthcoming*b*), p. 11.
[29] Cf. Ibid.

satisfy the ability platitude. If, however, the ability platitude is developed along the lines of the ability condition as derived from the achievement platitude and AAP, arguably, appearances are misleading. For, arguably, Henry fails to satisfy the ability condition. To see why this is so, notice, first, that in developing the ability condition we have given a somewhat more detailed account of the notion of a set of relevant cognitive abilities at issue in the ability condition. Recall that we pointed out, first, that the notion of a set of relevant cognitive abilities is to be understood in the same way as the notion of a set of relevant abilities at issue in other achievements; second, that having a set of abilities relevant for securing a given achievement requires that, in one's circumstances, one's abilities reliably lead one to success; and, third, that whether or not one possesses a set of abilities relevant for securing a given achievement may depend heavily on one's (external) circumstances. That said let us now ask whether, in his circumstances, Henry has a set of cognitive abilities relevant for the cognitive achievement involved in knowing that he is facing a barn. It is plausible that the answer to this question is that Henry does not have such a set of abilities. After all, the abilities he will exercise in forming a belief that he is facing a barn will involve certain discriminatory abilities. At the same time, it is also plausible that in Henry's circumstances—in which barn façades prevail—any set of abilities relevant for knowing that he is facing a barn will involve an ability to discriminate barns from barn façades.[30] However, *ex hypothesi*, Henry does not possess such an ability. He only possesses abilities that allow him to discriminate barns from other mundane objects such as houses, cows, hills etc. While it is plausible that Henry exercises these latter abilities[31], and while, if we accept the analogy with

---

[30] Or, to be more precise, any set of abilities relevant for knowing *by looking* that one is facing a barn will involve such a discriminatory ability. Since Henry is forming a belief *by looking* the assumption that, in Henry's circumstances, any set of abilities relevant for knowing that he is facing a barn will involve such a discriminatory ability ought to be uncontroversial.

[31] We can now also see that we had better not place a reliability condition even on *cognitive* ability possession *simpliciter* of the kind envisaged by Sosa and Greco (cf. p. 164, n.8). After all, if we did place such a condition on cognitive ability possession *simpliciter*, then, since the cognitive abilities Henry uses to form a belief about the presence of barns in Normal Barn County do not reliably lead him to the truth

Archie the archer, we will have to grant that Henry hits upon the truth through the exercise of these latter abilities, it is not the case that Henry satisfies the ability condition. In order to satisfy the ability condition Henry needs to hit upon the truth through the exercise of a set of abilities that, in his circumstances, is relevant for the cognitive achievement involved in knowing. In order to do so, he needs an ability he does not have—*viz.* one enabling him to discriminate between barns and barn façades. Thus, even if we have to acknowledge that Henry hits upon the truth through the exercise of a set of abilities, he does not satisfy the ability condition.

But what about Archie the archer? It would seem that Archie not only succeeds through the exercise of ability but that he also secures the relevant athletic achievement when he hits the target. However, that means that he must have succeeded through the exercise of a set of relevant abilities. But how is that possible provided that the case of Archie is formally analogous to Henry's case and that Henry is in a circumstance in which he lacks one ability any set of abilities relevant for the cognitive achievement he is after will involve? The answer, it seems to me, is that even in otherwise parallel situations, different types of achievements may render different types of abilities relevant for securing the respective achievements. Thus, it would seem that, while, in Archie's case, any set of abilities relevant to the athletic achievement involved in hitting the target will involve all sorts of physical and cognitive abilities, it is not the case that any such set of abilities involves an ability that enables him to discriminate between sabotaged and non-sabotaged targets. And why should it? After all, the achievement involved in hitting the target is athletic. As opposed to the achievement Henry aims for it does not consist in the successful discrimination of a certain thing.

---

when in Barn Façade County, Henry loses and hence cannot exercise them when in Barn Façade County. Since this is implausible, we had better not place a reliability condition on cognitive ability possession *simpliciter* of the kind envisaged by Sosa and Greco.

Another question one may ask is how Archie can be credited with an achievement given that he is in circumstances in which he is surrounded by sabotaged targets. After all, it was argued earlier on that in order to be in a position to secure a given achievement one must be in a set of circumstances in which one is reliably successful. However, since Archie is surrounded by sabotaged targets, it seems that Archie is not reliably successful. If so, it is not clear that, on my premises, Archie can be credited with an achievement. In response to this objection we may venture to argue that the set of circumstances affords different individuation in Henry's and in Archie's case. The fact that in Henry's case the achievement at issue will involve the exercise of discriminatory abilities may be used to argue that the set of circumstances needs to be individuated rather widely. After all, discriminatory abilities quite obviously depend on what goes on in the wider environment. As opposed to that, one may think that since none of the abilities that Archie exercises depend in the same way on what goes on in his wider environment, circumstances can be individuated more narrowly.[32]

---

[32] Notice that Greco (2007*b*) provides a rather different response to Pritchard's objection. Greco claims that since ability possession is relative to circumstances and since there is a reliability condition on ability possession, Henry is in circumstances in which he does not possess his barn-spotting abilities. If so, he cannot hit upon the truth through the exercise of his barn-spotting abilities either. Hence, he does not satisfy the ability condition. However, this response does not seem to be very promising. After all, as I have argued in section 1 the conception of ability possession on which the argument relies is not very plausible. And as we have already seen the implausibility also manifests itself in the present case (p. 186, n. 30): It would seem that in forming his belief, Henry is exercising the abilities that allow him to acquire knowledge in circumstances in which there aren't any barn façades around. That is to say, however, that he exercises *his* barn-spotting abilities. If so, however, he must also possess them in his present circumstances. The problem, it seems, is that his barn-spotting abilities are not good enough to give him knowledge in the present circumstances.

In fact, Greco's response is much less clear than presented here. Greco maintains that when in Barn Façade County, Henry does not have the "knowledge-relevant" abilities because abilities are relevant to circumstances. Now depending on how the idea of knowledge-relevant ability is spelled out, Greco's response is more or less similar to mine. If the idea of knowledge-relevant ability is unpacked as the idea of an ability strong enough to come to know, then, of course, Henry does not possess the knowledge-relevant ability. This response would be similar to mine in that Greco can allow that Henry exercises the same abilities as in Normal Barn County—that is, if he allows that whatever abilities Henry exercises constitute a knowledge-relevant ability in Normal Barn County but not in Barn Façade County. (Notice, however, that I do not construe abilities as abilities to secure an achievement but rather as abilities to secure success. Whether or not the abilities one possesses are good enough for achievement depends on the circumstances.) If, on the other hand, the idea of knowledge-relevant ability is unpacked as the idea of an ability to hit upon the truth, then Greco will be forced to accept that when in Barn Façade County, Henry does not possess—and hence cannot exercise—the abilities he has in Normal Barn County. Since elsewhere (2007*a*) Greco appears to construe abilities in the latter way, this is how I have represented him above.

Given that all this is so, the defender of the ability condition as developed from the achievement platitude and AAP will not only successfully defuse Pritchard's objection. He will also have good reasons to convince those who were inclined to accept both the ability platitude and the idea that the ability platitude explains our intuitions in Gettier cases to accept his condition as a plausible development of the ability platitude. After all, the ability condition provides an appealing solution to an otherwise difficult problem for anyone who wanted to claim that the ability platitude explains our intuitions in Gettier cases.

## 5.2   A Lackey-Style Objection: Not all Knowledge is Achievement

While Pritchard worries that the ability condition is too weak to explain our intuitions in all Gettier cases, there is another argument, due to Lackey, that can be developed so as to suggest that the achievement platitude is too strong to be plausible. In its original version Lackey's argument aims at showing that the view that knowledge is true belief for which the agent deserves credit is false. To be more precise, Lackey wants to convince us that one can know a proposition without deserving credit for one's true belief. However, it is not hard to see that her argument will also pose a problem for the achievement platitude. After all, it is plausible that people deserve credit for their achievements. Thus, if one can know without deserving credit, then it would seem that one can know without having secured an achievement. At the heart of Lackey's argument is the following case:

> Having just arrived at the train station in Chicago, Morris wishes to obtain directions to the Sears Tower. He looks around, approaches the first adult passerby that he sees, and asks how to get to his desired destination. The passerby, who happens to be a Chicago resident who knows the city extraordinarily well, provides Morris with impeccable

directions to the Sears Tower by telling him that it is located two blocks east of the train station. Morris unhesitatingly forms the corresponding true belief.[33]

Lackey maintains that, while Morris's true belief certainly qualifies as knowledge, he does not deserve credit for it. Rather, the credit goes to the Chicago resident who knows the city as well as he does. Alternatively, that Morris forms a true belief is not his achievement. Rather, the achievement is on the part of the Chicago resident. If Lackey's claims are true, then it would seem that the achievement platitude is in trouble.[34]

Somewhat surprisingly, however, it is Lackey herself who provides the key to responding to this objection. Lackey considers the following riposte on behalf of the defender of the view that knowers deserve credit for their true beliefs: In selecting the source for testimony Morris secures at least some minimal discriminatory success. After all, Morris does not ask a tree, a child, or an obviously confused person for directions. The minimal discriminatory success secures that he deserves credit for his testimonial knowledge in at least a minimal sense. Moreover, it would seem just as fair to say that in securing the minimal discriminatory success Morris also secures at least a minimal cognitive achievement. Now Lackey claims that this riposte won't do the job for the defender of the views under consideration because "…Morris could have just as easily

---

[33] Lackey (*forthcoming*), p. 10.

[34] Lackey also provides a couple of other cases that she claims will do the same job. These cases are cases in which, supposedly, the subject knows while, at the same time, no one deserves credit for her true belief. Hence, no achievement at all is secured. In the first case a virus leads a subject to form a reliable belief for which she has some independent background evidence. Lackey claims that in such a case, intuitively, the subject's belief qualifies as knowledge. At the same time, since the virus leads the subject to the truth and the virus is not in the ballpark of deserving credit/securing achievements, no one deserves credit/has secured an achievement. (cf. ibid., pp. 16-7). I must confess that I simply do not share Lackey's intuitions here. If Lackey were correct in claiming that the subject knows, then it would seem that she will have to admit that in a parallel case in which a benevolent demon plays the role of the virus she also knows. It seems to me, however, that certainly, in such a case the subject does not know. The second case Lackey considers is the case of what she calls 'innate knowledge'. She claims that since 'innate knowledge' is certainly possible we will have to accommodate it in our theory of knowledge. However, subjects do not deserve credit for 'innate knowledge'/instances of 'innate knowledge' are not achievements (cf. ibid., p. 18). Again, however, Lackey's claim that 'innate knowledge' is possible is far from uncontroversial. If internalism about knowledge turned out to be true, for instance, it is not at all clear that 'innate knowledge' will be so much as possible. Moreover, it is unclear why we would have to say that what Lackey calls 'innate knowledge' is really best construed as innate *knowledge* rather than innate *true belief*. If the defender of the achievement platitude can argue that what Lackey calls 'innate knowledge' may be construed as innate true belief, however, he will, once more, be off the hook.

approached a competent-looking compulsive liar or a directionally challenged speaker as he did an honest, knowledgeable, Chicago resident…"[35] The idea here is, I take it, that in order to secure the relevant minimal discriminatory success that would earn Morris credit, Morris would have to be able to discriminate competent adult humans not only from children, the obviously confused and so on but also from the likes of competent-looking compulsive liars and the directionally challenged. However, Morris cannot make these discriminations. In consequence, he does not secure the relevant minimal discriminatory success that would earn him credit for the true belief he acquires. By the same token he cannot be said to have secured an achievement either.

To see what is wrong with Lackey's argument, notice, first, that Lackey's central assumption that Morris *could just as easily* have encountered a competent-looking compulsive liar or a directionally challenged person is a modal claim—it is a claim about what could easily have happened. According to a standard possible-worlds semantics of the relevant notion of easy possibility, Lackey's assumption is tantamount to the assumption that there is some nearby possible world at which Morris encounters a competent-looking compulsive liar etc. Whether this is so, however, will depend on facts about the circumstances Morris is in—for instance, on whether there are there compulsive liars in the vicinity. If there aren't any compulsive liars etc. in the vicinity, then it seems, first, that Lackey's assumption is not true and second, that Morris's abilities to discriminate competent-looking adult humans from trees, children, the obviously confused and so on may be enough to allow him to secure the cognitive achievement involved in knowing. To see why this is so, recall that it depends on circumstances which abilities are needed to secure an achievement. If circumstances are hospitable, then some achievements may be easy to obtain. Thus, in circumstances in which no compulsive liars are around, Morris's discriminatory abilities may be good

---

[35] Ibid., p. 13.

enough to give him an easily obtainable minimal cognitive achievement. It is also instructive to compare the case of Henry: In Normal Barn County the achievement involved in knowing that he is facing a barn is relatively easy to obtain. Accordingly, the discriminatory abilities he possesses—that is, abilities to discriminate barns from other mundane objects such as houses, cows, hills etc.—allow him to secure this achievement.

On the other hand, if Lackey's assumption is true, then there must be some facts about Morris's circumstances that make it true—for instance, that compulsive liars are in the vicinity. However, it is independently plausible that introduction of such facts raises the bar for the achievement involved in obtaining testimonial knowledge in such a way that Morris's discriminatory abilities will no longer be good enough to allow him to secure this achievement. To see why this is so, compare once more the case of Henry: When in Barn Façade County the bar for the achievement involved in knowing that he is facing a barn is raised in such a way that Henry's abilities will no longer be good enough to secure this achievement.

Another way of describing Lackey's error is the following: In adding the assumption that Morris could just as easily have encountered a compulsive liar etc., Lackey surreptitiously gives the case the structure of a Gettier case. Certainly, however, the fact that the achievement platitude predicts that Morris lacks knowledge in such a case can hardly be said to count against it. On the contrary, it would seem that, if anything, it provides further support for the achievement platitude. Without the additional assumption, however, the case does not pose a problem for the achievement platitude. The riposte Lackey suggests on behalf of the defender of the achievement

platitude can be made to work in such a way that Morris secures the minimal cognitive achievement needed to rescue the achievement platitude.[36]

It is noteworthy, however, that Pritchard has recently argued that Lackey's argument can be revived so as to make the present response unavailable to the defender of the achievement platitude. To begin with, Pritchard points out that achievements require not only that the achiever deserve credit for his success but also that the success be *creditable* to the achiever, that the achiever get *primary credit* for his success. Applied to the achievement platitude, that is to say that knowers must not only deserve credit for their true beliefs but also that their hitting upon the truth must be creditable to them, that they get *primary credit* for hitting upon the truth. However, even if we grant that, in Lackey's case, Morris makes some successful discriminations and may therefore deserve some credit, he does not get *primary credit* for hitting upon the truth. For, certainly, the primary credit goes to the Chicago resident. Given that this is so, however, Morris does not secure a cognitive achievement when he forms a true belief concerning

---

[36] Again, Greco's response to Lackey's objection is slightly different from mine. Greco maintains that a receiver of testimony may deserve credit in the same way as a football player who receives an excellent pass and the goes on to score an easy goal may deserve credit for scoring the goal (cf. Greco (2007*a*), pp. 13-4)). It seems to me that as stated the analogy does not work. After all, scoring an easy goal from a great pass seems to be analogous to making an easy deduction from some received testimony rather than to receiving the testimony. In order to get the analogy to work, Greco would have to argue that the receiver of a pass deserves credit for the (ball-handling etc.) abilities exercised in receiving a pass. The question will then be whether a receiver of a pass must always exercise some abilities in receiving a pass. It is not clear that he must: Thus consider the following case: A football player is standing on his own on the field. His eyes are closed (perhaps he is in pain). Someone passes the ball to him. The pass is so perfect that it stops right at the receiver's foot. The receiver opens his eyes, sees the ball right in front of him and plays on. It looks as though in the reception of the pass, no ability has been exercised. Accordingly, no credit appears to be deserved for the reception of the pass. If so, Greco's analogy won't do the job.

Now one might think that I have shot myself in the foot here. After all, given that there is an analogy between the receiver of testimony and the receiver of a pass it now looks as though it should also be possible to receive testimony without deserving credit (and hence without securing an achievement). However, I am not convinced that even the analogy between the receiver of a pass and the receiver of testimony holds. For, I am not sure what the testimonial analogue of the case in which the receiver of the pass need exercise no abilities whatsoever would look like. Indeed, I have doubts whether there could even be such an analogue. After all, it seems that one will always have to exercise some abilities in receiving testimony. Perhaps the closest we can get is a case in which one overhears a fraction of a conversation between people one does not know. Even in this case, it seems that one exercises some discriminatory abilities. One discriminates the voice of an adult from the voice of the child the bark of a dog etc. If the circumstances are hospitable enough that may be enough to earn one some credit and even to secure at least a minimal achievement. For these reason, I think my response to Lackey's case is preferable to Greco's.

the whereabouts of the Sears Tower. Thus, the pressure is still on the defender of the achievement platitude.[37]

However, Pritchard's improved version of Lackey's argument strikes me as unsound. To see why consider the following case: I drive from Aberdeen to London. As I drive past Edinburgh I see a sign which reads "London 332 miles". Upon reading the sign I form a true belief that the distance between Edinburgh and London is 332 miles. Intuitively, my belief qualifies as knowledge. More importantly, however, it seems overwhelmingly plausible that, in hitting upon the truth, I secure a cognitive achievement. But now notice that if this is correct, then something must be wrong with Pritchard's improved argument. After all, if we agree that the primary credit for Morris's hitting upon the truth goes to the Chicago resident, then we must certainly also agree that the primary credit for my hitting upon the truth goes to the surveyor who measured the distance between London and Edinburgh. If so, however, either Pritchard's claim that achievements require that the success be creditable to the achiever or his claim that Morris does not get primary credit for hitting upon the truth must be false. Thus, Pritchard's improved argument of Lackey's case does not go through. The defender of the achievement platitude is once more off the hook.

In summary, then, this chapter has, first, provided reason to adopt the achievement platitude into the minimalist framework. It has been argued that the achievement platitude is not only independently plausible but also that it entails the anti-luck platitude[38] and therefore promises to explain a variety of phenomena—including the problematic Gettier case discussed in the last chapter. Moreover, one of the most pressing objections to the achievement platitude, which focuses on testimonial knowledge, has been warded off. It has also been argued that the ability condition can

---

[37] Cf. Pritchard (2007), pp. 14-15.
[38] If the ability condition captures the sense of the notion of relevant ability in the ability platitude, then there is also reason to believe that the achievement platitude entails the ability platitude (cf. also p. 167, n. 14).

be derived from the achievement platitude and the achievement ability principle (AAP). Since we found the both the achievement platitude and AAP to enjoy some independent plausibility, the argument provided some motivation for accepting the ability condition. More importantly, however, it was argued that the ability condition entails an anti-luck principle—*viz.* KLP—and that the type of luck at issue in KLP is the type of luck at issue in the anti-luck platitude. Given that this is so, there is reason to believe that the ability condition explains the phenomena that motivated the anti-luck platitude which, in turn, provides very strong support for the ability condition. Moreover, it has been argued, that even those who are willing to adopt the ability platitude into the minimalist framework will do well to develop the platitude along the lines of the ability condition as it will allow them to avoid what may otherwise become a difficult problem for them. Since we found that there is excellent reason to believe that the ability condition is consistent with the other platitudes of the minimalist framework, this chapter has also made a case for the ability condition.

# Conclusion

With Gettier's refutation of the received view the quest for the correct analysis of knowledge was kicked off. New and ever more complicated sets of necessary and sufficient conditions for knowledge were offered and refuted by new and ever more ingenious Gettier-style counterexamples. At the same time, the debate over the justification condition at issue in the received view was sparked and developed into one of the fiercest debates in the history of epistemology. Both developments gave us reason to be pessimistic about the prospects for identifying a set of individually necessary and jointly sufficient conditions for knowledge. In reaction to these developments and the pessimism it gave rise to, this thesis has explored two alternative approaches to the traditional analytical project—*viz.* deflationary and minimalist approaches to epistemology. Deflationary conceptions of knowledge were construed as maintaining that there is nothing epistemic (such as justification) that all subjects to which the predicate 'knows that *p*' applies have in common that explains why it applies to them. While there are two approaches in the epistemological literature that can justifiably be described as deflationary in this sense, none of them turned out to be satisfactory. In particular, both approaches failed to provide satisfactory analyses of some Gettier and lottery cases. Accordingly, the move was made to minimalist approaches to epistemology, approaches, that is, that aim to characterise the concept of knowledge by a series of platitude about knowledge. I have defended one version of such a minimalist approach to epistemology. I started out with a set of platitudes relating knowledge to truth, belief, good informants, informative speech acts, assertion, competent deduction and luck. In the last chapter, however, it was argued that there is good reason to add a platitude relating knowledge to cognitive achievements and perhaps even a platitude

relating knowledge to abilities. The minimalist framework I would like to offer, then, contains the following platitudes:

> The Factivity Platitude: One cannot know falsehoods.
> The Belief Platitude: One can know only what one also believes.
> The Good Informant Platitude: Someone who knows $P$ is typically—that is, provided that he is willing to part the information, has not lost his credibility etc.—a good informant concerning $P$.
> The Informative Speech Act Platitude: If one performs any speech act of the class of informatives, then one thereby represents oneself as knowing its content.
> The Assertion Platitude: If one asserts that $p$, then one thereby represents oneself as knowing that $p$.
> The Closure Platitude: One knows the conclusions of one's competently executed deductions provided that one knows the premises and comes to believe the conclusions on the basis of one's deductions.
> The Anti-Luck Platitude: Knowledge excludes epistemic luck.
> The Achievement Platitude: Knowledge is a cognitive achievement.
> (The Ability Platitude: If one knows that $p$, then the truth of one's belief that $p$ must have been ensured through the exercise of one's relevant cognitive abilities.)

Apart from just defending this minimalist framework, this thesis has also explored ways in which it can be developed into a more substantive conception of knowledge. In this vein I presented a number of arguments suggesting that knowledge entails safety from error one the one hand and, on the other hand, that the following ability condition on knowledge holds:

*The Ability Condition*    One knows that $p$ only if the truth of one's belief that $p$ has been ensured through the exercise of a set of relevant cognitive abilities.[1]

Moreover, I also scrutinized the relation between the ability condition and luck. In this regard, the most important argument was that the ability condition entails the knowledge luck principle KLP:

*Knowledge Luck Principle (KLP)*    One knows that $p$ only if, relative to one's cognitive performance, it is not a matter of luck that one has formed a true belief.

---

[1] I take it that the ability condition is a substantive condition on knowledge even if we are willing to accept the ability platitude because it gives a substantial account of the notion of relevant cognitive abilities at issue in the platitude.

It was argued that the type of luck at issue in KLP is the type of luck at issue in the anti-luck platitude. Given that this is so, a number of positive consequences ensue: First, since the achievement platitude entails the ability condition and hence KLP we have now reason to believe that the achievement platitude entails the anti-luck platitude. Second, since the ability condition entails KLP, we have now reason to believe that it explains the phenomena that motivated the anti-luck platitude. Third, Pritchard's argument from the veritic luck principle and a modal conception of luck to the safety condition can be recovered. It is noteworthy, however, that the argument takes quite a different shape. It proceeds from the ability condition to KLP and then moves from KLP and the weakened modal conception of luck, LE*, according to which absence of luck entails modal stability but not vice versa, to the corresponding safety condition.

Another welcome consequence of the minimalist framework is that it holds out the hope of explaining why knowledge valuable to us. After all, it is plausible that achievements are valuable. Moreover, as Pritchard has argued, achievements are valuable *for their own sake*.[2] Given that this is so, then, since by the lights of the minimalist framework here defended knowledge is a cognitive achievement, knowledge is valuable for its own sake. Accordingly, the minimalist framework holds out the hope of explaining the value of knowledge.[3]

Having pointed to some of the consequences of the minimalist framework defended here, let me now, finally, return to an issue introduced in the introduction of this thesis—*viz.*, whether the minimalist framework I have developed allows us to make progress with respect to the problems for the analytical project that resulted from Gettier's refutation of the received view. Does the minimalist framework defended here allow us to address either or both of the issues concerning the anti-Gettier condition and

---

[2] Cf. Pritchard (2007*a*), p. 10.
[3] The thesis that the kind of conception of knowledge presented here solves the value problem has been defended in more detail in Greco (2007*a*) and (forthcoming) and in Riggs (2002) and (forthcoming).

the status of the internalist justification condition? As regards the anti-Gettier condition, the arguments presented in the previous chapters provide reason to believe that the minimalist framework allows us to make progress at least in the sense that it tackles the problem at its roots. That is, in contrast with the greatest part of the proposed analyses in the post-Gettier literature, the minimalist framework does not merely address the latest Gettier-style counterexample and then goes on to hope for the best. Instead, the minimalist approach allows us to start from a plausible—albeit for the purposes of traditional analysis in and of itself not directly helpful—diagnosis of what the crux is in Gettier cases—*viz.* luck.[4] It is then argued that the type of luck at issue in Gettier cases is the type of luck at issue in KLP. Since the condition developed from the minimalist framework, that is, the ability condition, entails KLP, there is then independent reason to believe that the ability condition will explain ignorance in Gettier cases. In this way, there is reason to believe that the minimalist framework does allow us to make progress towards solving the Gettier problem.

But what about the debate over the status of the justification condition at issue in the received view? It would seem that none of the arguments presented in this thesis bear in any obvious way on the question as to whether we had better retain this justification condition. However, even if the arguments in and of themselves do not bear on this issue, there is something in the minimalist framework that might allow us to address it. Notice that in introducing the achievement (and/or ability) platitude, in accepting the ability condition, and in viewing them as doing the explanatory job that it was heretofore thought conditions such as the justification or the safety condition would do, the primary focus of epistemic evaluation was shifted from beliefs to subjects. That is to say, in order to assess whether a given belief is in the ballpark for knowledge, no

---

[4] This diagnosis is in and of itself not directly helpful for the purposes of traditional analysis in that it cannot be exploited for the purposes of traditional analysis in a direct way: Analysing knowledge as non-lucky true belief does not make for a satisfactory analysis of knowledge. I will come back to this issue in due course.

longer do we have to answer the question as to whether it, *the belief*, met certain further to be specified conditions such as a justification or safety condition. Rather, the pertinent question now was whether the *subject* secured an achievement and whether the *subject* has hit upon the truth through the exercise of a set of relevant cognitive abilities.[5] There is reason to believe that this shift in the primary focus of epistemic evaluation can be brought to bear on the issue of the status of the justification condition as it figures in the received view. To begin with, if the focus of epistemic evaluation is shifted from beliefs to subjects in this way, it is no longer clear that an additional condition on knowledge that continues to evaluate beliefs rather than subjects is still well motivated. On the contrary, a theory of knowledge that endorses the achievement (and/or the ability) platitude, the ability condition and the internalist justification condition at issue in the received view would seem to be an odd hybrid that we may expect to have little appeal. In this way, a defender of the minimalist framework has reason to reject that internalist justification as it appears in the received view is requisite for knowledge.

---

[5] Notice that in so shifting the primary focus of epistemic evaluation, the minimalist framework defended here is rendered akin to so-called virtue theories of knowledge. (Virtue theories of knowledge have been defended in Sosa e.g. (1991*b*), Greco e.g. (2000), Code e.g. (1987), Montmarquet e.g. (1993), and Kvanvig e.g. (1992), and Zagzebski e.g. (1996)). Virtue epistemologists typically also accept a further thesis about knowledge—*viz.*, that a belief qualifies as knowledge only if it has been produced by an intellectual virtue—where intellectual virtues are construed as stable and reliable traits that make up the cognitive character of its bearer. It is noteworthy that the ability condition defended here is independent of this further thesis. After all, it is plausible that one can exercise an ability out of character, as it were. This seems to me to be true, for instance, of the sceptical detective Green we encountered in chapter IV, section 2.1.2. Green knows that Hynes is the murderer because he evaluated the relevant forensic evidence and formed a true belief on its basis on this particular occasion although he normally requires more than the forensic evidence in order to form a belief—recall that, in addition, he needs evidence that the suspect has a motive for the deed. In this case, it seems to me, Detective Green has formed a belief as a result of the exercise of a set of cognitive abilities relevant for securing the cognitive achievement involved in knowing. However, due to some trait that makes up his cognitive character—*viz.*, that he sets the bar for belief-formation very high—these abilities do not normally lead to the formation of a belief. The unique circumstances disable his character traits and allow him to form a belief as the result of the exercise of his abilities anyway. Thus, in the case of Detective Green, the predictions of the ability condition and typical virtue theories of knowledge may differ. The defender of the ability condition may allow that Green knows that Hynes is the murderer because he formed a belief on the basis of the exercise of a set of relevant cognitive abilities. As opposed to that, we must expect that typical virtue theories will predict that Green's belief does not qualify as knowledge since Green forms a belief in spite of his otherwise more sceptical character. Since intuition sides with the ability condition's predictions here, the case also indicates that the ability condition has an edge over typical virtue theories of knowledge. (For a similar argument see Lackey (2007).)

Does that mean that the defender of the minimalist framework may now hope to have closed the book on the internalism/externalism debate? It seems to me that the answer to this question is 'no'. There is reason to believe that the debate between internalists and externalists (or something very similar) will reappear at a different level of theorising—*viz.* in the debate over the nature of the abilities at issue in the kind of cognitive achievement involved in knowing/in the debate over the nature of the abilities at issue in the ability condition (and/or the ability platitude). For it is certainly possible to construe the relevant abilities as involving a reflective element, for instance, in which case the resulting conception of knowledge will have a strong internalist flavour to it. Alternatively, the abilities can be construed in such a way that even creatures with no capacity for reflection at all may possess them, in which case the resulting conception of knowledge will almost certainly be externalist in nature.[6] Even if the defender of the minimalist framework has not succeeded in closing the book on the internalism/externalism debate, the shift of that debate to the debate over the nature of the cognitive abilities at issue in the cognitive achievement involved in knowing/the ability condition (and/or ability platitude) brings with it a major advantage for the analytical project. For the internalism/externalism debate is now no longer a debate on what conditions the correct analysis of knowledge features but rather a debate on how one of the conditions for knowledge is best spelled out in detail. That means, however, that we can now hope to state the conditions for knowledge without having to settle the debate between internalists and externalists.

But now notice that if the arguments I have provided are sound, then it is not only the case that the minimalist framework admits of development in such a way as to allow us provide a substantial condition on knowledge that, *arguably*, solves Gettier problem. There is also reason to believe that the conditions for knowledge can be stated

---

[6] Pritchard (in Pritchard (2005), ch. 7) argues that this is exactly what happens in the debate between reliabilist and responsibilist virtue theories of knowledge.

without settling the internalism/externalism debate. In consequence, if the arguments I have provided are sound, it looks as though we should now be in a position to provide an analysis of knowledge in terms of necessary and sufficient conditions. If all this is indeed correct, then the analysis will take roughly the following shape:

*Analysis of Knowledge*  One knows that *p* iff the truth of one's belief that *p* has been ensured through the exercise of a set of relevant cognitive abilities.

Two points need to be kept in mind here: First, the above analysis of knowledge will be successful only if (but also *if*) the relevant arguments are indeed sound. That is, it will be successful only if (but also if) the crux with Gettier cases really is luck and the argument that if the crux with Gettier cases is luck then the ability condition will predict ignorance in Gettier cases is valid. While by addressing the Gettier problem in this way there is reason to believe that the minimalist who also accepts the ability condition as a plausible development of the minimalist framework does better than the advocates of the various proposed analyses in the post-Gettier literature, the failures of these proposals highlight that one cannot be cautious enough when one professes to have identified a set of individually necessary and jointly sufficient conditions for knowledge. For that reason I would like to stress that this conclusion leaves me somewhat uncomfortable—it seems more like an unexpected and somewhat unwelcome consequence of something I have accepted earlier on—and I would much prefer to state it in conditional form.

  At the same time, second, even if the conclusion arrived at (in its detached version) is true—that is, even if the analysis of knowledge provided is correct—there are still substantial epistemological issues to be settled—such as the issues over how the nature of the cognitive abilities at issue in the analysis is to be spelled out in detail. The analysis arrived at is, in minimalist spirit, still fairly minimalist. The analytical project has been disentangled from other substantive disagreements about the nature of

knowledge. Thus, while the fact that I am committed to the analysis of knowledge just presented leaves me rather uncomfortable, the discomfort is somewhat alleviated by the fact that the analysis proposed is fairly minimal. It is noteworthy here that there is a well-known analysis that is minimal enough to be widely accepted amongst epistemologists—*viz.* the analysis of knowledge as non-accidental or non-lucky true belief (where the non-accidentality condition is the anti-Gettier condition). This analysis allows for so much disagreement that it has been accepted by most epistemologists.[7] In the same way I hope that the analysis presented here would be minimal enough (by allowing for enough disagreement) to gain equally widespread acceptance.[8]

However, let us suppose for a moment that the analysis of knowledge just presented is flawed. Let us suppose, instead, that Williamson's suspicions turn out to be true and knowledge does not admit of analysis in terms of individually necessary and jointly sufficient conditions. If so, then we still have the minimalist framework defended here that will allow us to locate the concept of knowledge in the conceptual landscape. Moreover, if the minimalist framework is rich enough, that is, if it contains either all platitudes about knowledge or enough platitudes to allow derivation of all the other platitudes about knowledge, then there is reason to believe that the minimalist framework pins down the precise location of the concept of knowledge in the conceptual landscape. If so, then we can also provide an alternative definition of knowledge, a definition that does not provide a set of individually necessary and jointly sufficient conditions for knowledge. Instead knowledge can be defined *functionally* in

---

[7] It is surprising that this analysis of knowledge is so widely accepted because it does nothing to address the debate between internalists and externalists. Since, first, this analysis retains beliefs as the primary focus of epistemic evaluation, and since, second, the non-accidentality condition is supposed to be the anti-Gettier condition, it remains an open question whether the three conditions presented ought to be supplemented by an internalist justification condition. Accordingly, it remains an open question whether the three conditions really present a satisfactory analysis of knowledge.

[8] In this way, the analysis arrived at here can be viewed as offering a framework for further epistemological research. In this regard it is similar to the project Greco pursues in 'Agent Reliabilism' (1999) in which he offers a virtue theoretic conception of knowledge as a general framework for epistemology. Notice, that the details of the two approaches differ dramatically in the details.

the following way: Take the conjunction, *A*, of all platitudes about knowledge (or, alternatively, the set of platitudes from which all other platitudes about knowledge can be derived). The conjunction will contain both *K* terms—that is, terms that the functional definition is to define such as "know", "knower" etc.—and *O* terms—that is, terms in the old, already understood vocabulary such as "truth", belief" etc. *A* can the be written down in one sentence:

$A(K_1, \ldots, K_n, O_1, \ldots, O_n)$

Now replace the *K* terms with variables and existentially quantify over each variable. The resulting sentence, the so-called Ramsey-sentence, can be used to provide a functional definition of the terms "know", "knower" etc.: For instance, where "$k_1$" is a variable standing for "know", "knows" can be defined roughly in the following way:

(FAK)  S knows that $p =_{df.} \exists k_1, \ldots, \exists k_n [A(k_1, \ldots, k_n, O_1, \ldots, O_n)$ & S stands in relation $k_1$ to $P]$[9]

Of course, if knowledge does not admit of analysis in the traditional sense, we have little to no reason to believe that the minimalist framework defended here will be so rich as to contain all the platitudes about knowledge or even enough platitudes to allow derivation of all the other platitudes about knowledge (although, of course, it just might). Consequently, we have little to no reason to believe that Ramseyfying the platitudes of the minimalist framework presented here will allow us to provide a correct functional definition of knowledge. However, the functionalist approach still gestures at how the minimalist framework (if suitably extended) can be used to arrive at a definition of knowledge even if knowledge does not admit of analysis in the traditional sense. Therefore, whether or not knowledge admits of analysis in the traditional sense,

---

[9] Again, the inspiration for the functional definition of knowledge comes from the theory of truth—or, to be more precise from Michael Lynch's alethic functionalism. (Lynch has defended a functionalist approach to truth in Lynch (2000), (2001*a*), (2004) and (2005).) This passage follows Lynch's (2001*a*, p. 732) functional definition of truth rather closely.

the minimalist framework is well suited to make a contribution to answering the question as to what knowledge is.

# References

Armour-Garb, B.P. and Beall, J.C. (2005), 'Deflationism: The Basics' in Armour-Garb, B.P. and Beall, J.C. (eds.) (2005).

Armour-Garb, B.P. and Beall, J.C. (eds.) (2005), *Deflationary Truth*, Peru/Illinois: Open Court.

Bach, K. and Harnish, R. (1979), *Linguistic Communication and Speech Acts*, Cambridge/Mass.: Harvard University Press.

BonJour, L. (1985), *The Structure of Empirical Knowledge*, Cambridge/Mass.: Harvard University Press.

—— and Sosa E. (2003), *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, Oxford: Blackwell.

Chisholm, R. (1977), *Theory of Knowledge*, 2nd Edition, Englewood Cliffs, NJ: Prentice-Hall, Inc.

—— (1989), *Theory of Knowledge*, 3rd edition, Englewood Cliffs/NJ.: Prentice-Hall.

Code, L. (1987), *Epistemic Responsibility*, Hanover/NH: University Press of New England.

Cole, P. and Morgan, J. (eds.) (1975), *Syntax and Semantics,* Vol. 3, New York: Academic Press.

Comesaña, J. (2005), 'Unsafe Knowledge' in *Synthese* 146: 195-404.

Conee, E. and Feldman, R. (1985), 'Evidentialism' in Conee, E. and Feldman, R. (2004).

—— (2001), 'Internalism Defended' in Conee, E. and Feldman, R. (2004).

—— (2004), *Evidentialism. Essays in Epistemology*, Oxford: OUP.

Craig, E. (1990), *Knowledge and the State of Nature. An Essay in Conceptual Synthesis*, Oxford: Clarendon.

Dancy, J. (ed.) (1979), *Papers on Language and Logic*, Keele: Keele University Press.

—— (1985), *Introduction to Contemporary Epistemology*, Oxford: Blackwell.

DePaul, M. and Zagzebski, L. (eds.) (2003), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford: OUP.

DeRose, K. (1991), 'Epistemic Possibilities' in *Philosophical Review* 100: 581-605.

—— (2002), 'Knowledge, Assertion, and Context' in *Philosophical Review* 111: 167-203.

Dretske, F. (1970), 'Epistemic Operators', in *Journal of Philosophy* 67: 1007-1023.

—— (1971), 'Conclusive Reasons' in *Australasian Journal of Philosophy* 49: 1-22.

Engel, M. (1992), 'Is Epistemic Luck Compatible with Knowledge?' in *Southern Journal of Philosophy* 30: 59-75.

Feldman, R. (2005), 'Justification is Internal' in Steup, M. and Sosa, E. (eds.) (2005).

Floridi, L. (2004), 'On the Logical Unsolvability of the Gettier Problem' in *Synthese* 142: 61-79.

Foley, R. (1996), 'Knowledge is Accurate and Comprehensive Enough True Belief', in Kvanvig, J. (ed.) (1996).

—— (2002), 'What must be added to true belief in order to have knowledge? Answer: More true beliefs', *Rutgers Epistemology Conference*, April 2002, (manuscript online at: www.philosophy.rutgers.edu/EVENTS/EPIS2002/foley.pdf).

Frankfurt, H. (1969), 'Alternate Possibilities and Moral Responsibility' in *Journal of Philosophy* 66: 829-39.

Gettier, E. (1963), 'Is Justified True Belief Knowledge?', in *Analysis*, 23: 121-3.

Goldman, A. (1967), 'A Causal Theory of Knowing' in *Journal of Philosophy* 64: 355-72.

—— (1976), 'Perceptual Knowledge and Discrimination' in *Journal of Philosophy* 73: 771-791.

——— (1999), 'Internalism Exposed' in *Journal of Philosophy* 96: 271-93.

Greco, J. (1994), 'Review: *The Intellectual Virtues and the Life of the Mind* by Jonathan L. Kvanvig' in *Philosophy and Phenomenological Research* 54: 973-76.

——— (1999), 'Agent Reliabilism' in Tomberlin, J. (ed.) (1999).

——— (2000), *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge: CUP.

——— (2003), 'Knowledge as Credit for True Belief' in DePaul, M. and Zagzebski, L. (eds.) (2003).

——— (2005), 'Justification is Not Internal' in Steup, M. and Sosa, E. (eds.) (2005).

——— (2007), 'The Nature of Ability and the Purpose of Knowledge' *manuscript*.

——— (forthcoming), 'The Value Problem' in Haddock, A., Millar, A. and Pritchard D. (eds.) (forthcoming).

Grice, P. (1991), *Studies in the Way of Words*, Cambridge/Mass.: Harvard University Press.

Guthrie, W.K.C. (ed. and trans.) (1956), *Protagoras and Meno*, London: Penguin.

Haddock, A., Millar, A. and Pritchard D. (eds.) (forthcoming), *The Value of Knowledge*, Oxford: OUP.

Harman, G. (1973), *Thought*, Princeton/NJ: Princeton University Press.

Harper, W. (1996), 'Knowledge and Luck' in *Southern Journal of Philosophy* 32: 79-84.

Hawthorne, J. (2004), *Knowledge and Lotteries*, Oxford: OUP.

——— (2005), 'The Case against Closure' in Steup, M. and Sosa, E. (eds.) (2005).

——— and Szabó Gendler, T. (eds.) (2005), *Oxford Studies in Epistemology*, Vol. 1, Oxford: OUP.

Heider, F. (1988), *The Notebooks* (ed. by Benesh-Weiner, M.), München-Weinheim: Psychologie Verlags Union.

Hetherington, S. (1998), 'Actually Knowing' in *Philosophical Quarterly* 48: 543-69.

——  (1999), 'Knowing Failably' in *Journal of Philosophy* 96: 565-87.

——  (2001), *Good Knowledge, Bad Knowledge. On Two Dogmas of Epistemology*, Oxford: Clarendon.

——  (ed.) (forthcoming), *Epistemology Futures*, Oxford: OUP.

Horwich, P. (1990), *Truth*, Cambridge: Blackwell.

Janoff-Bulman, R. (1992), *Shattered Assumptions: Towards a New Psychology of Trauma*, New York: Free Press.

Johnson, J.T. (1986), 'The knowledge of what might have been: Affective and attributional consequences of near outcomes' in *Personality and Social Psychology Bulletin* 12: 51-62.

Kelp, C. and Pritchard, D. (forthcoming), 'Two Deflationary Approaches to Fitch-Style Reasoning' in Salerno, J. (ed.) (forthcoming).

Kirkham, R. (1984), 'Does the Gettier Problem Rest on a Mistake?' in *Mind* 93: 501-13.

Kornblith, H. (2001), *Epistemology: Internalism and Externalism*, Oxford: Blackwell.

Kvanvig, J. (1992), *Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Contemporary Epistemology*, Savage/Md.: Rowman & Littlefield.

——  (ed.) (1996), *Warrant in Contemporary Epistemology: Essays in Honour of Plantinga's Theory of Knowledge*, Savage: Rowman & Littlefield.

Lackey, J. (forthcoming), 'Why We Don't Deserve Credit for Everything We Know' in *Synthese*.

Lehrer, K. (1965), 'Knowledge, Truth and Evidence' in *Analysis* 25: 168-75.

Lycan, W. (forthcoming), 'On the Gettier Problem Problem' in Hetherington, S. (ed.) (forthcoming).

Lynch, M.P. (2000), 'Alethic Pluralism and the Functionalist Theory of Truth' in *Acta Analytica* 15: 195-214.

—— (2001*a*), 'A Functionalist Theory of Truth' in Lynch, M.P. (ed.) (2001*b*).

—— (ed.) (2001*b*), *The Nature of Truth*, Cambridge/Mass.: MIT Press.

—— (2004), 'Truth and Multiple Realizability' in *Australasian Journal of Philosophy* 82: 384-408.

—— (2005),'Alethic Functionalism and Our Folk Theory of Truth', in *Synthese*, 145: 29-43.

McDowell, J. (1995), 'Knowledge and the Internal' in *Philosophy and Phenomenological Research* 55: 877-93.

—— (2002), 'Knowledge and the Internal Revisited' in *Philosophy and Phenomenological Research* 64: 22-30.

McGinn, C. (1984), 'The Concept of Knowledge' in McGinn, C. (1999).

—— (1999), *Knowledge and Reality*, Oxford: Clarendon.

Montmarquet, J. (1993), *Epistemic Virtue and Doxastic Responsibility*, Lanham/Md.: Rowman & Littlefield.

Moore, G.E. (1925), 'A Defence of Common Sense' in Moore, G.E. (1959).

—— (1939), 'A Proof of the External World' in Moore, G.E. (1959).

—— (1959), *Philosophical Papers*, London: Allen & Unwin.

Moser, P. (ed.) (2002), *The Oxford Handbook of Epistemology*, Oxford: OUP.

Nagel, T. (1979*a*), 'Moral Luck' in Nagel (1979*b*).

—— (1979*b*), *Mortal Questions*, Cambridge: CUP.

Neta, R. and Pritchard, D. (forthcoming), 'McDowell and the New Evil Genius' in *Philosophy and Phenomenological Research*.

Neta, R. and Rohrbaugh, G. (2004), 'Luminosity and the Safety of Knowledge' in *Pacific Philosophical Quarterly* 85: 396-406.

Nozick, R. (1981), *Philosophical Explanations*, Oxford: OUP.

Plato (1956), 'The Meno' in Guthrie, W.K.C. (ed. and trans.) (1956).

Radford, C. (1966), 'Knowledge—by examples' in *Analysis* 27: 1-11.

Pritchard, D. (2001), 'The Opacity of Knowledge' in Shaeffer, H.B. (ed.) (2001).

—— (2004*a*), 'Epistemic Deflationism' in *Southern Journal of Philosophy* 42: 1-32.

—— (2004*b*), 'Epistemic Luck' in *Journal of Philosophical Research* 29: 193-222.

—— (2005), *Epistemic Luck*, Oxford: OUP.

—— (2007), 'The Value of Knowledge', *manuscript*.

—— (forthcoming*a*), 'Anti-Luck Epistemology', in *Synthese*.

—— (forthcoming*b*), 'Greco on Knowledge: Virtues, Contexts, Achievements' in *Philosophical Quarterly*.

—— and Smith, M. (2004), 'The Psychology and Philosophy of Luck' in *New Ideas in Psychology* 22: 1-28.

Radford, C. (1966), 'Knowledge—by examples', in *Analysis*, 27: 1-11.

Riggs, W. (2002), 'Reliability and the Value of Knowledge' in *Philosophy and Phenomenological Research* 64: 79-96.

—— (forthcoming), 'Why Epistemologists Are So Down on Their Luck' in *Synthese*.

Ross, M. and Ross, G. (eds.) (1990), *Doubting: Contemporary Perspectives on Skepticism*, Dordrecht: Kluwer.

Russell, B. (1948), *Human Knowledge: Its Scope and Limits*, New York: Simon and Schuster.

Salerno, J. (ed.) (forthcoming), *New Essays on the Knowability Paradox*, Oxford: OUP.

Sartwell (1991), 'Knowledge is Merely True Belief', in *American Philosophical Quarterly*, 28: 157-65.

—— (1992), 'Why Knowledge is Merely True Belief' in *Journal of Philosophy*, 89: 167-80.

Schaffer, J. (2005), 'Contrastive Knowledge' in Hawthorne, J. and Szabó Gendler, T. (eds.) (2005).

Searle, J. (1975), 'Indirect Speech Acts' in Cole, P. and Morgan, J. (eds.) (1975).

Shaeffer, H.B. (ed.) (2001), *Essays in Philosophy, Vol. 2: The Internalism/Externalism Debate in Epistemology*, Berlin: Humboldt University Press.

Shope, R. (1978), The Conditional Fallacy in Modern Philosophy' in *Journal of Philosophy* 75: 397-413.

—— (1983), *The Analysis of Knowing: A Decade of Research*, Princeton/NJ.: Princeton University Press.

Slaght, R.L. (1977), 'Is Justified True Belief Knowledge?: A Selective Critical Survey of Recent Work' in *Philosophy Research Archives* 3: 1-135.

Slote, M. (1979), 'Assertion and Belief' in Dancy, J. (ed.) (1979).

Sosa, E. (1988*a*), 'Knowledge in Context, Skepticism in Doubt' in Tomberlin, J. (ed.) (1988).

—— (1988*b*), 'Beyond Skepticism, to the Best of Our Knowledge' in *Mind* 97: 153-89.

—— (1991*a*), 'Intellectual Virtue in Perspective' in Sosa, E. (1991*b*).

—— (1991*b*), *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: CUP.

—— (1999), 'How to Defeat Opposition to Moore' in J. Tomberlin (ed.) (1999).

—— (2002), 'Tracking, Competence, and Knowledge' in Moser, P. (ed.) (2002).

Stanley, J. (2005), *Knowledge and Practical Interest*, Oxford: OUP.

Statman, D. (1991), 'Moral and Epistemic Luck' in *Ratio* 4: 146-156.

—— (ed.) (1993), *Moral Luck*, Albany: State University of New York Press.

Steup, M. (2004), 'Internalist Reliabilism' in *Philosophical Issues* 14: 403-25.

—— and Sosa, E. (2005), *Contemporary Debates in Epistemology*, Oxford: Blackwell.

Stoljar (1997), 'The Deflationary Theory of Truth' in Zalta, E.N. (ed.), *The Stanford Encyclopedia of Philosophy*,

URL = [http://plato.stanford.edu/archives/fall1997/entries/truth-deflationary/](http://plato.stanford.edu/archives/fall1997/entries/truth-deflationary/).

Teigen, K.H. (1995), 'How Good is Good Luck? The Role of Counterfactual Thinking in the Perception of Lucky and Unlucky Events' in *European Journal of Social Psychology* 35: 156-71.

Tomberlin, J. (ed.) (1988) *Philosophical Perspectives, 2: Epistemology*, Atascadero/CA.: Ridgeview.

—— (ed.) (1999), *Philosophical Perspectives, 13: Epistemology*, Oxford: Blackwell.

Unger, P. (1968), 'An Analysis of Factual Knowledge' in *Journal of Philosophy* 65: 157-170.

—— (1975), *Ignorance: A Case for Scepticism*, Oxford: Clarendon.

Vogel, J. (1990), 'Are there Counterexamples to the Closure Principle?' in Ross, M. and Ross, G. (eds.) (1990).

Weatherson, B. (2003), 'What are Good Counterexamples?' in *Philosophical Studies* 115: 1-31.

Weinberg, J.M., Stich, S., and Nichols, S. (2001), 'Normativity and Epistemic Intuitions' in *Philosophical Topics* 29: 429-60.

Weiner, M. (2005), 'Must We Know What We Say?' in *Philosophical Review* 114: 227-251.

Williams, M. (1996), *Unnatural Doubts. Epistemological Realism and the Basis of Scepticism*, Princeton/NJ: Princeton University Press.

Williamson, T. (1996), 'Knowing and Asserting' in *Philosophical Review* 105: 489-523.

—— (2000), *Knowledge and Its Limits*, Oxford: OUP.

Wright, C. (1994), *Truth and Objectivity*, Cambridge/Mass.: Harvard University Press.

Zagzebski, L. (1994), 'The Inescapability of Gettier Problems' in *Philosophical Quarterly* 44: 65-73.

—— (1996), *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: CUP.