

University of Groningen

Luck and justification

de Grefte, Job

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:
2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

de Grefte, J. (2017). *Luck and justification: Internalism, externalism and kinds of luck*. [Thesis fully internal (DIV), University of Groningen]. Rijksuniversiteit Groningen.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

LUCK AND JUSTIFICATION

© Job de Grefte, Groningen, 2017

Cover image: *Allegory of Fortune*, c. 1530, by Dosso Dossi (c. 1490-1542). Oil on canvas, 179,1 x 217,2 cm.

Not only does the painting nicely illustrate the separate but related nature of chance (figure on the left) and fortune (figure on the right), it was recognized as an important, long lost allegorical scene of the Ferrarese master Dosso Dossi after having been found at a flea-market and strapped unwrapped to the roof of a car, making it a matter of luck that we can marvel at it today.

Typesetting: This document was typeset using (a modified version of) the typographical look-and-feel classicthesis developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*".

ISBN: 978-90-367-9995-9
(e-book) ISBN: 978-90-367-9995-6



rijksuniversiteit
 groningen

Luck and Justification

Internalism, Externalism and Kinds of Luck

Proefschrift

ter verkrijging van de graad van doctor aan de
Rijksuniversiteit Groningen
op gezag van de
rector magnificus prof. dr. E. Sterken
en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op

donderdag 7 september 2017 om 16.15 uur

door

Job Arent Maria de Grefte

geboren op 22 juni 1989
te Amersfoort

Promotores

Prof. dr. C. Dutilh Novaes

Prof. dr. A.J.M. Peijnenburg

Beoordelingscommissie

Prof. dr. D. Pritchard

Prof. dr. S.C. Goldberg

Prof. dr. B.P. Kooi

CONTENTS

ACKNOWLEDGEMENTS	vii
1 INTRODUCTION	1
1.1 Luck and Justification	2
1.2 Methodology	4
1.3 Structure	8
2 INTERNALISM AND EXTERNALISM	13
2.1 The Internalism/Externalism-debate	14
2.2 Accessibility-internalism	21
2.3 Mentalism	33
2.4 Reliabilism	43
2.5 Virtue Epistemology	52
2.6 Concluding Remarks	57
3 A MODAL ACCOUNT OF LUCK	59
3.1 Pritchard's Modal Account of Luck	59
3.2 An Alternative Modal Account of Luck	67
3.3 Rival Accounts of Luck	81
3.4 Objections To the Modal Account	90
3.5 Concluding Remarks	98
4 EPISTEMIC LUCK	99
4.1 Benign Forms of Epistemic luck	100
4.2 Malignant Forms of Epistemic Luck	108
4.3 Reflectively Lucky Knowledge	128
4.4 Concluding Remarks	136
5 INTERNALISM AND LUCK	139
5.1 Accessibilism	142
5.2 Mentalism	150
5.3 Concluding Remarks	162

6	EXTERNALISM AND LUCK	165
6.1	Reliabilism	166
6.2	Virtue Epistemology	173
6.3	Externalism and reflective luck	177
6.4	Factive Accounts of Justification	180
6.5	Safety and luck	186
6.6	The Internalism/Externalism-Debate Revisited	190
6.7	Concluding Remarks	195
7	JUSTIFICATION AND LUCK	197
7.1	Type I and Type II justification	199
7.2	False Beliefs and Veritic Risk	207
7.3	Good and Bad Veritic Luck	215
7.4	Scepticism	219
7.5	Concluding Remarks	224
8	CONCLUSIONS	227
	BIBLIOGRAPHY	233
	SUMMARY/SAMENVATTING	247

ACKNOWLEDGEMENTS

*No man is an island, entire of itself; every man
is a piece of the continent, a part of the main;*

— John Donne

Wat voor mensen geldt, geldt in dit geval ook voor proefschriften; zij bestaan niet in een vacuüm, maar vinden hun plaats in een netwerk van menselijke relaties. Ik wil de volgende personen in het bijzonder bedanken voor de verbindingen die ze tot stand hebben gebracht.

Catarina, ik wil jou ontzettend bedanken voor je bereidheid mij onder je hoede te nemen en de manier waarop je dat gedaan hebt. Ik heb mij vanaf het begin af aan welkom en gewaardeerd gevoeld. Bedankt voor je bereikbaarheid, waar je ook bent. Bedankt voor je brede interesse en voor het vertrouwen in mijn keuze voor de epistemologie. Bedankt voor het lezen van mijn talloze drafts. Bedankt voor je altijd nuttige feedback, voor je creativiteit en je waakzaam oog op het gevaar van antropocentrisme in mijn proefschrift. Bedankt voor je leesgroep en je hulp bij het zoeken naar nieuwe verbindingen. Bedankt voor de diners, en voor de karaokeavonden. Ik heb meer van je geleerd dan ik zou kunnen benoemen.

Jeanne, ook jou ben ik ongelooflijk dankbaar. Voor je vertrouwen in mij, gedurende het project, maar ook toen ik nog student was. Voor je actieve betrokkenheid bij mijn project als tweede promotor. Voor het feit dat ik altijd bij je terecht kon. Voor het op de rem trappen als ik te snel wilde. Voor het geduld met mijn talloze drafts. Voor het steeds weer opnieuw en even nauwkeurig doorspitten van mijn manuscripten. Voor je integriteit. Voor het zijn van

een baken van rust. Voor je waardevolle commentaar. Zonder jou had mijn proefschrift er heel anders uitgezien.

Duncan, you too prove Donne to be right. For even though you were literally on an island, I could not have felt more connected to you. Not only would this thesis not have been possible without the towering achievement of your work, you were willing to discuss it at length with me in person on numerous occasions, even though you were under no obligation to do so. I enjoyed these conversations to the fullest, for they were as pleasant as they were stimulating. I want to thank you for your continuing support and encouragement.

I also would like to thank the other members of my reading committee. Sandy, I feel privileged that you agreed to assess my thesis, and I am honoured that you approve of its contents. Thank you for your providing me with feedback that is as sharp and on-point as it is helpful and elaborate. Thank you for taking the time. Barteld, jij ook bedankt voor je nuttige en uitgebreide commentaar op mijn manuscript. Naast de inhoudelijke punten heeft het ettelijke spelfouten weten te voorkomen. Bedankt voor je geduldige uitleg en je kritische blik op mijn argumenten. Ze zijn er beter door geworden.

Tony, thank you for your support in especially the first part of my project. You were extremely helpful when I was just starting to get to grips with the debate between internalists and externalists. Your comments have set the project on its current course, and without them this thesis would have looked substantially different. I hope to see you soon so that we can discuss the final product.

Er zijn twee netwerken die hier in het bijzonder vermeld dienen te worden. Ten eerste de maandelijkse Work-in-Progress bijeenkomsten. Het commentaar dat tijdens deze bijeenkomsten geleverd wordt is van een ongekend hoog en kritisch niveau, en mijn verschillende hoofdstukken hebben er veel aan gehad op dit podium besproken

te worden. Het was een eer deze bijeenkomsten voor te zitten, en ik wil dan ook alle deelnemers bedanken, maar in het bijzonder Bouke en Leon, die altijd ongeëvenaard scherp commentaar geleverd hebben, ongeacht het onderwerp. Ten tweede verdient ook de PromotieClub Cognitieve Patronen (of hoe zij ook mogen heten) vermelding. Na presentatie voor dit specialistengezelschap komt men uitgeput thuis, maar met een veel beter paper.

Leon, bedankt voor je vriendschap en voor je niet aflatende discussiedrift. Als academische broertjes zaten we vaak in dezelfde situatie, en ik heb veel steun aan je gehad. Geen theorie was jou onbekend, geen filosoof ongelezen, en zo ontstonden er spontaan verbindingen tussen Hegel en Hintikka. Dank voor deze gesprekken, voor de sushi en voor de karaoke.

Marga, ook jou wil ik expliciet bedanken. Voor je snelle respons en behulpzaamheid. Meer dan eens maakte jij mogelijk waar ik al niet meer op gerekend had. Jouw deur staat altijd open (letterlijk), en onlangs heb je zelfs die deur van jou opgegeven voor meer collegieruimte op onze faculteit. Het is bijzonder een directeur te zien met zó'n hart voor de zaak. Chapeau.

Ik bedank ook mijn paranimfen, Coos en Sander. Coos, jij bent in de jaren dat we samen aan het promoveren waren een gewaardeerde collega geworden, maar bovenal een goede vriend. Samen hebben we aardig wat mijlpalen meegemaakt (huwelijk, samenwonen, huis kopen, en nu ook promoveren), iets dat ons altijd weer gespreksstof verschafte voor bij de koffie. Niet dat die gesprekken altijd over serieuze zaken gingen: vaak kwam Boer Zoekt Vrouw voorbij of het verschil tussen in en óp de bibliotheek zitten. In ieder geval klonk jouw schaterlach herhaaldelijk, iets waar niet alleen ik, maar ook toevallig aanwezige omstanders zich vrolijk over konden maken. Omdat onze onderwerpen zo dicht tegen elkaar aanlagen hebben we inhoudelijk ook veel aan elkaar gehad, en ik kijk er dan ook

naar uit onze samenwerking in de toekomst voort te zetten. Bedankt voor alle mooie momenten.

Sander, ook met jou heb ik veel gelachen en ook van jou heb ik veel geleerd. Onze Lucky Queergress leesgroep was fantastisch, en het was een eer je paranimf te zijn. Ik hoop dat ik het net zo goed als jij zal doen, tijdens mijn verdediging, maar ook tijdens het vervolg van mijn carrière. Je bent een goede vriend en een voorbeeld voor mij, bedankt daarvoor.

Pieter en Tom, ook jullie ontzettend bedankt voor de talloze diners, borrels en gezellige (beschonken) avonden. Voor de nuttige zowel als de nutteloze discussies. Voor de roddel en de achterklap. Voor de tips en het advies. Bedankt, jongens, voor het *plezier*.

Ook mijn overige collega's bedank ik van harte. De filosofiefaculteit aan de Oude Boteringestraat is een fantastische plek om te werken, en dat ligt met name aan de mensen die er rondlopen.

Er is ook leven buiten de filosofie dat onlosmakelijk met dit proefschrift is verbonden. Bijzondere vermelding verdienen mijn ouders, schoonouders, broers en zus. Bedankt voor alle hulp en voor de stimulerende omgeving.

Ik bedank ook mijn vrienden graag, want zonder vrienden kan ik niet.

En tenslotte, Madelein, jou ben ik onbeschrijflijk dankbaar, en het meest van allen. In jouw dankwoord gaf je mij een quote van Shakespeare, ik geef je er een van Wittgenstein terug:

Wovon man nicht sprechen kann, darüber muss man schweigen.

INTRODUCTION

We all know many things. We know, for example, that the earth is round, and that Paris is the capital of France. We know that we have hands, and we know that it is our own face that stares back at us from the mirror each morning when we get out of bed. We further know that $1 + 1 = 2$, and that nothing can be both green and red all over.

We do not rest content with what we know: we perpetually want to know *more*. Collectively, we invest vast sums of money in the pursuit of knowledge in the form of science budgets. Individually, we read books, magazines, blogs, newspapers. We watch films, series, documentaries and instruction video's. We travel the world. All at least partly in the pursuit of knowledge.

We talk *a lot* about what we know. According to the Oxford English Dictionary, the verb 'know' is the eighth most frequently used verb in the English language.¹

What this all shows is that the concept of knowledge plays an extremely important part in our lives. We care about what we ourselves know, but also about what others know. We value knowledge.

It is no wonder, then, that among the oldest questions of philosophy is the question what it is exactly that we are so interested in, the question: "What is knowledge?". It is this question, ultimately, that the present study is concerned with.

¹ See (*What can the Oxford English Corpus tell us about the English language?*, n.d.). This is not a peculiarity of the English language: across the world, usage of verbs with the same meaning is extremely frequent (e.g. Davies, 2005; MasterRussian.com, 2016).

1.1 LUCK AND JUSTIFICATION

While some philosophers have claimed otherwise (e.g. Williamson, 2000), significant progress has been made in the philosophical study of knowledge since antiquity. For example, we know that the kind of knowledge in play in the examples at the beginning of this chapter is of a particular kind: it is *propositional* knowledge. Propositional knowledge, sometimes also called 'knowledge-that', is knowledge *that* a certain proposition is true. It is knowledge *that* the cat is on the mat, or *that* the square root of 9 is 3. It is knowledge concerning facts, knowledge *that* the world is one way rather than another.

Propositional knowledge is distinct from other forms of knowledge, such as knowledge by acquaintance. The latter kind of knowledge is exemplified when it is true that I know Mary. When I know Mary, there is no particular proposition that I know to be the case, no proposition that can be embedded in a that-clause such that it is true that I know that ... is the case. I simply know who Mary is. Similarly, philosophers have distinguished knowledge-that from knowledge-how, the kind of knowledge that you have when you know *how* to do something, although this distinction is more controversial.² In any case, the kind of knowledge we will be concerned with throughout this study is propositional knowledge.

In analytic epistemology, propositional knowledge is widely held to be factive: you cannot know something unless it is true. Thus, one cannot know that the earth is flat if in fact it is round. One cannot know that that the capital of France is Nice if in fact it is Paris.

It is also widely held that knowledge requires belief on the part of the knowing agent. So one cannot know that the earth is round if one does not believe this to be the

² Cf. (Stanley & Williamson, 2013).

case. Of course, not everything one believes is true, so one may believe something to be the case without knowing it.

So far, the sketched necessary conditions on knowledge are relatively uncontroversial.³ Knowledge requires both truth and belief. However, knowledge requires more: one can believe truly that the capital of France is Paris on the basis of a lucky guess. Such beliefs, even if true, do not constitute knowledge. So, for knowledge it matters how one's beliefs are *formed*. Not every formation of true beliefs results in knowledge.

Many have labelled the property of true beliefs that turns them into knowledge the property of *justification*. The notion of justification as it features in epistemology—*epistemic* justification—should be distinguished from other kinds of justification, such as moral or pragmatic justification. One can be morally justified in believing that all people are equal, even if one believes so on the basis of no evidence whatsoever, so that one's belief is not *epistemically* justified. So these other forms of justification do not entail epistemic justification. Nor does epistemic justification entail that one is morally or pragmatically justified. As William James famously argued, one can be pragmatically justified in believing one will get better, even if this belief flies in the face of all available evidence (1897).

I said that whether one is epistemically justified depends on how one forms one's beliefs. If one does so in an appropriate way, and one ends up with a true belief as the result of this, one's belief may qualify as knowledge. As Gettier famously showed, however, satisfaction of these three conditions (truth, belief, and justification) does not *guarantee* knowledge (1963). That is, having a justified true belief is not sufficient for having knowledge. The reason for this seems to be that even if one's true belief is justified, it may

³ As anything in philosophy, these conditions are not completely uncontested. For a recent challenge to the factivity of knowledge, see (Hazlett, 2010). For a challenge to the belief-condition, see (Radford, 1966).

still be a matter of *luck* that it is true. So while the justification condition was conceived to exclude ill-formed true beliefs such as lucky guesses from counting as knowledge, Gettier showed that justification nevertheless failed to exclude all problematic forms of luck from knowledge.

This is not the only way to interpret the import of Gettier's paper. However, it will serve as a useful way of introducing our present topic.⁴ For the main project that I shall be engaged in, in the coming chapters, is that of *refining* this standard story. In particular, I will investigate the relation between epistemic justification and luck in order to see the extent to which epistemically justified belief can still be lucky.

It should be noted, right from the start, that I have benefited much from previous studies into the nature of luck and its place in epistemology. In particular, I draw on the work of Duncan Pritchard. One may view the present study as a continuation of his general anti-luck epistemology by investigating the role justification can play in the elimination of luck.

1.2 METHODOLOGY

How will I do so? This question may be interpreted in two ways: a question as to the methodology of the study, and a question as to its structure. This section regards the first, the next section the second interpretation.

This study is an exercise in conceptual analysis. But there are different ways to understand what is involved in that method. As I understand it, the method of conceptual analysis aims to uncover necessary and sufficient conditions for concepts to apply.⁵ These concepts are taken

⁴ For some examples of this kind of narrative, see (Dancy, 1985, p. 134; Booth, 2011, p. 39; Pritchard, 2005a).

⁵ For a defence of a similar account of conceptual analysis, see (Jackson, 2000).

from ordinary language, and so agreement with ordinary language is an important criterion of adequacy for the proposed analysis. The method proceeds by identifying an intuitively plausible set of necessary and sufficient conditions, and modifying them according to the extent that counterexamples can be found. Hopefully, the result will be a set of conditions that is immune from counterexamples, but history teaches that the perpetual modification of set in response to new counterexamples may also be a real possibility.⁶ As such, the conceptual analyses in this book should not be taken as final, except in the sense that they are the most complete analyses that I am presently able to devise.

As said, the concepts that I will conceptually analyse are taken from ordinary language, and our analysis should thus generally accord with ordinary language. But ordinary language is vague and inconsistent. One of the virtues of conceptual analysis is that it brings out *where* exactly vagueness and inconsistency lie.

However, the method of conceptual analysis treats vagueness and inconsistency in different ways. For while our proposed set of necessary and sufficient conditions may include vague notions, it may not be inconsistent. So we may say that it is necessary for a giant to be tall (tall is a vague predicate), but we may not say that it is necessary for a giant to be both tall and not tall, for that reduces our talk of giants to meaningless gibberish. When doing conceptual analysis, there is always a presumption in favour of the claim that the concepts most important to us are meaningful in a literal sense. Our method serves to provide a clear and consistent account of this meaning, seeing how it 'hangs together' with other concepts of importance.

⁶ Compare, in this respect, the method set out in Lakatos' famous "Proofs and Refutations" (1976).

It is because of this that conceptual analysis will be revisionary to a certain extent. If the content of our concepts is inconsistent, conceptual analysis can provide reasons for revision by identifying the necessary and sufficient conditions for *closely related* coherent concepts. If we find, for example, that we can save much of what we want to say about knowledge by using a slightly different concept, knowledge₁, where knowledge₁ is consistent, but our ordinary concept of knowledge is not, this may provide reason to adopt the former in place of the latter.

The method of conceptual analysis is often contrasted to the method of (Carnapian) explication.⁷ Without going into too much detail, the essential difference is that the former aims to provide analyses of the concepts we actually use, whereas the latter aims to construct new concepts that are meant to replace the old ones in certain contexts. Thus, whereas agreement with ordinary usage is an important criterion of adequacy for conceptual analysis, it is less important for explication. It should be noted, however, that conceptual analysis and explication are different ends on a single methodological spectrum. For as we have seen, theoretical considerations (consistency) may play a role in conceptual analysis, and similarly, agreement with ordinary language is among Carnap's criteria for the adequacy of explication. These criteria pull in opposite directions (Dutilh Novaes & Reck, 2017). For often the theoretical considerations will motivate a move further away from the pre-theoretical concept. So while theoretical virtues and agreement with ordinary language may pull in different directions, both methodologies accord some weight to both these aspects. The difference is therefore one of emphasis. Conceptual analysis tends to accord relatively

⁷ See, for example, (Beaney, 2016). In Beaney's terminology, the distinction of relevance here is between two forms of analysis, Carnap's method of explication and Oxford linguistic analysis.

greater weight to agreement with ordinary language, while explication accords greater weight to theoretical virtues.

Seen in this light, there is no essential tension between both projects. We would like to reach an ideal trade-off between agreement with ordinary language and theoretical virtues such as consistency and fruitfulness. What this ideal trade-off is will depend partly on one's objectives. Is one trying to get clear about what it is that we are talking about, or is one looking for a concept that would fit a particular theoretical role? Both projects seem to me worthwhile. While the present investigation leans more towards the former, I certainly do not think this is the only valuable approach to the present subject.

Conceptual analysis draws on intuitions. Both in the formulation of the conditions and of the counterexamples that are raised against them, an appeal is often made to our intuitions. The epistemic value of intuitions has been contested, however. In particular, some philosophers have argued that intuitions are unreliable, or relative to cultural factors.⁸

I don't think intuitions are infallible, nor that they are always reliable. But they can provide direction to investigations, such as conceptual analyses, that are not possible otherwise. Conceptual analysis is to an important extent an investigation into what is consistent with what, and such an investigation will have to refer to facts about what is possible and what is impossible. Intuition is an important, perhaps *the* most important guide we have to these

⁸ See, for example, critique from experimental philosophy, such as (Weinberg, Nichols, & Stich, 2001). In a different vein, Herman Cappelen argues that neither intuitions nor conceptual analysis play the central philosophical role that I here make it out to play (Cappelen, 2012). For discussion of his arguments, I refer the reader to the *Philosophical Studies* symposium devoted to his book, especially to (Chalmers, 2014; Weatherson, 2014).

facts.⁹ That does not mean of course, that every appeal to intuition is as good as the other. Some intuitions are more broadly shared and stronger than others. The former will make for better arguments than the latter. Also, intuitive judgements provide stronger arguments the better they accord with other intuitions, or the better they can be embedded within a large theoretical framework. I take the construction of such frameworks to be among the central tasks of philosophy.

1.3 STRUCTURE

This study is structured in the following way. In Chapter 2, I provide an overview of the internalism/externalism-debate about epistemic justification. Because it is impossible to provide a complete survey of this extensive debate, I focus on four main positions: two internalist (accessibilism, mentalism), two externalist ones (reliabilism, virtue epistemology). I provide an overview of their general commitments, as well as their main motivations. In particular, I argue that these different internalist and externalist theories of justification are motivated by four distinct *concepts* of justification. Whereas accessibilism is motivated by a deontological concept of justification, mentalism is defended by an appeal to an evidentialist concept of justification. On the externalist side, we will see that reliabilism can be defended by an appeal to the truth-conducive concept of justification, whereas agent reliabilism, despite its similarity in name to reliabilism, takes a radically different perspective by conceiving of justification as a particular kind of *competence*. This raises the question whether participants

⁹ See (Chalmers, 1996, 2002, 2010; Menzies, 1998; Yablo, 1993) for defences of this claim. I am using the notion of intuition here broadly, to include both the notions of imagination and conceivability.

to the debate are disagreeing with each other substantially or merely verbally.

My answer to this question draws on my account of the relation between luck and justification, and will have to wait until Chapter 6. First, in Chapter 3, I defend and develop a particular modal account of luck, based on Duncan Pritchard's Modal Account of Luck (MAL). After first providing an overview of MAL, I continue to develop and improve the account in various ways. In particular, I argue that we should incorporate a significance condition on luck, and recognize the fact that luck depends on *proportions* of possibilities. I also argue that there are various ways in which the degree of luck to which an event is subject depends on subjective factors. This gives rise to a problem for accounts of luck parallel to the reference class problem for accounts of probability. I also discuss the two main rivals to a modal account of luck, the lack-of-control account and the probabilistic account of luck, and argue that we should prefer my modal account of luck to either of these alternatives.

In Chapter 4, I focus on various kinds of *epistemic* luck. These come in varieties compatible with knowledge ('benign' forms of luck) and varieties incompatible with knowledge ('malignant' forms). I discuss the following benign forms of epistemic luck: content luck, capacity luck and evidential luck. I discuss two potentially *malignant* kinds of luck: veritic luck and reflective luck. I also discuss two more recent additions to the anti-luck literature: environmental and intervening luck. I draw on Sosa's distinction between 'animal' and reflective knowledge to argue that only veritic luck is incompatible with *any* kind of knowledge. The upshot of this chapter is thus that different forms of luck are compatible with different grades of knowledge.

In Chapter 5, I investigate the relation between internalism about epistemic justification and the two clearly malignant kinds of epistemic luck discussed in the previous

chapter: veritic and reflective luck. In particular, I argue that the two main internalist theories of justification in contemporary epistemological literature are incompatible with reflective luck but not with veritic luck. The findings of this chapter support a novel interpretation of the import of Gettier cases. As we said above, many epistemologists think that Gettier showed that justification could not deliver on the promise of excluding (problematic forms of) luck from knowledge. This chapter shows that we can provide a more nuanced picture. For internalist justification eliminates reflective luck but not veritic luck. The fact that reflective luck is incompatible with the possession of *reflective knowledge* but compatible with the possession of animal knowledge, raises questions about the function of internalist justification. On this picture, there is a kind of knowledge that is incompatible with the kind of luck that internalist justification eliminates. Perhaps the kind of knowledge that philosophers like Plato have been interested in has always been reflective knowledge. For that kind of knowledge, the elimination of reflective luck is necessary, and the results of Chapter 5 indicate that an internalist justification condition is well-suited for that task. What Gettier draws our attention to, then, is precisely the fact that there are different forms of malignant epistemic luck, forms of luck that require different forms of justification to eliminate.

In Chapter 6, I argue that prominent externalist accounts of justification are incompatible with veritic luck but compatible with reflective luck. Further, I consider some implications of this view. First of these, I discuss Zagzebski's claims about the inescapability of Gettier cases for non-factive accounts of justification. I argue that factive accounts of justification are neither necessary nor sufficient to escape Gettier problems. Second, I argue that if externalist concepts of justification are incompatible with veritic luck, this means that externalists about justification can do

without a separate safety condition on knowledge. Since internalist concepts of justification do little to eliminate veritic luck, this will not be true for those who endorse an internalist concept of justification. Finally, I argue that the findings from this chapter and the previous one provide the basis for a novel characterization of the internalism-/externalism-debate about epistemic justification in terms of luck. I argue that this conception of the internalism-/externalism-debate has some advantages over existing ways of characterizing the debate, most notably that it is more informative about the commonalities among internalist and externalist accounts respectively. This way of looking at the internalism-/externalism-debate also allows us to identify a substantial core of disagreement between internalists and externalists (*pace* William Alston (2005)).

After having thus provided a characterization of the relation between various existing accounts of justification and luck, I present my own account of justification in Chapter 7. I distinguish two kinds of justification: Type I and Type II justification. Type I justification requires the elimination of veritic luck. Type II justification requires the elimination of both reflective and veritic luck. I argue for a more general condition on justification in terms of the notion of veritic *risk*, a notion closely related to the notion of veritic luck. I further distinguish between good veritic luck and bad veritic luck, and argue that while all epistemic justification is incompatible with good veritic luck, justification is compatible with bad veritic luck. I finally consider a possible charge of scepticism, and some ways to respond to this charge.

That draws the study to an end. In Chapter 8, I close with some concluding remarks.

INTERNALISM AND EXTERNALISM

What is knowledge? Traditionally, philosophers have thought that in order to know, your belief must be *justified*. But over the course of time opinions have come to diverge widely over how we should understand this claim.

This chapter is about one of the central debates about the nature of epistemic justification: the internalism/externalism-debate. As of late, this debate has not been getting the best of press. In fact, the deplorable state of the debate—evinced by the seemingly inverse proportionality between the progress made and the number of pages written on the subject—has led prominent epistemologists like William Alston to urge us to abandon the concept of epistemic justification altogether (2005). In this study I resist such a pessimistic meta-induction. More concretely, one of my aims is to shed new light on what is at stake in the internalism/externalism-debate by drawing on recent work on *epistemic luck*. My claim will be that doing so will allow for a new perspective on the internalism/externalism-debate itself. Even if internalists and externalists are partly talking past each other, we can still say something interesting about the relation between their concepts of justification and various forms of epistemic luck.

Before I can argue for this claim, however, I need to provide an overview of the internalism/externalism-debate. That is what I will do in this chapter. This overview is not meant to be exhaustive. Rather, I focus on what I take to be the four most prominent theories of epistemic justification: two internalist, and two externalist theories. I will set out their general commitments, and explain how they are

motivated. This will provide precise targets for discussion in the coming chapters.

Of course, this methodology has its limitations. In particular, one may worry that what we find to hold for these archetypes of internalism and externalism need not necessarily hold for other possible internalist or externalist theories of justification. I agree. The simple fact, however, is that the last decades of epistemological research have produced such a vast amount of different theories of epistemic justification that any attempt to be complete in the amount of space appropriate for a project like this would be doomed to fail. Choices thus need to be made. In choosing my archetypes, I have tried to provide a selection of theories that is in my view representative of the debate, by including only concepts that are, or have been, widely endorsed and are clearly distinct from each other.

The chapter is structured in the following way. In Section 2.1, I briefly explain the general features and context of the internalism/externalism-debate about epistemic justification. In Sections 2.2 and 2.3, I discuss the two most prominent internalist theories of justification of the last decades of epistemological literature: theories that I shall name ‘accessibility-internalism’ and ‘mentalism’, respectively. In Section 2.4 and 2.5, I do the same for two prominent externalist theories of justification: ‘reliabilism’ and ‘virtue epistemology’. As said, for all these accounts, I will discuss their central tenets, as well as their main motivations. I will close with some concluding remarks in Section 2.6.

2.1 THE INTERNALISM/EXTERNALISM-DEBATE

What is the internalism/externalism-debate about? In this section, I provide an overview of its general structure. But first, a brief historical tale.

Epistemology is first and foremost the study of knowledge. However, epistemologists' concern with epistemic justification is as old as epistemology itself. The reason for this is that most philosophers studying knowledge have thought (and most still think) that epistemic justification is *necessary* for knowledge. Already in Plato, for example, we find that 'mere' true belief does not suffice for knowledge (Plato, 1973). In the *Theaetetus*, Socrates asks us to imagine a judge, who forms the belief that a particular defendant is innocent on the basis of mere hearsay. Such a judge, Socrates submits, does not come to *know* the defendant is innocent by forming her belief in the way she does, even if it turns out her belief is true. Why not? Because it seems that a true belief based on weak or otherwise defective evidence (such as mere hearsay) is akin to a lucky guess, and we do not think lucky guesses can generate knowledge.¹

If a true belief does not suffice for knowledge, what else is needed? In the *Theaetetus*, Socrates discusses the notion of a *rational explanation* for the truth of the belief in question. This concept can be regarded as a precursor to our modern concept of epistemic justification. While Socrates ultimately dismisses the option, many have thought that suitably reformulated, a notion of epistemic justification *could*, together with true belief, be used to define necessary and sufficient conditions for knowledge. This account of knowledge as justified true belief is known as the classical, or 'tripartite' account of knowledge.

In 1963, however, Edmund Gettier famously argued against the classical account by providing examples of justified, true beliefs that (allegedly) fail to be knowledge.² Among other things, Gettier's short but hugely influential

¹ This will prove important; we will have much more to say about this claim in later chapters.

² Traces of an early challenge the classical account can, according to some, be found in the work of the 11th century Indian philosopher Sriharsa (Matilal, 1986). Another early version is provided by Russell (1948).

paper ignited the internalism/externalism-debate about epistemic justification.

The notion of justification that is present in Gettier cases is typically of the internalist kind. Epistemologists who wanted to save the tripartite definition of knowledge therefore tried to modify the requirements of justification in a way that would render justified true belief immune to Gettier cases. Since the requirements made by these theories often referred to facts that were external to the believing subject, they became known as 'externalist' theories of justification. Thus, according to this story, did Gettier give birth to the internalism/externalism-debate.³

Whether the above story is accurate or not, Gettier cases *can* be used to support externalism, as we will see below. What matters here is that our story provides an entry point to what is at issue in the internalism/externalism-debate. As the story brings out, the debate between internalists and externalists is essentially a debate about what sort of things determine whether a belief is justified. The things that determine justification are commonly called 'justifiers' or 'J-factors'. Somewhat trivially, the internalism/externalism-debate is essentially a debate about whether justifiers are 'internal' or not. Of course, for a theory of justification to be of any value, the relevant sense of 'internal' would need to be spelled out, and this is exactly what various theories of justification have tried to do. Importantly however, the debate is thus not a debate about what it *is* to be epistemically justified in the first place. Rather, it is a debate about what kinds of facts justification supervenes upon. However, as we will see below, different views regarding the latter issue are motivated by different views regarding the former.

None of these, however, had the impact of Gettier's classic paper "Is justified true belief knowledge?" (1963).

³ This story is based on (Poston, 2016).

Within the internalism/externalism-debate, we can distinguish between strong and weak versions of both internalism and externalism. Strong internalism is the position that all justifiers are internal, whereas weak internalism is committed only to the claim that *some* justifiers are internal. Since externalism is usually formulated as the denial of internalism, there are strong and a weak forms of externalism as well. Strong externalism is the denial of weak internalism and consequently holds that all justifiers are external (again, in a sense yet to be specified). Weak externalism is the denial of strong internalism and holds that *some* justifiers are external.⁴ There are thus two possible relevant oppositions between internalism and externalism: the opposition between strong internalism and weak externalism, as well as the opposition between weak internalism and strong externalism. In the literature, the internalism-/externalism-debate is usually depicted as the opposition between strong internalism and weak externalism:

Epistemic internalism is the view that a thinker's epistemic status depends wholly on matters which are 'internal' to that thinker, rather than at least partially on matters which are 'external' to her, such as her relations to her environment. Let epistemic externalism be the denial of epistemic internalism. (Brown, 2007, pp. 13-14)

Internalism asserts that justification is internally determined, whether by evidence possessed, or by coherence among beliefs, or by some other internal condition. Externalism

⁴ Of course, on the current formulation, weak internalism and weak externalism are logically *equivalent* on the assumption that a justifier is either internal or external. Nevertheless, a 'weak' theory may allow for more or less internal justifiers, and we may classify it as weakly internalist or externalist accordingly.

about justification is readily understood as the denial that internal factors are sufficient. Something external has an independent role in justifying beliefs. Justification does not supervene on the internal alone. (Conee, 2004, p. 48)

The choice to focus on the distinction between strong internalism and weak externalism, rather than on the distinction between weak internalism and strong externalism is easily explained: the most widely endorsed internalist theories of justification are of the 'strong' kind, and the most widely endorsed externalist theories of the 'weak' kind. In what follows I will therefore use the term 'internalism' to denote strong internalism, and the term 'externalism' to denote weak externalism. We should keep in mind, however, that strong externalism and weak internalism are also theoretical possibilities.

Thus formulated, the internalism/externalism-debate is thus a debate about the kinds of facts that justification supervenes upon (I call the collection of these facts the 'supervenience base' of epistemic justification). If internalism is right, this supervenience base is provided by the internal alone. If externalism is right, the base will include justifiers that are external to the subject as well.

Saying that internalism is the view that justification depends only on matters which are internal to the thinker is not very informative. The characterization of externalism as the denial of this claim even less so. It could be taken as an indication of the state of the internalism/externalism-debate that this is all that can be said, in general, about internalism and externalism. As I shall argue in later chapters, however, this is *not* all that can be said. But that will have to wait. In the next sections we will see that specific internalist and externalist theories flesh out the specific sense in which justification does or does not depend on the internal alone in different ways.

The internalism/externalism-debate is thus a debate about the question whether epistemic justification supervenes on the internal alone. To get a good handle on this question, however, we need to draw a distinction between propositional and doxastic justification. Propositional justification is a property of a proposition relative to an agent (Bergmann, 2006, p. 4). Different epistemologists will have different views as to what propositional justification amounts to, but one example would be to say that proposition p is propositionally justified for agent S at time t if and only if p is sufficiently supported by the total evidence available to S at t .⁵ What is important is that one does not need to believe a proposition in order for one to have propositional justification for believing it. Nor does one need to be capable of believing a proposition in order for that proposition to be propositionally justified. Thus, some very long and complex logical tautology may be propositionally justified for me, even if I am psychologically incapable of entertaining the proposition, let alone believing it.

Doxastic justification, by contrast, is a property of held *belief*. According to a standard view, agent S is doxastically justified in believing p if and only if S has propositional justification R for p and S bases her belief that p on R .^{6,7}

5 See (Neta, 2007).

6 See (Korcz, 2000; Kvanvig, 2003; Pollock & Cruz, 1999; Swain, 1979). John Turri holds the unorthodox view that propositional justification should be explained in terms of doxastic justification instead of the other way around (Turri, 2010).

7 Propositional and doxastic justification are sometimes distinguished from *personal* justification, which is, as the name indicates, a property of a person rather than a belief. When we say that Sally was justified in believing that her house was on fire, we are talking about personal justification (a property of Sally), whereas when we say that Sally's belief that her house was on fire was justified, we are talking about doxastic justification (a property of Sally's belief). The relationship between personal and doxastic justification is controversial, however. Some epistemologists think that S is personally justified in believing p if and only if

How to interpret this 'basing' requirement is a matter of considerable controversy. Roughly, we may say that a belief B is based on reason R if and only if R is the reason *for which* B is held.

In a sense, both propositional and doxastic justification are 'epistemic'. To be as clear as possible, however, I will reserve the term 'epistemic justification' for doxastic justification only. This reflects standard practice. Swain, for example, has a necessary basing requirement in his definition of epistemic justification (Swain, 1981, p. 134). Similarly, Korcz claims: "Merely possessing good reasons for a belief is not sufficient to make one epistemically justified in holding it. In addition, those reasons must be the reasons *for which* the belief is held", thus implicitly equating epistemic justification and doxastic justification (2000, p. 525, my italics). Also, most accounts of epistemic justification that we will discuss below refer to actually held beliefs rather than to propositions. In addition, they refer to the methods that produced these beliefs.⁸ With the notable exception of Feldman and Conee, whose views we will discuss below, most epistemologists thus seem to reserve the term epistemic justification for doxastic, rather than propositional justification.

Ted Poston argues for the contrary view that internalism should be understood as a thesis about propositional rather than doxastic justification by noting that the best account of the basing relation is formulated in causal terms (Poston, 2016). Since these causal relations typically do not

S's belief that *p* is doxastically justified (Kvanvig & Menzel, 1990). Other epistemologists think there are important differences between personal and doxastic justification (Engel, 1992b; Littlejohn, 2009, 2012). I do not intend to take a stand on this issue, which does not bear on the claims I want to make.

⁸ In this study use the terms 'methods' and 'processes' interchangeably. Alvin Goldman draws a distinction between them (Goldman, 1986, p. 93–95; Goldman, 2015, p. 142). The distinction is not important for our present purposes, however, so I will leave it aside.

supervene on the internal, we should regard the internalist thesis as a thesis about propositional rather than doxastic justification.

However, there are other accounts of basing available to the internalist besides the causal account. Not all of these refer to facts external to the believing subject.⁹ So the internalist is not *forced* to restrict her thesis to propositional justification. As we will see, Conee and Feldman's usage of the term 'epistemic justification' is non-standard. Most epistemologists, including various internalists, refer to doxastic justification when they speak of epistemic justification. In this thesis I will follow the majority and refer to doxastic justification when I use the term 'epistemic justification'.

2.2 ACCESSIBILITY-INTERNALISM

On the present understanding, internalism holds that all factors that determine justification are internal to a believer. That means that justification, according to the internalist, supervenes on what is internal: no two believers could be internally alike yet differ in the justificatory status of any of their beliefs. As we said, this formulation of the internalist thesis is not very informative, since it is not clear what the relevant sense of 'internal' is. The literature provides us with two distinct interpretations, leading to two seemingly distinct kinds of internalism: 'accessibility-internalism' (or 'accessibilism') and 'mentalism'.

Historically speaking, the most widely accepted internalist requirement on justification has been that we should have a special kind of *cognitive access* to our justifiers. While not especially prevalent in contemporary epistemological literature, versions of this form of internalism have been endorsed by such eminent epistemologists as Bonjour (1980),

⁹ See, for example, Adam Leite's internalist account of the basing relation (2008).

Chisholm (1977) and Ginet (1975), among others. Because this type of internalism is defined in terms of the accessibility required to one's justifiers, it is commonly called 'accessibility-internalism'.¹⁰

What sort of access is required by the accessibility-internalist for a belief to be justified? Different writers give different formulations:

Every one of every set of facts about S's position that minimally suffices to make S, at a give time, justified in being confident that *p* must be *directly recognizable* to S at that time. (Ginet, 1975, p. 34, italics in original)

Internalism ... treats justifiedness as a purely internal matter: if *p* is justified for S, then S must be aware (or at least immediately capable of being aware) of what makes it justified and why. (Bach, 1985, p. 250)

The things we know are justified for us in the following sense: *we* can know what it is, on any occasion, that constitutes our grounds, or reason, or evidence, for thinking that we know. (Chisholm, 1977, p. 17, italics in original)

The general idea behind these different formulations seems to be that justification requires some special sort of *direct* or *immediate* access to the justifiers of one's belief.¹¹ Let us call this special sort of access 'reflective access'. We then define accessibility-internalism in the following way:

¹⁰ The term 'accessibility-internalism' is from Conee and Feldman (2001). Pryor calls this form of internalism 'simple internalism' (Pryor, 2001).

¹¹ Note here that it is possible in principle to combine such an access requirement with external conditions on justification. For example, William Alston, an externalist about epistemic justification, nevertheless writes: "[A] justifying ground is the sort of thing that, in general, and when nothing interferes is available for citation by the subject" which leads him to conclude, "the concept [of justification] should include

ACCESSIBILITY-INTERNALISM: S is justified in believing p only if p 's justifiers are reflectively accessible to S.

When do you have reflective access to a justifier for your belief? Reflective access is the sort of access that you have by engaging in introspection. Accordingly, you have reflective access to a justifier for your belief that p if and only if you can become aware of this justifier by introspection alone. Reflective access is opposed here to empirical access. Some facts we have access to only through our senses. Such facts are said to be empirically accessible. Other facts we have access to without relying on our senses, and these kinds of facts are said to be reflectively accessible.

What can we say about the kinds of facts reflectively accessible? Traditionally, it has been thought that by engaging in reflection we can become aware of our *mental states*. But which mental states are reflectively accessible? Here it proves useful to distinguish occurrent from dispositional mental states. First, occurrent mental states are states one is currently in. Thus, if I am currently experiencing pain, then my pain-state is an occurrent mental state. Similarly, we can say that my belief that p is occurrent if and only if I am currently aware of my propositional attitude that p is true. Dispositional mental states, on the other hand, are—as the name suggests—*dispositions* to form occurrent mental states. As dispositions, these mental states require conditions under which they are actualized (become occurrent). Thus, I have a disposition to feel pain just in case there are certain circumstances under which I normally will be in an occurrent pain-state. Similarly, I can have a dispositional belief that p if there are certain circumstances

the requirement that the justifier be accessible to the subject" (Alston, 1988b, p. 274-275). Alston's view is only weakly internalist, because he also requires that the grounds on the basis of which a subject believes be *adequate*, which is an external fact.

under which I will normally form the occurrent belief that *p*.¹²

The distinction between occurrent and dispositional mental states is relevant because it helps us to see which mental states are reflectively accessible. As we said, something is reflectively accessible if one can become aware of it by reflection alone. It is clear that occurrent mental states are reflectively accessible: it seems that just by turning my attention inward, I can become aware of which mental states I am currently in, at least on a sufficiently course-grained level. We also saw that dispositional mental states are states that under certain circumstances will become occurrent. Therefore, it is possible to have reflective access to our dispositional mental states as well: for every dispositional mental state, there will be circumstances in which I have reflective access to it. Something is accessible to one if it is possible for one to have access to it. Both our occurrent and our dispositional mental states are thus reflectively accessible.

Note the distinction between access and accessibility here. Of course, in circumstances where my dispositional mental states do not manifest themselves, I lack current access to them. But they might still be accessible if there are conditions under which I *do* have access to them. If accessibility-internalism were to require access to one's justifiers, only one's occurrent mental states would qualify. But that would be much too strong a requirement on justification. Take my (occurrent) belief that the Pythagorean theorem is true. I take it that this belief is justified. Nevertheless, as it happens, I do not occurrently believe anything that supports it. The point generalizes. I often form beliefs for which I do not pause to actively consider their basis, among these many of my perceptual beliefs. Intu-

¹² For further discussion of the distinction between occurrent and dispositional mental states, see (Armstrong, 1973; Audi, 1994; Price, 1969; Schwitzgebel, 2015; Searle, 1992).

itively, these beliefs are justified, but if justifiers were restricted to occurrent mental states, we would have to say that they are not. As formulated, accessibility-internalism evades this problem: even if I am not currently aware of the bases for my belief in the Pythagorean theorem, there may still be circumstances in which I am, and this will allow my belief to be justified.

Not only occurrent and dispositional mental states are said to be reflectively accessible, however. We also can come to know certain necessary truths through introspection.¹³ These necessary truths are called *a priori* for the very fact that we can come to know them independently of any contact with the external world, through the method of introspection. This set of *a priori* truths is said to include facts about arithmetic, logic and certain philosophical truths (Swinburne, 2001, p. 10).

It is sometimes said that reflective access is infallible, that one cannot be mistaken about one's own mental states, or about *a priori* truths. On the present account of justification, that would mean that we could not be mistaken about the justifiers in our possession for a particular belief. That is, for each and every one of our justified beliefs we could be certain which justifiers we have for them. This is the way accessibility-internalism is classically understood: the descriptions of internalism by Bach, Chisholm and Ginet above all support such a strong reading. This interpretation of accessibility-internalism is problematic, however, since infallibilism about reflective access is rather implausible. First, concerning *a priori* truths, such as truths of arithmetic or philosophy, it seems fairly obvious that we can be mistaken about them, especially when the truths in ques-

¹³ Although Kripke famously argued that not *all* necessary truths are discoverable through introspection (Kripke, 1980). Examples that are not include truths like "Water is H₂O" and "Obama is a person" (where Obama is picked out by pointing). We know these truths through some means other than introspection.

tion are rather complex. Second, as indicated by A.J. Ayer's famous 'problem of the speckled hen', we can also be mistaken about which mental state we are in (Ayer, 1940). Suppose we take a look at a speckled hen. One's experience of the hen features many speckles. But how many? One can consider this question and if pressed, produce an answer, such as "48". But it seems clear that one can be wrong about the number. Importantly, we are asking for a judgement about one's conscious experience of the hen—hence about one's mental state—and not about the hen itself. The problem is thus not just that one may be wrong about how many speckles the actual hen has, but rather that one may be wrong about how many speckles one's *visual experience* of the hen has. If it is assumed that we have reflective access to our occurrent mental states, and if this access is assumed to be factive, our judgement should be flawless. Yet we may be mistaken, it seems.

The implausibility of infallibilism does not impugn accessibility-internalism as we have defined it above, however, for two reasons. First, it is dubious that cases like the one above are cases of *justified* belief. That means we have to conceive of accessibilism in a weaker way than epistemologists like Bach, Ginet and Chisholm have understood that criterion. It is not the case that we can always know, on reflection alone, what the justifiers for one's belief are. But we may still require this for *justified* beliefs. Nothing in the claim that justification supervenes on what is reflectively accessible entails that reflection must *always* generate justified beliefs.

Second, our definition does *not* require that one's justifiers are reflectively accessible *as justifiers*. That would be to over-intellectualize epistemic justification. Not many people outside of philosophy have a concept of epistemic justification in the first place, let alone the concept of what it is to be a justifier. Presumably they do have justified beliefs. Thus any plausible reading of accessibility-internalism

should only require that one can become aware, by reflection alone, of the facts that justify one's belief. One does not need to become aware of the fact *that* these facts are justifiers.¹⁴

What are the motivations for accessibility-internalism? First, and foremost, the position is often defended by an appeal to a specific account of the nature of epistemic justification itself. The concept of justification that defenders of accessibility-internalism have traditionally appealed to is known as the *deontological* concept of epistemic justification.¹⁵ The fundamental idea behind the deontological concept is that epistemic justification shares important properties with other forms of justification, such as moral justification. Proponents of this view maintain that, like moral justification, epistemic justification can be analysed in terms of duties, although, in this case the relevant duties will be intellectual rather than moral.

Moral justification is related to the concepts of requirement and permission, obligation and duty, praise and blame, and, importantly, to specific moral *norms*. How are these concepts related? First, an action is morally justified just in case it is permitted by the moral norms. For example, you are morally justified in giving up your job to help the poor if you are permitted to do so by the relevant norms, even if these norms do not require this of you. However, you cannot be morally justified in giving up your job if doing so is *not* permitted by the relevant moral norms. Second, you are morally obliged to ϕ if and only if $\neg\phi$ is prohibited. Moral duties are obligations: it is your moral duty to ϕ if and only if it is your moral obligation to ϕ . We morally praise and blame people, finally, insofar as they are morally justified.

¹⁴ For a similar claim, see (Alston, 1988b).

¹⁵ The following discussion owes much to William Alston (cf. Alston, 1985, 1988a).

The deontological concept of epistemic justification embeds epistemic justification in an analogous structure. Thus, on this conception, the belief that p (Bp) is epistemically justified if and only if Bp is permitted by the epistemic norms. Thus, we arrive at the following definition for deontological justification:

DEONTOLOGICAL JUSTIFICATION: S is justified in believing p if and only if S does not violate any epistemic norms in believing p .

The deontological conception links epistemic justification to permission by epistemic norms. What is the content of these epistemic norms? According to BonJour, "One's primary epistemic duty ... includes both seeking the truth and avoiding error" (BonJour, 2002, p. 236).¹⁶ Since the goals of seeking truth and avoiding falsehood will often pull in opposite directions (the former is promoted by low evidential standards, the latter by high standards, for example), the two goals are to be weighed against each other. How exactly to weigh them is presumably sensitive to the context and contents of the beliefs in question. If falsity entails catastrophe, it is probably better to maintain high evidential standards. If nothing much hangs on it, we can be more gullible. However these norms are to be weighed against each other, it seems clear that epistemically speaking, *ceteris paribus*, we should believe when things are true, and disbelieve when things are false. We are epistemically praiseworthy, then, if we believe in such a way that our beliefs tend to be true, and blameworthy when we believe in such a way that they are not.

¹⁶ See also (Alston, 1989, p. 116). Specifications of the norm outside a specifically deontological context can be found in (David, 2001), (BonJour, 1985, p.7-8).

According to many, accessibility-internalism follows from the deontological concept of justification.¹⁷ The deontological argument for accessibility-internalism can be captured in the following way:

- P₁ Justification is solely a matter of meeting the applicable epistemic obligations.
- P₂ Obligations apply to one only if one has control over whether one complies with them or not.
- P₃ If meeting some epistemic obligation involves what is reflectively inaccessible, we lack control over whether we meet the norm.
- C₁ The epistemic obligations that apply to us do not involve what is reflectively inaccessible.
- C₂ Justification supervenes on what is reflectively accessible.

Here, P₁ states the deontological concept of justification. P₂ tells us that some conscious control over whether one satisfies a norm is required for that norm to apply. This is a special version of the famous 'ought-implies-can'-principle: not only do norms imply that one *can* comply with them, they also imply that one can *choose* to comply with them. For it makes no sense to say that someone should ϕ if she has no way of *choosing* to ϕ . That we tend to think of norms this way is borne out by the previously mentioned connection between norms, or obligations, and blameworthiness. Violating a norm or an obligation entails being blameworthy, but we do not consider someone blameworthy who did not choose to violate the norm. Culpable, maybe, but not blameworthy. Thus I am not blameworthy for driving

¹⁷ See, for example, (Alston, 1989; Bonjour, 1980; Goldman, 1999; Plantinga, 1993b).

through a red light if I did not decide to do so (perhaps I was distracted or incapacitated, or simply colour-blind).¹⁸

P₃ then says that if the satisfaction of the epistemic norms depends on factors that are reflectively inaccessible it will be *impossible* for us to consciously choose whether or not we satisfy them. The idea here is that a conscious choice to satisfy the norm requires that one can assess whether the conditions for the satisfaction of the norm are met. If these conditions are reflectively inaccessible, we cannot assess whether they obtain or not. Together, these premisses entail C₁ and C₂, and so they support accessibility-internalism.

Externalists will, of course, try to deny either one or several of the premises of the above argument. Here my aim is to provide a brief overview of some of the most important motivations for internalism. Whether the argument above succeeds or not, it is generally considered to be one of these.¹⁹

A second consideration in favour of accessibility-internalism is provided by Descartes' intuition that the main aim of epistemology is to give us epistemic *guidance*.²⁰ More recently a view like this has been endorsed by Laurence Bonjour, who writes:

[T]he central rationale for internalism, at least as I conceive it, ... arises when I ask simply whether or not I have good reasons for thinking that my various beliefs are true, understanding this question in a global way in which all of my beliefs (and ways of arriving at beliefs) are in question. (Bonjour, 2002, p. 237).

18 Of course, being distracted or incapacitated can *itself* be something for which I am blameworthy. But that is a different issue.

19 See, for example (Alston, 1989, pp. 198–200). The specific formulation of the argument above is my own.

20 See his "Rules for the Direction of the Mind" (Descartes, 1931).

Thus, for Bonjour an epistemic theory is adequate just in case it helps us to answer the question whether we have any good reasons for our beliefs, understanding this question in a broad sense. This criterion entails an accessibility constraint on the justifiers of one's beliefs, because justifiers that are reflectively inaccessible will be irrelevant for answering this question: "[o]ne immediate upshot of this is that the "internal" of "internalism" means primarily that what is appealed to for justification must be *internal to the individual's first-person cognitive perspective*" (Bonjour, 2002, p. 238, *italics in original*). Further, for Bonjour this conception of epistemology provides a rationale specific to accessibility-internalism rather than other forms of internalism, such as mentalism (to be discussed below) because properties of mental states that are reflectively inaccessible will likewise be irrelevant to epistemology.

Finally, internalism can be defended with the help of some 'intuition-pumps': thought-experiments or imaginary cases that serve to bring out our intuitions regarding a given concept.²¹ Two famous such cases should be mentioned here. First there is the 'new evil demon problem'.²² To see the problem, we first start by considering the possibility that everything that one perceives or thinks is actually produced by an evil demon manipulating the senses and one's cognitive faculties, producing exactly the same non-factive mental states as that we are currently in. As a result of the demon's tinkering, most of these are false. We then compare this possible world to our (supposedly) actual world in which everything is as it seems. According to the intuition the argument draws on, we feel that the demon's victims would be *as* justified as we are in the actual world. This would mean that justification does not supervene on anything other than our non-factive mental states,

21 See (Dennett, 2013, 1980) for more on the role of thought-experiments as intuition-pumps.

22 This famous case originates in (Lehrer & Cohen, 1983).

since, as the case is set up, everything except these mental states is different between these two worlds.

It is important to note that the intuition that the evil demon's victims are as justified as we are can—and has been—explained by an appeal to the deontological concept of epistemic justification:

Justification is a normative concept. It is an evaluation of how well one has pursued one's epistemic goals. Consequently, if we have reason to believe that perception, for example, is a reliable process, then the mere fact that it turns out not to be reliable, because of some improbable contingency, does not obliterate our justification for perceptual belief. (Lehrer & Cohen, 1983, p. 193)

Thus, the extent to which one is moved by the new evil demon case will presumably depend on the extent to which one accepts the deontological concept of epistemic justification. This is another way in which the deontological concept of epistemic justification can be used to motivate internalism.²³

Our second case concerns Bonjour's reliable clairvoyant Norman:

[Norman,] under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for

²³ For an interesting recent externalist response, see (Littlejohn, n.d.).

or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable. (BonJour, 1980, p. 62)

BonJour stresses that intuitively, Norman's beliefs are not justified, because Norman "has no reason at all for thinking that the belief is true. And the suggestion here is that the rationality or justifiability of Norman's belief should be judged from Norman's own perspective, rather than from one that is unavailable to him" (BonJour, 1980, p. 64). On the assumption that Norman's 'perspective' is constituted by what is reflectively accessible to him, the case thus seems to support accessibility-internalism.

This concludes our overview of accessibility-internalism. Much more could be said about the positions and arguments presented in this section, and some more will be said in later chapters. For now this overview will suffice. Let us continue, therefore, to the other prominent form of internalism in contemporary epistemological literature.

2.3 MENTALISM

Like accessibility-internalism, mentalism holds that justification supervenes on the internal. But mentalism has a different criterion for the internal than accessibility-internalism. What is specifically 'internal' about accessibility-internalism is the *relation* that is required to one's justifiers, whereas mentalism manifests its internalist character in its conception of what kind of things can be justifiers. It is for this reason that mentalism is sometimes called 'grounds-internalism'. In particular, mentalists hold that only mental states can be justifiers.

MENTALISM: The justificatory status of a person's doxastic attitudes strongly supervenes on

the persons occurrent and dispositional mental states, conditions and events. (Conee & Feldman, 2001, p. 2)

While the accessibilist form of internalism can perhaps be traced back to Descartes' insistence on indubitable beliefs, mentalism is a relatively recent development.²⁴

Mentalism and accessibility-internalism are usually taken to constitute substantially different theories of justification (Brown, 2007; Conee & Feldman, 2001; Pappas, 2014). It is usually argued that mentalism and accessibility-internalism are logically distinct: not all of our mental states are reflectively accessible (think of repressed memories, implicit biases, sub-personal mental states etc.), and conversely, it is at least logically possible that we could have reflective access to things other than mental states (Brown, 2007; Pappas, 2014).²⁵ Because of this, it should be possible to be justified according to accessibility-internalism, but not according to mentalism, and *vice-versa*.

While the theses are logically distinct, there are some reasons for doubting they have different extensions. First, mentalism, as defended by Conee and Feldman, consists of the claim that justification strongly supervenes only on one's *occurrent and dispositional* mental states. Whereas the

²⁴ Mentalism is most famously defended by Conee and Feldman (2001). Cf. (Wedgwood, 2002). The claim that accessibility-internalism can be traced back to Descartes is found in many places, among which the recent discussion between Bonjour and Sosa regarding internalism and externalism (Bonjour & Sosa, 2003, p. 7-8).

²⁵ Richard Fumerton thinks we can have "facts about which propositions make probable others, before our consciousness" (Fumerton, 1988, p. 163). Also, epistemological disjunctivists hold that we can have reflective access to factive reasons, like *seeing that p* (cf. McDowell, 1994; Pritchard, 2012b). While these factive reasons are mental states, usually mentalism is understood as the thesis that justification supervenes upon our *non-factive* mental states (cf. Goldman & Olsson, 2009, p. 312). We will come back to epistemological disjunctivism in Chapter 7.

above argument for the difference between accessibility-internalism and mentalism is based on a conception of mentalism that allows *all* of the subject's mental states to count as justifiers, this does not seem to be true to Conee and Feldman's conception. Since we have already argued that occurrent and dispositional mental states are reflectively accessible, it seems that on this formulation mentalism *does* imply accessibility-internalism.

This claim can be strengthened by considering the way mentalism is defended. I will focus on the most famous defence of mentalism, by Conee and Feldman (2001).

Conee and Feldman's case for mentalism has two parts. The first part consists of a number of contrast cases that are meant to show that justificatory differences imply mental differences. This argument is *abductive* because the general conclusion is meant to be the best explanation of the cases. They supplement it in the second part of their defence with responses to some possible objections to their view. Here I will focus on the first part of their defence.

As said, Conee and Feldman conclude from the exposition of their cases that the best explanation for the differences in justification between these cases is a difference in total mental state between the believing subjects.²⁶ I will argue, however, that the relevant mental differences all involve differences in *reflectively accessible* mental states. Because of this, the examples support the claim that justification supervenes on the *reflectively accessible* just as much as they support the mentalist claim. Here are the cases:

CASE 1: Bob and Ray are sitting in an air-conditioned hotel lobby reading yesterday's newspaper. Each has read that it will be very warm today and, on that basis, each believes that it is very warm today. Then Bob goes out-

²⁶ While I do believe all cases support the claims I want to make, I will focus on some of them to keep the discussion manageable.

side and feels the heat. They both continue to believe that it is very warm today. But at this point Bob's belief is better justified. (Conee & Feldman, 2001, p. 3).

Here the experience of the heat that contributes to Bob's justification is not only reflectively accessible, but reflectively *accessed*, since Bob's experience of the heat is an occurrent mental state of him.

CASE 2: A logic Teaching Assistant and a beginning logic student are looking over a homework assignment. One question displays a sentence that they both know to express a truth and asks whether certain other sentences are true as well. The TA can easily tell through simple reflection that some of the other sentences express logical consequences of the original sentence and thus she is justified in believing that they are true as well. The student is clueless. (Conee & Feldman, 2001, p. 4).

In this example it is explicitly stated that the TA has better justification because she can tell by simple reflection which sentences follow from which. Her better justification is thus due to a difference in reflectively accessible information.

CASE 3: Initially, Smith has excellent reasons to believe that Jones, who works in his office, owns a Ford. Smith deduces that someone in the office owns a Ford. The latter belief is true, but the former is false. Smith's reasons derive from Jones pretending to own a Ford. Someone else in the office, unknown to Smith, does own a Ford. The fact that Jones is merely simulating Ford ownership keeps Smith from knowing

that someone in his office is a Ford owner, but it does not prevent Smith from being justified or diminish his justification. At a later time Smith gains ample reason to believe that Jones is pretending. At that point Smith is not justified in believing either that Jones owns a Ford or that someone in his office owns a Ford. (Conee & Feldman, 2001, p. 4).

What explains the difference in justification is that Smith gains reasons to believe Jones is pretending. If this reason were not reflectively accessible to Smith however, it seems hard to see how his belief change would be a reasonable one. Again, the difference in justification is due to a change in mental state that is reflectively accessible.

In each case, the relevant mental states that made the difference for justification are reflectively accessible. Thus, these cases do not favour the claim that justification supervenes on one's total mental states *over* the claim that justification supervenes only on one's reflectively accessible mental states. If, as Conee and Feldman presuppose, we can extrapolate from these cases, then we could also conclude that only reflectively accessible mental states are relevant for epistemic justification. Thus, mentalism would seem to imply accessibilism.

What about the implication in the other direction? As we said above, we have reflective access to our occurrent and dispositional mental states, as well as to *a priori* truths. Since *a priori* truths are not themselves mental states, it may seem that their inclusion in the supervenience base of accessibilist justification makes that theory substantially different from mentalism.

The above reasoning does not go through, however. Why so? Let us first consider the nature of supervenience claims. A claim that A supervenes on B is a claim that there can

be no A-difference without a B-difference, or, equivalently, that every A-difference implies a B-difference.

Both mentalism and accessibilism are supervenience claims.²⁷ Thus, mentalism is the claim that there can be no justificatory difference without a mental difference, and accessibilism is the claim that there can be no justificatory difference without a difference in what is reflectively accessible.

Whether these claims amount to the same thing depends on what we understand the relevant mental states to be, as well as on what we consider to be reflectively accessible. Concerning the first, we have seen above that mentalism, at least as it is formulated by Conee and Feldman, in terms of occurrent and dispositional mental states, and as it is supported by their examples, seems to be committed to the claim that the mental states on which justification supervenes are all reflectively accessible. Regarding the second, we have seen above that there are two primary factors said to be reflectively accessible: mental states, and *a priori* truths.

Understood in this way, mentalism amounts to the claim that there can be no justificatory difference without a difference in one's reflectively accessible mental states, whereas accessibilism amounts to the claim that there can be no such difference without a difference in *either* one's reflectively accessible mental states *or* in the *a priori* truths reflectively accessible to one.

Crucially, on the assumption that all *a priori* truths are necessary and reflectively accessible, there can be no difference between two subjects as regards the *a priori* truths that are reflectively accessible to them.²⁸ If this is so, then every difference in justification, by the accessibilist lights,

²⁷ That is, considering these theories in their purest form, bracketing the concepts of justification used to motivate them—see below.

²⁸ Remember that *a priori* truths are *defined* as those necessary truths that are reflectively accessible.

must be accompanied by a difference in the reflectively accessible mental states of the subject, which is equivalent to the mentalist claim, as we have interpreted it. Understood as supervenience claims, and under the present assumptions, accessibilism thus implies mentalism.

This receives further support from classical accessibilist passages like the following:

It is not the fact that *there is* smoke rising from the forest that justifies S in thinking there is smoke rising from the forest, but rather such facts as that *S is confident that he sees* smoke, *S has no reason to mistrust* his sight on this particular matter at this particular time, and *S seems to remember* that he has come to know that virtually always when there is smoke of the sort he sees there is fire. (Ginet, 1975, p. 34, *italics in original*)

Here Ginet identifies as the reflectively accessible facts on which justification supervenes precisely those mental states that Conee and Feldman identify as relevant for justification. Even George Pappas, who otherwise stresses the difference between mentalism and accessibility-internalism, says that “we may safely say that if accessibility internalism is true, this is evidence in favor of mentalism” (Pappas, 2014).

Of course, Conee and Feldman distance themselves explicitly from any assumption to the effect that “the special kind of access on which many internalist theories rely can reach only mental items, and perhaps all mental items, or at least all that might be counted as playing a role in justification” (Conee & Feldman, 2001, p. 2). I think Conee and Feldman are right to stress that to identify mentalism with accessibilism requires one to make this assumption. However, the considerations above support the plausibility of

this assumption and, as we have seen, the assumption can in fact be motivated by considering Conee and Feldman's own examples. This gives us reason to think the distinction between accessibility-internalism and mentalism is not as clear-cut as it is sometimes presented.

Nevertheless we will treat mentalism and accessibility-internalism separately. The reason for this is twofold. First, one need not make the assumptions above about the supervenience bases of mentalism or accessibilism. One could adopt a mentalism, for example, for which justification does not supervene on the reflectively accessible mental states only, or an accessibilism that maintains that things other than mental states and *a priori* truths are reflectively accessible. Second, and perhaps more interestingly, mentalists tend to supplement their supervenience claim with a conception of epistemic justification that is quite different from the deontological conception that motivates accessibility-internalism. For Conee and Feldman support a theory of justification that they call *evidentialism* (1985). Evidentialism essentially consists of the following thesis:

EVIDENTIALISM: Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* if and only if having *D* toward *p* fits the evidence *S* has at *t*. (Feldman & Conee, 1985, p. 15).

What is one's evidence? While the notion of evidence in everyday life seems to primarily refer to items like blood-stained knives, witness testimonies or database entries, Conee and Feldman's view is that one's evidence consists of one's mental states. On this interpretation, evidentialism clearly supports mentalism.²⁹

There are further questions to be asked. For example, what does it mean for a doxastic attitude *D* to 'fit' a given

²⁹ We may ask whether one's evidence consists of all, or only of some of one's mental states. See Section 5.2 for discussion.

body of evidence e ? Here, I will not attempt to provide a complete overview of evidentialism.³⁰

What is crucial, however, is that, as we already mentioned in Section 2.1, Conee and Feldman's conception of epistemic justification is a conception of *propositional* justification, rather than of doxastic justification. As we said, however, we are discussing the internalism/externalism-debate concerning *doxastic* justification. *This* kind of justification Conee and Feldman label 'well-foundedness':

WELL-FOUNDEDNESS: S 's doxastic attitude D at t toward proposition p is well-founded if and only if

1. having D toward p is justified for S at t ;
and
2. S has D toward p on the basis of some body of evidence e , such that
 - a) S has e as evidence at t ;
 - b) having D toward p fits e ; and
 - c) there is no more inclusive body of evidence e' had by S at t such that having D toward p does not fit e' . (Feldman & Conee, 1985, p. 24)

Well-foundedness differs from evidential justification in that it not only requires that the relevant belief satisfies the evidentialist criteria, but that it *fits* the evidence on the basis of which it is formed, and there is no more inclusive body of evidence available to the subject that the belief does not fit.³¹ As is apparent from its definition, well-foundedness is a notion that applies to actually held dox-

³⁰ For such overviews, see (Conee & Feldman, 2004; McCain, 2014).

³¹ This latter requirement can be seen as a no-defeater condition, since a belief will not fit a body of evidence that includes defeaters for it.

astic attitudes, rather than merely possibly believed propositions.

Whether well-foundedness supports mentalism understood now as a claim about *doxastic* justification depends on how the notion of 'basing' is unpacked. If one specifies this in such a way that it refers exclusively to mental states, well-foundedness will supervene on the mental. If one opts for a causal conception of basing, then it will not. Consequently, this view may or may not qualify as a strong form of internalism regarding doxastic justification. What is important now is that we have a sketch of the mentalist picture of epistemic, that is doxastic, justification, and how it is supported.

We have seen how mentalism can be supported by an evidentialist concept of justification. A further motivation for mentalism draws on the assumption that justification is closely related to the notion of *rationality*.³² For example, we would not normally consider wildly irrational persons to be justified in their beliefs. Similarly, we consider most of our rationally held beliefs to be justified. Further, it is natural to spell out the notion of rationality in terms of coherence or consistency amongst one's beliefs. But consistency and coherence supervene on one's mental states, so the assumption that one's justified beliefs are one's rationally held beliefs supports the mentalist claim that justification supervenes on one's mental states.

Finally, mentalism can be supported by the 'new evil demon problem' discussed above. For if the beliefs of victims of the new evil demon are as justified as our beliefs are, and under the assumption that the deceived and non-deceived counterparts are mental duplicates, it follows that justification supervenes just on one's mental states. Impor-

³² For a defence of internalism (of the mentalist variety) along these lines, see (Wedgwood, 2002). Wedgwood's defence can be seen as a recent development of the Platonic idea that knowledge requires rational explanation (Section 2.1).

tantly, however, this argument only goes through under certain assumptions, which we will discuss below.

With this brief sketch of mentalism and its main motivations in mind, we will continue to the most prominent externalist theories of justification from the last decades.

2.4 RELIABILISM

One of the most prominent externalists is Alvin Goldman, who argues for the following necessary and sufficient conditions on epistemic justification:

RELIABILISM: If [and only if] *S*'s believing *p* at *t* results from a reliable cognitive belief-forming process (or set of processes), then *S*'s belief in *p* at *t* is justified.³³ (Goldman, 1979, p. 96)

Reliability of a method of belief-formation is not a mental state. Nor does it supervene on the totality of one's mental states, as the new evil demon problem seems to imply. Thus, Goldman is not a mentalist. Furthermore, reliability is typically regarded as reflectively *inaccessible*, which means he is not an accessibility-internalist either.

What does it mean for a belief to be reliably produced? According to the standard story, the cognitive process that caused the belief must have a 'tendency' to produce true beliefs (Goldman, 1979, p. 96). This tendency can be understood as a high frequency (either actual or hypothetical)

³³ In his 1979 paper, Goldman formulates this base clause as a sufficient condition, even though it is clear from the rest of the paper that he intends it to be necessary as well. In any case, at other places he explicitly endorses reliability as a necessary and sufficient condition for justification: "*S*'s belief in *p* is justified IFF it is caused (or causally sustained) by a reliable cognitive process, or a history of reliable processes" (Goldman, 1994, p. 309). Also, later in the 1979 paper, Goldman modifies his formulation somewhat. These complications are left out of the present presentation for reasons of clarity.

of true beliefs in the actual world, or in nearby possible worlds as well. In an early paper, Goldman explicitly endorses the latter claim:

[A] cognitive mechanism or process is reliable if it not only produces true beliefs in actual situations, but would produce true beliefs, or at least inhibit false beliefs, in relevant counterfactual situations. (Goldman, 1976, p. 771)

The non-modal interpretation of reliability should be rejected for independent reasons. Consider the following example, adapted from Greco (1999):

Reasoning according to the Gambler's Fallacy, Rene believes that roulette numbers which have not come up for long strings are more likely to come up next. However, Rene has a demon helper, which has decided on a whim to arrange reality so as to make Rene's beliefs all come out true.

Greco's assumption is that we would not consider Rene's beliefs justified. Still, Rene's method of belief-formation is reliable in the actual world. However, since the demon helper has decided on a whim to help Rene, we may suppose that he will not be helping Rene in most of the nearby possible worlds. In that case, while Rene's beliefs have a strong tendency to be true in the actual world, they do not have this tendency in most nearby possible worlds. The modal interpretation can thus explain why we would not consider Rene's beliefs to be justified, whereas the non-modal interpretation cannot.

Now that we have seen how reliability is to be understood, we may ask what can be said in favour of reliabilism as a theory of epistemic justification. As before, a principal reason for adopting reliabilism is provided by a

specific concept of epistemic justification, a concept that is known as the 'truth-conducive' concept of epistemic justification.³⁴

According to the truth-conducive concept of justification, "[o]ne is justified in believing that *p* only if that belief that *p* was formed in such a way as to make it at least very likely that the belief is true, or, as is sometimes said, only if it was formed in a 'truth-conducive' way" (Alston, 1989, p. 144). Thus, on the truth-conducive concept of justification, beliefs are justified only insofar as they are likely to be true. The truth-conducive concept of justification supports reliabilism because beliefs that are produced by a reliable method will be likely to be true. If likely truth is thus the primary criterion for justification, a necessary and sufficient reliability condition on justified belief will make sure that this criterion is met for all and only justified beliefs.

Why adopt the truth-conducive concept of justification? Laurence Bonjour (himself no friend of externalist theories of justification!) describes the motivation for a truth-conducive concept well:

[T]he goal of our cognitive endeavours is *truth* ... If truth were somehow immediately and unproblematically accessible ... then the concept of justification would be of little significance and would play no independent role in cognition. But this epistemically ideal situation is quite obviously not the one in which we find ourselves. We have no such immediate and unproblematic access to truth, and it is for this reason that justification comes into the picture. The basic role of justification is that of a *means* to truth ... If epistemic justification were not conducive to truth in this way, if finding epistemically justified beliefs did not substantially increase the

³⁴ Goldman himself adopts this concept of justification (2009).

likelihood of finding true ones, then epistemic justification would be irrelevant to our main cognitive goal and of dubious worth. (BonJour, 1985, pp. 7-8)

BonJour explains how the requirement of truth-conducivity falls out of the function of justification with respect to our ultimate cognitive aim. We want our beliefs to be justified, according to BonJour, because we want them to be true. Thus justification should be truth-conducive.

While the truth-conducive concept of justification supports reliabilism, internalists like BonJour have tried to accommodate the intuitions behind the truth-conducive concept of justification by relativising the truth-norm to the abilities of the believing subject. Above we saw that the deontological concept of justification relates justification to blameworthiness (or lack thereof), whereas the truth-conducive concept does not. This makes for an important difference between the two concepts of epistemic justification because blameworthiness presupposes control, and we may sometimes lack control over whether we form our beliefs in a truth-conducive way. On the deontological accounts, such beliefs should be epistemically permitted and thus justified. On the truth-conducive account, however, no such control is presupposed, and consequently these beliefs will not be justified if they are not formed in a truth-conducive way.

The truth-conducive concept is also distinct from the evidentialist concept of justification. For as we have seen, the evidentialist takes justification to consist in one's beliefs fitting one's evidence. But there is no guarantee that this will result in beliefs that are likely to be true. For example, one's beliefs may fit a body of highly misleading evidence very well without one's belief-forming method thereby being truth-conducive. This means there will be possible cases of evidentialist justification without truth-conducive

justification. There will also be cases of truth-conducive justification without evidentialist justification. Clairvoyance cases like Norman's case described above may be examples of this. Norman's beliefs will presumably be justified in the truth-conducive sense because they result from reliable clairvoyant powers. However, Norman's case is set up such that Norman does not have any evidence that support his beliefs, and so his beliefs will not be justified in the evidentialist sense.

Besides an appeal to the truth-conducive concept of justification, the defender of reliabilism has some other arguments for his position at her disposal. Most of these arguments are negative in the sense that they try to establish the plausibility of externalism by demonstrating the *implausibility* of internalism. As such, they usually do not provide arguments specifically meant to establish reliabilism rather than some other externalist account of justification. They only support externalist accounts in general. Insofar as they support reliabilism, they do so indirectly, by eliminating some of its competitors. However, as we will see, some of the arguments to be found in the literature fail even as arguments for externalism generally, because they attack too narrow a conception of internalism. Let us start with the ones that do succeed.

First of these, externalist concepts of justification can be motivated by the desire to avoid Gettier cases.³⁵ This argument can be used as an argument for externalism generally, but we will here consider a version that is specifically congenial to the reliabilist. As we saw, Gettier's cases hinge on the presence of justified true belief that fails to constitute knowledge. Importantly, however, the justification of the beliefs of the victims of Gettier cases typically satisfies internalist criteria for justification. That is, Gettier's

³⁵ See (Conee, 2004, p. 79; Plantinga, 1993a, p. 32; Pritchard, 2012a, p. 35; Sennett, 1992, p. 644; Sturgeon, 1993, p. 158,159).

victims are justified in both the accessibilist and the mentalist sense. However, it is not clear that reliabilist justification is similarly present. Gettier's two original cases concerned beliefs that were formed on the basis of an inference from false beliefs, and while this will require the defender of reliabilism to provide a criterion for the individuation of belief-forming methods, it is at least possible to maintain that this is an example of an unreliable belief-forming method. If so, the beliefs in Gettier's own examples do not present an argument against the tripartite analysis of knowledge *provided* that the justification condition is interpreted in a reliabilist way. Thus, one can motivate reliabilism by arguing that it might save the tripartite account of knowledge.

However, Linda Zagzebski has argued that no non-factive externalist theories of justification, including the reliabilist one, can evade all possible Gettier cases (Zagzebski, 1994). If this is true, then Gettier cases will not provide reliabilism with a distinctive advantage over internalist concepts of justification after all. As we will see in Chapter 4, however, Zagzebski's claim is too strong. What is required for a theory of justification to be immune to Gettier cases is to be incompatible with a certain form of *luck*, a criterion that may be met by some accounts of fallible justification. In Chapter 6, we will then see that reliabilism is incompatible with the specific form of epistemic luck that is problematic in Gettier cases, and thus that, contrary to Zagzebski's claims, reliably justified belief will be immune to Gettier cases.

Second, externalist concepts of justification are sometimes motivated by an appeal to sceptical arguments.³⁶ The sceptic starts with the assumption that two people can be mentally identical (with identical possibilities for reflective access), even though they are in radically different contexts.

³⁶ See (Pritchard, 2005a; Sosa, 1999).

In one of these contexts, one's mental states accurately represent one's surroundings (call this the 'good' case). In another, one's mental states may be wildly off the mark, perhaps due to the deception of a malevolent demon (call this the 'bad' case). In the good case one may know many things, whereas in the bad case one may hardly know anything. The sceptic then adduces a closure principle that says that if we know p , and we know p implies q , then we can know q . The problem then becomes apparent: most of our knowledge implies that we are not in the bad case (since in that case we have no knowledge). Closure thus dictates that if we know the things we *think* we know, we can know that we are not in the bad case.

Internalist accounts of justification have a hard time explaining such knowledge. Such knowledge would require justification for the belief that we are in the good case rather than in the bad case. But such justification cannot be internal, for as it is generally understood nothing internal distinguishes the good case from the sceptical scenario. It seems thus, that one cannot be internalistically justified in believing that one is in the good case. If knowledge requires justification, this means one cannot know one is in the good case on an internalist concept of justification. At this point, the closure principle returns with a vengeance, since not knowing that I am in the good case implies that I don't know all propositions that entail that I am in the good case, such as that I have two hands. The result of the adoption an internalist conception of justification in combination with the closure principle is thus a thoroughgoing scepticism.

Externalist accounts of justification fare better in this regard, for externalist justification *does* discriminate between the good and the bad case. In the good case, for example, my beliefs are by and large reliably produced, whereas they are not in the bad case. Thus, it is possible for me to have a true justified belief that I am not in the bad

case, if justification is determined by reliability of belief-forming method. I can then have most of my other knowledge too. Note that externalism need not maintain that it is always possible to have externalistically justified *second-order* beliefs to claim an advantage over internalism. For even if second-order knowledge is impossible, externalism would still rescue the possibility of external first-order knowledge.³⁷ Although crude, this sketch of the sceptical argument does show at least that externalist accounts of justification have a *prima facie* advantage over internalist accounts in dealing with the sceptical problem.

Next, I will discuss two arguments from the literature that ultimately fail as general arguments for externalism, since they are directed at too narrow a conception of internalism. Even if they both fail as arguments for externalism, it is instructive to consider them anyway, since they bring out some hitherto undiscussed advantages of mentalism over accessibility-internalism. The first of these derives from semantic externalism. As Hilary Putnam famously argued, the meaning of our concepts and thus our beliefs “just ain’t in the head”, but instead depends on factors external to the believing subject. (Putnam, 1975, p. 227). Semantic externalism poses a problem for the accessibility-internalist, who maintains that justification supervenes on what is reflectively accessible. For whether one is justified depends on the logical relations between the contents of various of the subject’s beliefs, which in turn depends on their respective *meaning*. For example, whether I am justified in believing that it is currently raining outside on the basis of my belief that I see droplets of water running down the window will depend on whether the latter be-

³⁷ But in principle the reliabilist could maintain that second-order knowledge can be had. Even if we can never have reflectively accessible *reasons* for thinking we are not in the bad case, our beliefs that we are not, and that our first-order beliefs are reliably produced could still themselves be the result of reliable processes.

lief makes the former more probable. But whether this is so will of course depend on the *contents* of those beliefs, that is, their meaning. Semantic externalism says that this content is partially constituted by facts that we only have empirical access to. Famously, the contents of these beliefs will depend on whether water is H₂O instead of XYZ, a fact to which I lack reflective access. Thus, semantic externalism implies that justification does not supervene on the reflectively accessible alone.

The argument is ultimately unsuccessful as a general argument for externalism, however, because it does not rule out mentalism (Brown, 2007). For mentalism itself does not require that we have reflective access to our justifiers. Semantic externalism is a theory about the meaning of our mental states. Mentalism is neutral on the question what the meaning of a mental state is, or how we can access that meaning. The only thing that seems necessary for mental states to be able to justify, is that they *have a meaning*, something which semantic externalism does not dispute.³⁸ At first sight, it thus seems that semantic externalism cannot be used in a straightforward way to support externalist theories of justification.³⁹

A second argument frequently found in the literature that fails to rule out mentalism starts from the observation that we regularly attribute knowledge to young children and animals. For example, we say things like “Little Jane knew that today was going to be a special day for her”, and “My dog Fido knows it is time for a walk when I put on my

³⁸ Note that if semantic externalism is true, mentalism and accessibilism will have different extensions, contrary to our earlier suggestion. Here and later on, I will remain agnostic about the truth of semantic externalism.

³⁹ Unfortunately I do not have the space here to discuss the implications of semantic internalism for the internalism/externalism-debate in epistemology in full. For a more detailed discussion leading to the same conclusion, I refer the reader to an excellent discussion by Jessica Brown (2007).

old boots". On the assumption that knowledge requires justification, this implies that these subjects all have justified beliefs. It can be doubted, however, that young children and animals have the capacity for reflection necessary for justification according to accessibility-internalism. If some of these subjects lack this capacity, then they are incapable of accessibilist justification.

If we want to save the possibility of young children and animal knowledge, it thus seems we need to abandon accessibility-internalism. It is unclear, however, what the bearing of such knowledge is on mentalism. For young children and animals may well have occurrent and dispositional mental states that fit their experiential evidence. If so, they are justified according to the mentalist. Again, we find that this argument fails as a general argument for externalism because it does not affect internalism of the mentalist variety.

2.5 VIRTUE EPISTEMOLOGY

Our final account of justification is based on a virtue-theoretic framework. There are various virtue-theoretic accounts of justification, some internalist, others more externalist. Here we consider perhaps the best-known and most extensively developed externalist virtue-theoretic account of justification: the account championed by Ernest Sosa and John Greco, among others (e.g. Greco, 1993, 1999, 2000; Sosa, 1985, 1991, 1993, 2007, 2009c, 2010). The view is characterized by Sosa as follows:

[W]e can ... evaluate [a belief] as 'epistemically justified', in one or another sense: 'competently adroit' perhaps (or reliably based, or counterfactually safe, etc.), or perhaps 'rationally justified' (coherently fitting, and held in part on that basis). (Sosa, 2009d, p. 114)

The first thing to note is that Sosa distinguishes two senses of justification: competently adroit belief, and 'rational justification'. Here we will focus on the first of these, since this is Sosa's particularly externalist concept of justification. In later chapters we will come back to the distinction. While the second sense of justification that Sosa identifies is stronger than the first, adroitness is required for either sense of epistemic justification. A belief is adroit, on Sosa's account, if and only if it manifests an epistemic competence, in particular a competence to attain true belief (e.g. Sosa, 2015, ch. 1). Thus, a belief is justified for Sosa only if it manifests a competence to believe truly.

We will not discuss Sosa's complete account of competence. What is relevant is the following:

What then is required for possession of a competence? Required for archery competence, ... [t]here must be a close enough sphere of possible worlds where one takes shots, varied enough across the relevant range, and these shots must easily enough succeed, extensively enough across the relevant range. (Sosa, 2015, p. 97)

Thus, for Sosa, one possesses a competence for X only if one could not have easily failed to X if one tried. Competences are relative to a set of appropriate conditions, however. An archer is competent if her shot could not too easily fail to hit the target. But archery competence does not require that one's shot could not easily fail to hit the target when shooting in the midst of a storm. For Sosa, competences are thus relative to certain 'appropriate circumstances'.

Justification on Sosa's conception requires *manifesting* a competence to attain true belief. When is a competence manifested?

Consider the archer who shoots with the unexpected gust about to cross the field and the guardian angel poised to intervene, unbeknownst to the archer. This archer does not earn proper credit for his success, which does not really manifest competence. And the reason for this, I suggest, is that the archer does not shoot when in appropriate shape, in an appropriate situation. (Sosa, 2015, p. 103)

According to Sosa, then, a competence is manifested only if exercised in the appropriate situation. Justification, we saw, requires the manifestation of a competence to believe truly. According to the above, this requires in turn that one could not have easily believed falsely in the appropriate situation, and that further, the situation one is actually in is indeed appropriate. If one is justified in this sense, one believes thus in a way that could not have easily produced false belief.

This strand of virtue epistemology is sometimes called ‘agent reliabilism’ because it restricts the methods of belief formation relevant to epistemic justification to *competences*, reliable methods of belief formation that have their basis in the ‘cognitive character’ of the believer. Thus, not all reliable processes can confer justification, on the present view, but only those that are in some sense rooted in the character of the believing agent. It is clear that this account is externalist, because whether one’s belief manifests an epistemic competence depends on features of the environment, which are neither mental states nor reflectively accessible.

I shall call the concept of justification present in the virtue-theoretic theory of justification a *competence based* concept of justification. The epistemic competences possessed by the agent are interpreted by the proponents of this kind of account as intellectual *virtues*. For this reason, Sosa labels his account a virtue epistemology. These

virtues are not to be understood in a narrow Aristotelian sense, however, for that would require them to involve deliberate choices, which the usual competences identified by agent reliabilists like Sosa, such as our perceptual belief-forming processes, usually do not.⁴⁰ Rather, they are to be seen as virtues in a broader teleological sense: they facilitate the fulfilment of some proper ends. Having a sharp blade can be said to be a virtue of a knife, since it facilitates the fulfilment of on the knife's proper ends (cutting), and so can having a reliable visual system be said to be an intellectual virtue, if one considers it a proper end of the intellect to learn the *truth* about one's environment.⁴¹

Because we have already discussed the general arguments for externalism, we can keep our discussion of the motivations for the competence based concept of justification relatively brief. One of the principal reasons for adopting such a view of justification rather than other externalist accounts, is that justification seems to require a certain 'direction of fit' between our beliefs and the world. In particular, justification requires the direction of fit to be such that our beliefs 'fit' the world, rather than the other way around. Consider the following pair of cases:

CASE 1: John has normal eyesight and forms the belief that there is a fir tree in front of him. His eyesight is reliable and he thus acquires a true belief.

CASE 2: Modified from Rene's case, discussed above. John is playing roulette and reasons ac-

⁴⁰ *Internalist* virtue theoretic accounts of justification do conceive of the virtues in a more Aristotelian sense, see (e.g. Baehr, 2011; Zagzebski, 1996). Also, the fact that Sosa does not think that deliberate choice is *necessary* for epistemic virtue does not mean that it is incompatible with it. According to Sosa, *some* intellectual virtues do require deliberate choice on the part of the agent (Sosa, 2015, Ch. 2).

⁴¹ This does not imply that learning the truth about one's environment is the *only* proper end of the intellect.

according to the gambler's fallacy. He thus believes that numbers that have not come up for a long time will be more likely to come up next. Unbeknownst to John, a helping angel, feeling sorry for his poor reasoning skills, decides to influence the roulette table in such a way that John's beliefs about the upcoming numbers will all be true. The helping angel's disposition to help John is stable, and so the angel will continue to help John across all nearby possible worlds.

Case 1 is an uncontroversial case of perceptual justification. It is not so clear, however, whether John's beliefs in Case 2 are justified. The virtue theorist can explain our hesitance here by noting that John's beliefs do not manifest any epistemic competence of his. The success of his beliefs is not creditable to him, but rather to the helping angel. Because we can fill out the details of the two cases in such a way that John's beliefs in Case 2 may actually be *more* reliable than his beliefs in Case 1, a competence bases view of justification here has a distinctive advantage over a purely reliabilist one.

In a more general sense, a virtue theoretic account of knowledge and justification has the advantage of being able to accommodate the sense in which knowledge and justification are *achievements* of some sort. It is able to do so because it takes as its central focus the agent and her epistemic competences. Achievements can then be defined as competences that successfully manifest themselves. In doing so, virtue epistemology is able to provide a unified picture of various kinds of *value*. For on a virtue theoretic account, both moral and epistemic value are the result of competences that successfully manifest themselves. They are distinguished by the fact that the moral competences

have a different aim (the good) than epistemic competences (the truth), but their structural analysis is the same.

The above approaches to epistemic justification are, I believe, the most prominent approaches in contemporary literature. More can, and will be said about them, in later chapters.⁴² Here we set out to provide a bird's eye view of the internalism/externalism-debate by considering a representative sample of the main positions within the debate. For this I believe the brief discussion above suffices.

2.6 CONCLUDING REMARKS

In this chapter I provided an overview of the internalism-/externalism-debate about epistemic justification. Because it is impossible to provide a complete survey here, I focussed on four main positions within that debate; two internalist, two externalist. I provided an overview of their general commitments, and their main motivations. In particular, I argued that these different internalist and externalist theories of justification are motivated by four distinct concepts of justification. Whereas accessibilism is motivated by a deontological concept of justification, mentalism is defended by an appeal to an evidentialist concept of justification. On the externalist side, we have seen that reliabilism can be defended by an appeal to the truth-conducive concept of justification, whereas agent reliabilism, despite its similarity in name to reliabilism, takes a radically different perspective by conceiving of justification within a virtue-theoretic framework.

What this means for the internalism/externalism-debate is that it is not clear whether the disagreement between the various positions discussed in this chapter is substantial or merely verbal. After all, a substantial disagreement presupposes a common concept of justification *about which*

⁴² See especially Chapters 5, 6.

different parties disagree. Willam Alston argues that there is no such common concept, and that we should consequently refrain from talking about epistemic justification altogether (Alston, 2005). In Chapter 6, I argue that we can resist this pessimistic conclusion by looking at the relations between the various theories of justification and various kinds of *luck*. To this end, we will start in the next chapter by developing a general account of luck.

A MODAL ACCOUNT OF LUCK

In the previous chapter I provided an overview of the main rival theories of epistemic justification. In order to investigate the relation between these accounts of justification and luck, I need an account of luck. In this chapter I provide and defend such an account.

The philosophical literature on luck features roughly three distinct accounts: a modal account, a ‘lack-of-control’ account and a probabilistic account. In this chapter I develop and defend a modified version of Duncan Pritchard’s Modal Account of Luck (MAL).

The chapter has the following structure. In Section 3.1, I provide a brief overview of MAL. In Section 3.2 I further develop the account and present a modified version that I will call the Alternative Modal Account of Luck (AMAL). In Section 3.3, I discuss the two main rival accounts of luck, the ‘lack of control’ account, and the probabilistic account, and provide some arguments to prefer AMAL over its alternatives. Finally, in Section 3.4, I discuss objections that have been raised against modal accounts of luck in general, and defend AMAL against these objections. I will close with some concluding remarks in Section 3.5.

3.1 PRITCHARD’S MODAL ACCOUNT OF LUCK

Arguably, the most comprehensive and illuminating treatment of luck in modal terms is provided by Duncan Pritchard (Pritchard, 2005a, 2005b, 2007, 2008a, 2014). We will therefore take Pritchard’s modal account of luck (MAL) as a starting point for further development.

One caveat should be noted before we begin, namely that Pritchard's account has evolved over time.¹ I will argue, however, that there are some points where Pritchard was wrong to modify his earlier views. Therefore, the alternative modal account of luck that I will develop in the next section will be closer to Pritchard's original modal account of luck than to his most recent one. In this section I will provide an overview of the most recent account (2014).

What distinguishes lucky events from other events that are significant to an agent? One natural suggestion is that a lucky event is an event that could have *easily* failed to occur. Paradigm cases of luck satisfy this requirement. If I win the lottery, I could have easily lost. When I find a treasure by randomly digging in my yard, I could have easily failed to find it. Bad luck works in the same way. If it is a matter of bad luck that I contracted a rare disease, then I could have easily *failed* to contract it. The interpretation of these 'easy possibilities for failure' in modal terms—in particular, in terms of *possible worlds*—is the fundamental basis on which MAL rests. Thus, Pritchard provides the following characterization of luck:

“roughly, what makes an event lucky is that while it obtains in the actual world, there are—keeping the initial conditions for that event fixed—close possible worlds in which this event does not obtain”. (2014, p. 599)

How does this claim explicate the idea that lucky events could have easily failed to occur? First, to say that an event could have failed to occur means that there is a possibility for it not to occur. These possibilities are interpreted using a possible worlds framework (cf. Lewis, 1973; Kripke, 1980). In this framework, to say that an event could have

¹ Cf. (Pritchard, 2014, 2005b)

failed to occur is to say that there are possible worlds where the event does not obtain.

Second, some possibilities (possible worlds) could more easily be realised than others (Lewis, 1973; Sainsbury, 1997). To take a philosopher's favourite example, it is possible that I am currently a brain in a vat (BIV) controlled by evil scientists. But presumably, this is less *easily* possible than the possibility that I would have arrived a few minutes later at work today.

We can thus order possible worlds in terms of how easily possible they are from the perspective of the actual world. For Pritchard, this ordering is determined exclusively by the *similarity* of these worlds to the actual world.² Thus, the more similar a world is to our actual world, the more easily possible it will be, from our perspective. Pritchard calls the degree of similarity between two worlds the *modal distance* between these worlds. The more similar a world is to another, the 'closer' it is to it, modally speaking. On this interpretation, lucky events are events that fail to occur in nearby possible worlds. To take our previous example, it is (presumably) less easily possible that I am a BIV than that it than that I would have arrived a few minutes later at work today, because (presumably) the closest world where I am a BIV is less similar to the actual world than the closest world where I arrive a few minutes later. If the difference is large enough, then on Pritchard's view, it would not be a case of luck that I am not in a sceptical scenario, but it would be a matter of luck that I arrived at just that particular time today.³

² I will disagree with this claim in the next section.

³ If one finds the second part of this claim implausible, that is because the event in question is insignificant to me. Or so I will argue in the next section.




Figure 3.1: Winning the lottery

On this conception, lucky events can be graphically represented.⁴ For example, Figure 3.1 depicts a case of winning the lottery, a paradigm lucky event. Here, the curve W encloses an area that represents the set of possible worlds in which one wins the lottery. The bold circle encloses all possible worlds, the set of nearby possible worlds is enclosed by a sphere S , and α represents the actual world. We see that W encloses only a small subset (the grey area) of S , which represents the fact that there are many close possible worlds where one does not win the lottery, which, on the present account, means that the lottery win is a clear case of luck.

The formulation above restricts the nearby possible worlds relevant for luck to those where the ‘relevant initial conditions’ for the event are the same as in the actual world. This is necessary, because whether an event is lucky depends on what we take to be the initial conditions for that event. Suppose I win the lottery. If the initial conditions for this event are such that I randomly bought a

⁴ The representations in this chapter are of my own construction. They are inspired by those in (Lewis, 1973).

ticket, this is clearly a case of luck. But if they are such that I rigged the lottery in my favour, it may well not be.

Whether an event is lucky thus depends on what the initial conditions are for that event. This means that not all nearby possible worlds will be relevant for the luckiness of an event, because in some of those the initial conditions for the event may differ. That just means that sometimes the initial conditions for an event could have easily been different than they are. Consider again my winning the lottery, but now suppose that while I actually played fair, I *almost* cheated, such that in many nearby possible worlds I cheat and therefore win the lottery.⁵ This situation is depicted in Figure 3.2. Again, α represents the actual world, the bold circle the space of all possible worlds, S the sphere of nearby worlds, and W the set of worlds in which one wins the lottery. However, we now introduce a second set of possible worlds, C , to represent the set of possible worlds in which one cheats. Because one almost, but not actually cheated, C does not contain α , but it does contain most of the worlds in S .

The fact that if one had cheated, one could have easily won the lottery is represented by the fact that many of the worlds in C are also in W . Crucially, however, the C -worlds are irrelevant for the luckiness of my win in the actual world. For that, only the nearby worlds (S -worlds) where I do *not* cheat are relevant. Only a small fraction of these are worlds in which I win (grey area), and thus we get the required result that my win is a case of luck, even if I almost cheated.

⁵ If one has trouble imagining how this is possible, suppose that I walked passed the desk of some betting official who had left on his computer unlocked, enabling me to easily change the winning numbers of the lottery. Just at the moment when I was about to enter the numbers, I got an urgent call, preventing me to actually go through with it.




Figure 3.2: Winning the lottery while almost cheating

How are we to determine, in each case, what the relevant initial conditions are for an event? Pritchard says the following:

Is there a general specification that one can offer of these “initial conditions”? Well, we can say this much: they need to be specific enough to pick out a particular kind of event that we want to assess for luckiness, but not so specific as to guarantee that this event obtains ... That’s quite vague, of course, but my suspicion is that we shouldn’t expect anything more detailed, in that we shouldn’t require a theory to be any more precise than the phenomena about which we are theorizing. For our purposes it is enough that we can pick out such initial conditions on a case-by-case basis (which I believe we usually can). (Pritchard, 2014, p. 599)

This statement is indeed somewhat vague, and in the next section I will comment on it. Here my aim is merely to present Pritchard’s views.

It should be noted that luck is a notion that admits of degrees: some events are more lucky than others. For

Pritchard, the degree of luck varies exclusively with the nearness of the nearest world in which a given event *E* fails to occur. The closer the nearest world in which *E* fails to occur, and in which the relevant initial conditions are the same as in the actual world (from now on, I will leave this clause implicit if it is clear from the context that it should be included), the more lucky *E* is. Winning a lottery when the lottery is rigged in one's favour is less lucky than winning it if it is not so rigged. In the second case, one could have easily lost: the nearest worlds where one loses are worlds in which the lottery machines generate a few different random numbers. If the lottery is rigged in one's favour, more would need to be different for one to lose.

Is there a threshold value for luck? That is, can we neatly divide events into lucky ones and non-lucky ones, based on the specific modal distance between the actual world and the nearest world where the event fails to occur? That seems implausible because modal distance is interpreted here in terms of similarity between worlds, and the notion of similarity is notoriously vague. Thus, the modal distance between worlds will be vague too. This does not mean, of course, that we cannot compare modal distances at all. *Sometimes* it will be clear whether a world is similar enough to the actual world to be easily possible, or whether one world is more similar to the actual world than another. But sometimes this will not be so clear. Thus, there will be borderline cases: worlds of which it is indeterminate whether they are easily possible. Consequently, it will often be difficult to determine *how* easily possible some event is, from the perspective of the actual world. In these cases, it may be indeterminate whether the failure of the event in these worlds suffices for the event to be a case of luck.

Undoubtedly, some will think this vagueness is a weakness of Pritchard's modal account of luck. However, I be-

lieve it rather to be a strength. Lewis famously defended his analysis of counterfactuals in terms of comparative similarity between possible worlds by saying that given that the notion of counterfactual is *itself* vague, we cannot expect to provide a precise definition of it. Rather, we should aim "to rest an unfixed distinction upon a swaying foundation, claiming that the two sway together rather than independently"(Lewis, 1973, p. 92). In a similar vein proponents of the modal account of luck can be seen to rest the 'unfixed distinction' between luck and non-luck on the 'swaying foundation' provided by the notion of 'easy possibility'.

This approach requires one to maintain that luck is itself a vague notion. It is easy to provide examples that support such an interpretation. Suppose I lose my wallet, only to discover ten minutes later while retracing my steps that it fell out of my pocket on the street. It is still there. Is it a matter of luck for me that it is? Are the worlds in which I don't find my wallet similar enough to the actual world to count as 'nearby'? Hard to say. The case seems to be a borderline case of luck.⁶ Further, was it a matter of bad luck for the dinosaurs that an asteroid hit the earth? Was it a matter of luck for the allied forces that WWII ended the way it did? Am I lucky not to have fallen ill this year? There are many borderline cases of luck, it seems, and any adequate account of luck should respect this fact.⁷ We should therefore not expect the concept of luck to be susceptible to precise definition. Nor should we *want* to make the concept of luck precise any more than we should want to make the concepts of similarity, tallness, or heap precise. For some

⁶ I owe this example to (Pritchard, 2005a, p. 132, fn. 12).

⁷ Note that the indeterminacy does not depend on any lack of information on our part (cf. Williamson, 1994). Even if we were in possession of all relevant information, it would, I presume, still be indeterminate whether it is a matter of luck that I retrieved my wallet, or that the meteor hit the earth.

events it is simply indeterminate whether they are lucky or not.

Some say that vague notions are unfit for clarification in principle, but this seems implausible. 'Tallness' is a vague concept: there are borderline cases where it is indeterminate whether someone is tall or not. Nevertheless, one can gain understanding of what it is to be a giant by being told that giants are tall. In a similar vein, we can gain understanding of the phenomenon of luck by being told that lucky events could have easily failed to occur, even if we cannot always determine whether an event could have easily failed to occur.

Moreover, modal analyses in terms of similarity between worlds have been useful in explicating concepts pertaining to a large number of philosophical issues. Lewis and Stalnaker modally explicate counterfactuals (Lewis, 1973; Stalnaker, 1968). Peacocke puts modal concepts to use in the analysis of freedom of will (Peacocke, 1999). In epistemology, defenders of a safety condition on knowledge (including Pritchard) use the concept of easy possibility in various ways to clarify the concept of knowledge (cf. Pritchard, 2005a; Sosa, 1999; Williamson, 2000). The fruitfulness of these projects warrants at least the attempt to use these concepts in an analysis of luck.⁸

We have provided an overview of Pritchard's modal account of luck. While I agree for the most part with the analysis, I will in the next section consider some ways of improving it.

3.2 AN ALTERNATIVE MODAL ACCOUNT OF LUCK

As I said above, Pritchard changed his views about the best way to capture the phenomenon of luck. Perhaps most

⁸ This is especially plausible considering the strong relation between knowledge and the absence of luck.

strikingly, and the first thing I want to focus on here, is that he formerly considered there to be a necessary *significance* condition on luck (cf. Pritchard, 2005a, 2005b, 2007, 2008a). That is, where Pritchard used to think that lucky events must be significant to someone, he now thinks that events that lack any significance can nevertheless be lucky.

The first thing to note about Pritchard's new position is that it deviates considerably from the way we ordinarily speak about luck. We do not say, for example, that small avalanches at the South Pole are lucky, or that heaps of sand on Mars luckily form certain patterns, even if these events could have easily failed to occur. In such cases we use different concepts to describe the events, such as 'accidental', 'unlikely', or 'improbable'. It is no accident that The Oxford English Dictionary defines luck as "the fortuitous happening of an event *favourable or unfavourable* to the interest of a person (emphasis added)."

While I take consistency with ordinary usage as a desirable feature of any philosophical theory, its importance may be outweighed by other considerations. Why did Pritchard decide to drop the significance condition? The main reason is this:

[w]e shouldn't expect an account of the metaphysics of lucky events to be responsive to such subjective factors as whether an event is the kind of thing that people care about enough to regard as lucky. (Pritchard, 2014, p. 604)

Pritchard thus wants to restrict his analysis of luck to its *metaphysics*. This metaphysical aspect of luck—that lucky events share a particular modal profile—has been part and parcel of the modal account of luck since its inception. What has changed, however, is that Pritchard now thinks that the significance of an event is relevant not to the phenomenon of luck itself but only to our *judgements* concerning luck.

There are some reasons to resist this view, however. First, the same event may be lucky for person A, but not for B (Rescher, 2001, pp. 19–20). Suppose I win the lottery. This is lucky for me, but the fact that I won is not lucky for you, if you have nothing to do with me and live at the other side of the world. We can explain this by saying that the event is significant to me, but not to you. Importantly, the difference does not, as Pritchard maintains, merely concern our judgements of luck: even impartial observers would judge the event to be a case of luck for me, but not for you.

My second reason for thinking that lucky events are necessarily significant to someone is that it is precisely this property that sets the notion of luck apart from different notions that share the same metaphysical profile. For example, what distinguishes concepts like ‘accidentality’ or ‘unlikelihood’ from luck is precisely the fact that lucky events are necessarily significant, whereas accidental or unlikely events need not be, even if they too concern events that could have easily failed to occur. When Pritchard drops the significance condition to focus on the metaphysics of luck, it seems that rather than analysing the concept of luck, he is analysing one of these related concepts instead.

Thus, while I think Pritchard is right (or roughly right, as we will see below) as far as the metaphysics of lucky events go, I think these *metaphysics* do not suffice for an adequate account of luck. Luck depends on the relation between an event and an agent, not just on the metaphysical properties of the event considered in isolation. It has a modal as well as a significance dimension.⁹

For these reasons, my account of luck includes a necessary significance condition:

⁹ For a different defence of the claim that luck requires significance, see (Whittington, 2016). I will discuss Whittington’s argument in the next chapter, as it relates particularly to *epistemic* luck.

SIGNIFICANCE: An event E is lucky for an agent S only if E is significant to S (or would be significant to S were she to be availed of the relevant facts).¹⁰

The qualification between brackets is necessary because it is clear that the event that Jane wins the lottery is lucky for Jane even if, at the moment, she forgot that she bought a ticket. Even if currently, she doesn't accord it any significance, she would if she were to be availed of the relevant facts.

I take the notion of 'significance' as an undefined primitive. What is clear is that events can be significant in a positive sense or in a negative sense. Winning the lottery is an example of the former, getting struck by lightning an example of the latter. The significance required for luck can be of either kind. This allows for a very natural interpretation of the difference between good and bad luck: an event is a case of good luck for agent S insofar as the event is positively significant to her, whereas it is a case of bad luck for S insofar as it has negative significance to her. When an event is neither a case of good luck nor a case of bad luck, I will refer to it as 'non-lucky'.

So, contrary to MAL, my alternative modal account of luck includes a necessary significance condition on luck. A second distinction between Pritchard's modal account of luck and my own consists in the interpretation of what it is for an event to be easily possible in the first place. As we saw, for Pritchard, an event E is easily possible if and only if E occurs in a world similar to the actual world. I agree with Pritchard that the similarity of the nearest world where the event obtains is relevant for how easily the event is possible, but I disagree that this modal distance is the only factor of relevance.

¹⁰ This condition is based on Pritchard's original significance condition on luck (2005a, p. 132).

In particular, it seems that how easily something could have occurred depends partly on the *proportion* of relevant worlds where that event occurs. Take, for example, the following case. I am toying with a random number generator. I set it so that it picks a random number out of a million. The number it generates is 34.058. Could it *easily* have generated any of the other numbers, say number 34.057? On the one hand, it seems we must say yes, since the nearest world where that number is generated is very much like our own (depending on how the numbers are generated just a few bits and bytes would have to differ). But on the other hand, there also seems to be a sense in which it is not at all easily possible for any of the particular numbers to be generated. After all, the generator could just as easily have spit out any of a million different numbers.

In the above case, our intuitions thus seem to pull in different directions. I think this shows that how easily something is possible depends not just on one, but on two factors: modal distance *and* relative proportion of possibilities. In the above case, the modal distance between the actual world and the world where the event in question obtains is very small. This explains our intuition that the event could have easily obtained. But the *proportion* of all easy possibilities where the event obtains is also very small, which explains our intuition that the event could *not* have easily obtained.

The relevance of this point for our present story is that if we want to maintain that degree of luck depends on how easily the relevant event could have failed to obtain, this will depend *not only* on the modal distance between the actual world and the nearest world where the event fails to obtain, but *also* on the proportion of easy possibilities where the event fails to obtain. The higher this proportion, the luckier the event. Consider again the above case, but now suppose I bet a large sum of money on the number 34.057 being generated. If it is, this is surely a case of luck:

the proportion of possibilities where this event fails to obtain is very high, *and* the nearest world where the event fails to obtain is very similar to the actual world. But now suppose that I bet on the number 34.057 *not* being generated. Suppose I am right. Is it a matter of luck for me that this event occurred? On the one hand, it seems it is, for the nearest world where my number is generated is very similar to the actual world. But on the other hand, the proportion of possibilities where the event fails to obtain is very low. There is a sense in which the event could have easily failed to occur, and there is a sense in which it could not. I think the correct explanation of such cases is that the modal closeness of the nearest world where the event fails to occur means that the event is somewhat lucky, but the relative proportion of nearby worlds where this is the case means that the degree of luck is not very high. In any case, it seems that the degree of luck in the case where I bet on my number being generated is far greater than in the case where I bet on it not being generated. We cannot explain this difference in luck exclusively in terms of modal distance, since the *nearest* worlds where the relevant events fail to obtain are equally close to the actual world in both cases. We thus need to appeal to relative proportions of possibilities to explain this difference.

The fact that degree of luck depends partially on proportion of relevant possibilities and not just on modal distance has some consequences for the interpretation of the relevant possibilities in question. As we have seen, Pritchard seems to draw on a broadly Lewisian conception of possible worlds.¹¹ I will now argue that such a conception of

¹¹ Pritchard is less than clear on what he considers to be the metaphysical nature of possible worlds. However, both the fact that he refers exclusively to Lewis' work and his examples suggest a realist interpretation. The main point that I develop here is that such a conception is unsuitable for an account of luck. This point stands irrespective of the account of possible worlds that Pritchard in fact endorses.

possible worlds is unsuitable for an analysis of luck that acknowledges the fact that luck depends partly on proportions of possibilities.

The problem is that on a realist picture of possible worlds such as Lewis' picture, we should distinguish uncountably many *distinct* nearby possible worlds where the relevant initial conditions for the event are the same as in the actual world. On Lewis's conception, possible worlds are *maximal* in the sense that they contain everything that is spatio-temporally related to each other. This means that *different* spatio-temporal relations entail different possible worlds. For example, in our number-generator case, there will be *distinct* worlds not just for each different generated number, but also for each different possible number of grains of sand in the Sahara, or for each slightly different distribution of all these grains of sand. In a similar vein, there will be different worlds for every different distribution of water molecules in the universe, for each different number of possible hairs on my head, etcetera. The point is that even if we restrict our attention just to the close worlds where the relevant initial conditions for the event are the same as in the actual world, there will still be uncountably many distinct such worlds where my number is generated, just as there are uncountable many distinct such worlds where it is not. This is problematic because if we want to do justice to our intuition that there are far *more* nearby possibilities where a number different from the one I picked is generated, we cannot simply appeal to relative proportions of such worlds.

There are at least two ways to avoid this problem. One solution would be not to talk of relative proportions but instead to define a *probability measure* over the relevant nearby possible world space. We would then be able to say that the difference in degree of luck between my number being generated and my number not being generated is due to the far smaller *probability* of the former compared to the

latter. On this interpretation, degree of luck depends not just on modal closeness, but also on the probability of the event in question. Interestingly, this brings MAL closer to Nicholas Rescher's probabilistic account of luck (Rescher, 2001, 2014). On Rescher's account, (good) luck ($\lambda(E)$) is a product of the significance of an event (Δ) and the probability that the event did not occur ($1 - p$):

$$\lambda(E) = \Delta(E) - \Delta(E)p(E) \text{ (Rescher, 2014, p. 624)}$$

On Rescher's account, the modal closeness of the nearest world where the event fails to occur has disappeared completely from the analysis of luck. In the next section, I argue this is a mistake. For now, let us look at a second way out of our difficulty.

This solution is to abandon Lewis' realist interpretation of possible worlds. On Lewis' conception possible worlds are maximal entities—complete alternate universes, or 'ways absolutely everything might have been'. It is this feature that gave rise to our problem. There are many alternative interpretations of possible worlds available in the literature. However, most of them still conceive of possible worlds as *maximal* entities, and these may thus generate analogous problems to the one under consideration. Particularly instrumental for our present purposes, however, is the recent interest in partially specified *possibilities* rather than maximally specific possible worlds.¹² Whereas possible worlds are maximal entities, possibilities represent ways *some* things could have been, leaving much else undetermined. An example would be the possibility that I throw a six with a fair die. This possibility does not entail anything about the weather at the time of my roll, or the

¹² The general modal semantics were originally proposed by Lloyd Humberstone (Humberstone, 1981). It has been applied and developed in various ways in (Edgington, 1985; Holliday, 2014; Ruffitt, 2015).

exact location of the die after it, whereas Lewisian possible worlds would be determinate in this regard.¹³

In the above example, while there are uncountably many distinct possible worlds where my number is generated, there need not be uncountably many distinct possibilities where it is. Depending on our specification, there may be more or less distinct possibilities. For example, without further specification, there are six distinct possibilities concerning the face on which a particular die lands, corresponding to its six faces. Thus, while there may be uncountably many distinct possible worlds where a particular die lands on a six, there is only one out of six *possibilities* where it does so, under this specification of the relevant possibilities. This saves our way of talking about relative proportions of possibilities, and so provides a second possible solution to our problem.¹⁴

The above points lead me to propose the following Alternative Modal Account of Luck (AMAL):

AMAL: Event E is lucky for agent S relative to initial conditions I iff

1. E is significant to S (or would be significant, were S to be availed of the relevant facts), and
2. E actually occurs, but could have easily failed to occur, given I.

¹³ Kripke called such state descriptions ‘miniature possible worlds’ (Kripke, 1980, p. 16). However, for Kripke, these miniature worlds are an abstraction: his modal operators still quantified over maximally specific possible worlds. In the present framework such miniature worlds are taken as the basic entities over which the modal operators quantify.

¹⁴ A possible problem for this approach is that even if we distinguish between just two possibilities: p and $\neg p$, standard interpretation of the modal operators requires us to treat $\Box p$ as distinct from $\Box\Box p$, and $\Diamond\Box p$, etc. We would then be back with an infinite number of distinct possibilities on every specification. I will ignore this complication here. Thanks to Barteld Kooi for bringing it to my attention.

AMAL thus departs from Pritchard's most recent formulation of MAL in two ways. First, it incorporates a necessary significance condition. In this sense, AMAL is closer to Pritchard's original formulation of MAL. Second, AMAL is formulated so as to allow for an interpretation of the notion of 'easy possibility' where how easily something is possible does not just depend on modal distance but also on the relative proportion of nearby worlds where the event in question obtains. Because of this, I can, whereas Pritchard cannot, explain the difference of degree of luck involved in cases like our number-generator cases above, where a difference of degree of luck cannot be explained in terms of modal distance alone.

There is a third way, however, in which I suspect my views on luck differ from Pritchard's views. This third difference concerns the specification of the 'relevant initial conditions' for an event.

As we have seen, Pritchard says that the relevant initial conditions for an event "need to be specific enough to pick out a particular kind of event that we want to assess for luckiness, but not so specific as to guarantee that this event obtains" (Pritchard, 2005a, p. 599). Admittedly, it is not quite clear from the context how strong Pritchard intends this requirement to be. I think it is good to stress, however, that I do not recognise *any* principal restriction on the relevant initial conditions of an event to be evaluated for luck. This is important because as we have seen, the degree of luck to which an event is subject will depend on how we specify these initial conditions. To take Pritchard's own example, winning the lottery is a case of luck if the relevant initial conditions just include buying a random ticket. However, "if one includes in the initial conditions for the event the demand that the balls fall into the lottery machine in a certain way, then one will no longer generate the desired result that the event is lucky" (Pritchard, 2005a, p. 131). That is, given that we specify in the initial condi-

tions for the lottery win in such a way that they include the number of the ticket actually bought, as well as the number actually drawn, it will no longer be easily possible to lose, *given* that these conditions are both met.

Pritchard seems to take this to motivate a restriction on admissible initial conditions. In his view, the above conditions are not the 'right' conditions for evaluating the luckiness of a lottery win. This betrays the view that luck is an objective phenomenon: for any event, it is a matter of fact whether it is a case of luck or not. My own view is that whether an event is a matter of luck depends on our description of that event, and there is no objective criterion for deciding what description we 'should' use.

There are several ways in which degree of luck may be influenced by non-objective factors. First, whether an event is lucky or not depends on how its relevant initial conditions are specified. Here I think our choice will often be determined by pragmatic factors. On my view, lottery wins are paradigmatic instances of luck partly because we are usually interested in lottery wins only under certain descriptions. Usually, we do not care whether my win is lucky given that a particular combinations of numbers was selected as the winning combination. This is so simply because we usually do not possess this kind of information at the time the assessment of luck is relevant, e.g. when we are deciding whether to buy a ticket or not.

Second, even *given* a set of initial conditions, subjective factors influence whether the event is lucky or not (or the degree of luck to which it is subject). For as we have seen, the degree of luck to which an event is subject also depends on how we specify the relevant possibilities. Here again we have a choice. And again this choice will depend on pragmatic factors such as the information at our disposal, and our interest in evaluating the event for luck in the first place.

Finally, our account of luck includes a significance condition, and I have argued that degree of luck depends on degree of significance. Of course, the degree of significance of an event is a subjective factor: there are some people for whom surprise Justin Bieber concerts are highly significant, but I am not one of them. So there are at least three ways in which the degree of luck to which an event is subject may depend on non-objective factors.

This non-objective aspect of luck is rarely noted, and so deserves to be stressed.¹⁵ Interestingly, considering these non-objective factors influencing degree of luck brings out a similarity between the concepts of luck and probability. As we argued above, degree of luck depends on our description of the event in question, and there will often be no objective criterion for deciding which description to choose. This reminds one of the ‘problem of the reference class’ in probability theory.¹⁶ For just as the degree of luck to which an event is subject depends on how the event is described, the probability of an event also depends on the description of the event rather than of the event itself.

The problem of the reference class is widely known and discussed in the theory of probability.¹⁷ As an example, suppose we try to determine the probability that Jane wins the lottery. If we describe Jane as one of the 10.000.000 ticket holders, this probability will be 0.000001. However, we may also describe Jane as ‘the winning ticket holder’. If we describe the case in this way, the probability that Jane will win is 1. Whether an event is probable or not depends crucially on how we describe the event in question. If there were an objective-criterion for deciding how to describe the

15 Nicholas Rescher can perhaps be seen to gesture in this direction, when he says that “[t]he idea of good (or bad) luck is inherently context relative” (2014, p. 623). This may be due, as we will see now, to the fact that he considers luck in a probabilistic framework.

16 Thanks to Jeanne Peijnenburg for bringing this to my attention.

17 See for example (Reichenbach, 1949; Hájek, 2007).

event, this would not be problematic, but the problem of the reference class arises precisely because there does not seem to be any objective criterion for choosing between rival descriptions. Similarly, there does not seem to be an objective criterion for deciding between rival specifications when we are considering the degree of luck to which an event is subject.

While the problem of the reference class is widely acknowledged in probability-theory, I am not aware of anyone noticing a similar problem for the theory of luck. That luck is subject to a reference class problem is to be expected, however, especially given that luck seems to depend partially on relative proportions of possibilities, which may be described in terms of probabilities. But even if it would not so depend, there would still be a problem of the reference class. For degree of luck also depends on our specification of the initial conditions for the event, and there may not be an objective criterion favouring one specification over another.

The above does not imply, of course, that there are no objective facts about luck at all, or that whether an event is lucky or not is a matter of opinion only. For it will be an objective matter of fact whether an event is lucky for a particular subject *given a particular description*. Here people can be wrong: it is possible to estimate that an event will not occur in a significant proportion of nearby possible worlds where the relevant initial conditions for that event are the same as in the actual world, when the event in fact *will* occur in most of these worlds. But the modal account of luck will not provide us with reasons for preferring one set of initial conditions over another, or one specification of the relevant possibilities over another.

Nor does the above imply that there are no reasons *at all* for favouring one description of an event over another. On the contrary, we can have many reasons, pragmatic, epistemic or otherwise, to favour some initial conditions over

others in our description of that event. For example, when we are deliberating whether to buy a ticket for the lottery, we are usually not in possession of the information that determines the winning ticket. From our epistemic perspective at that moment, a win will surely be lucky. The point is rather that for someone else, there may be equally valid reasons for preferring a different description.

Even if it does not make sense to talk about events being lucky *simpliciter*, we can thus still accommodate the fact that there are certain sets of relevant initial conditions for events that are of particular interest to us. Such will be our strategy in Chapter 4, when I discuss the various kinds of luck that may be of epistemic interest.

This concludes our exposition of the alternative modal account of luck that I will draw upon throughout this book. It differs from Pritchard's account of luck in three respects: i) it recognizes a necessary significance condition on luck, ii) it recognizes the influence of proportions of nearby worlds where the event in question fails to obtain, and iii) it classifies as a 'subjective' account of luck since I recognize no restrictions on the specification of an event's 'relevant initial conditions'. Despite these differences, however, it should be clear that the account owes much to Pritchard's account of luck. In particular, it is still motivated by the central idea that lucky events share a particular modal profile, that lucky events are events that could have easily failed to occur. For this reason, one might view AMAL not as an alternative to MAL, but rather as a refinement of it. A refinement, however, that I have argued is an improvement of the original.

With my account of luck thus stated, I will now continue to discuss its merits by comparing it to its two main rivals: the 'lack of control' account of luck and the probability account of luck. In the final section I will discuss some objections that have been raised to modal accounts in general, and see how the present one fares.

3.3 RIVAL ACCOUNTS OF LUCK

Lucky events have a modal dimension. This observation motivated the account developed in the previous section. A modal account of luck is not the only available option, however. In this section I review two alternative accounts of luck. I argue that the modal account developed in the previous section should be preferred.

3.3.1 *The Lack of Control Account of Luck*

The first rival of the modal account has been named the 'Lack of Control' account of luck (LC). It can be captured by the following principle:

LC: An event is lucky for a given agent, *S*, if and only if the occurrence of such an event is beyond—or at least significantly beyond—*S*'s control (Lackey, 2008, p. 256).¹⁸

Many authors agree that lack of control is necessary for luck (e.g. Greco, 1995; Riggs, 2007; Statman, 1991; Zimmerman, 1987). The motivation for this view is clear. Paradigm cases of luck, such as winning the lottery, or finding a treasure are uncontroversially beyond our control. While I may decide whether or not to buy a ticket, I cannot decide to *win* the lottery. When I *do* have such control—suppose I rigged the lottery—then my win would not be a case of luck. Luck, at first sight, seems to come and go with the absence of control.

What is it to have control over the occurrence of an event? Neil Levy writes about the control necessary to rule out luck: "an agent has direct control over *E*'s occurrence

¹⁸ Lackey ultimately thinks LC is wrong. I nevertheless use her formulation because it is particularly clear.

when he can bring about E's occurrence by virtue of performing some basic action which (as he knows) will bring about E's occurrence" (Levy, 2011a, p. 19). Thus I have direct control over whether to put on my sunglasses just in case there is some basic action that I can perform that will result in my putting on the sunglasses, and I know that performing this action will bring about that result.

As formulated by Lackey, LC entails that a lack of control is *sufficient* for luck. This is highly implausible, however. As we have seen, lucky events are *significant* to someone. That somewhere in India two children are playing a game of chess is completely out of my control, yet it makes no sense to say that I am lucky that this event occurred. For this reason, Wayne Riggs, a prominent defender of a lack of control account of luck, adds a significance condition like the one discussed above to his account (Riggs, 2009).

Even with a significance condition added, LC will not be sufficient for luck, however, since there are many events that are both significant to us *and* completely out of our control, but which are not in the slightest cases of luck. Take the event of the sun rising tomorrow. This event is both completely beyond my control and highly significant to me. Yet it sounds very strange to say that it is lucky for me that it did (Latus, 2003, p.467). Similarly, that wine tastes the way it does to me, or that the U.S. did not start a war with Europe yesterday are events that are both significant to me and out of my control. Yet these are not cases of luck.

Wayne Riggs argues that there are cases where such events *are* lucky (Riggs, 2009). He presents a case of two explorers that are captured by tribesman. Just before they will be murdered, a solar eclipse happens, and this leads the tribesmen to set the explorers free. The first, who was not expecting this eclipse at all, cries out that the eclipse must have been one of the luckiest events in his life. The second, who *was* expecting the eclipse, responds that it was

not a case of luck for him at all, since he had been expecting the eclipse all along, and suspected it would be enough to set them free. Riggs wants both explorers to be right here, and as a result incorporates another necessary condition on luck:

B: S did not successfully exploit E for some purpose (Riggs, 2009, p. 219).

Together with LC and a significance condition, B is supposed to be sufficient for luck. Adding B allows Riggs to say that both explorers are right in their judgements about luck. But it is not clear how adding this condition helps establish the sufficiency of the account of luck. Even if I do not exploit the sun's rising this morning, it is still not lucky for me that it did. So Riggs' conditions are still not sufficient for luck.

Even with the added conditions, LC is not sufficient for luck. Could LC then at least be regarded as necessary for luck? Jennifer Lackey (2008) provides the following counterexample (which I compressed somewhat):

DEMOLITION WORKER: Ramona is a demolition worker, about to press a button that will blow up an old abandoned warehouse. Unbeknownst to her, however, a mouse had chewed through the relevant wires in the construction office an hour earlier, severing the connection between the button and the explosives. But as Ramona is about to press the button, her co-worker hangs his jacket on a nail in the precise location of the severed wires. As it happens, the hanger on which the jacket is hanging is made of metal, and it enables the electrical current to pass through the damaged wires just as Ramona presses the button and demolishes the warehouse.

Lackey argues that in this case, Ramona has control over the explosion, but that this event is nevertheless a case of luck.¹⁹ While I agree with Lackey that LC is not necessary for luck, I do not think that DEMOLITION WORKER brings this out. To see this, we need to separate the lucky events in DEMOLITION WORKER from the non-lucky ones.

We can distinguish various separate events in Lackey's case. For example, instead of focussing on the explosion, we can focus on Ramona's co-worker hanging his jacket in the right spot to close the electronic circuit. We may assume this event is lucky for Ramona, for it seems to satisfy LC (and, we may suppose, the other conditions necessary for luck). Ramona does not have control over where her co-worker chooses to hang his jacket. This event thus does not provide a counterexample to the lack of control account. But neither is it supposed to. The supposed counterexample to the lack of control account is the event of the building exploding. It is this event that Ramona *has* control over, and it is this event, that, as Lackey argues, is still a matter of luck.

Why does Lackey think that the explosion is a matter of luck for Ramona? Lackey submits that "What DEMOLITION WORKER shows, then, is that although an event may be within a given agent's control, that the agent has such control can itself be largely a matter of luck" (Lackey, 2008, p. 259). Lackey then infers from the fact that it is a

¹⁹ Lackey argues that Ramona has control over the explosion because "Ramona's pressing of the button—which is an activity that she could have refrained from engaging in—is what is directly responsible for the explosion" (Lackey, 2008, p. 259). Note that it is not clear whether this suffices to satisfy our definition of control above. In particular, it is not clear whether Ramona *knows* that her action will set off the explosion, due luck involved in that belief being true. For the moment I will grant Lackey's assumption, however. See also Chapter 4 for more on knowledge-precluding luck.

matter of luck that Ramona has control over the explosion, that it is a matter of luck that the building exploded.

This inference is invalid, however. From the fact that it is lucky that one has control over event A, it does not follow that A itself is lucky. Suppose I am almost run over by a car while crossing the street, such that it is lucky that I am still alive. Since being alive is required for having control over my feet, this entails it is a matter of luck that I have control over my feet. On Lackey's picture this would mean that if I then, shortly after my near-accident, move my feet to return to the side walk, these movements will all be cases of luck, since it is a matter of luck that I have the required control over my legs. This seems counter-intuitive.

We can explain where the inference goes wrong by stressing again the importance of the specification of the relevant initial conditions for the event we want to assess for luck. *Given* that I have control over my legs, the resultant movements are not lucky. The same holds for Lackey's case. *Given* that the mouse and the co-worker's actions coordinated to close the electrical circuit, the explosion will not be a case of luck. Moreover, Ramona only has control over the explosion in Lackey's sense *given* that these conditions are met. The moral to be drawn is that we cannot conclude from the fact that Ramona was lucky to have this control in the first place, that the event that resulted from this control was a case of luck.²⁰ Lackey's counterexample thus fails because it relies on an invalid inference. Pending other counterexamples, it seems that lack of control is indeed necessary for luck.

The fact that luck requires lack of control is compatible with our present modal account of luck, however.²¹ For events that could have easily failed to occur will automat-

²⁰ The same argument is made by Neil Levy (2011a, p. 22-23).

²¹ Pritchard provides a similar argument for this claim (Pritchard, 2005a, 2014). The differences between our accounts do not impact the validity of the argument.

ically be out of one's control, according to the definition above. If an event could have easily failed to occur, the subject will not know there is a basic action she can perform that will achieve the desired result. So all events judged lucky on the modal account will satisfy the lack of control condition. Properly understood, the lack of control account and the modal account are quite compatible.

3.3.2 *The Probability Account of Luck*

The lack of control account is not the only alternative to the modal accounts of luck. We already met another prominent account of luck: the probability account. According to the probability account, luck is inversely related to probability: the more probable an event is, the less lucky it is. Stephen Hales writes that the probability account of luck is especially prominent among mathematicians and scientists, and can be traced back to de Moivre (Hales, 2016).

If by saying that a man has had good Luck, nothing more was meant than that he has been generally a Gainer at play, the Expression might be allowed as very proper in a short way of speaking ... (de Moivre, 1738, p. iii)

While it is not easy to extract a definition of luck from the passage above, it is clear that de Moivre considers luck to attach to events that defy the odds. A more recent philosopher advancing a probability account of luck in a more explicit way, is Nicholas Rescher (2001, 2014).

As we saw before, Rescher defines the measure of luck of an event E ($\lambda(E)$) as the product of the 'stake' of E ($\Delta(E)$), and the 'risk' of the event (probability of the event failing to occur) ($1 - p(E)$). This results in the following formula for luck:

$$\lambda(E) = \Delta(E) - \Delta(E)p(E) \text{ (Rescher, 2014, p. 624)}$$

Rescher's variable 'stake' ($\Delta(E)$), the difference in value between the event occurring and it not occurring can naturally be interpreted as a measure of its significance. Thus, both an increase in stake and an increase in risk will raise the luckiness of an event, on Rescher's conception. Events that are certain to occur ($p(E) = 1$) are completely non-lucky, as are events that are completely insignificant.

As a probability account, Rescher's account will face the reference class problem. This need not be a problem, but it does mean that luck will be to some extent a subjective phenomenon on the probability account of luck.

What can be said in favour of the probability account of luck? As we have seen, the probability account seems to be able to accommodate many of the points made in Section 3.2. It includes a significance condition, as well as a probability condition. It can thus accommodate both the fact that degree of luck depends on degree of significance and on probability. Accordingly, gets most of the cases right. Our paradigmatic cases of luck, finding a treasure and winning the lottery, can be plausibly regarded as events with high stakes as well as low probabilities. They will thus rightly be identified as cases of luck according to the probability account of luck.

Contrary to the account proposed in the previous section, Rescher's probability account does not take into account the influence of modal distance on luck. I think this is a mistake for the following reasons.

First, suppose we take a closer look at the event of the sun rising tomorrow. We said above that such events like should not be counted as lucky events. However, a pure probability account would have to judge it *somewhat* lucky, albeit only to a very small degree. This is so, because there is presumably *some* small probability that the sun will not rise tomorrow, and the event is thus somewhat risky. Given that the stakes are clearly very high here, this means that the event will be somewhat lucky on the probability ac-

count. I think this is still the wrong result, and it is avoided by adopting a modal account of luck, for while it is possible that the sun does not rise tomorrow, we may assume this is not *easily* possible, on any plausible description of the case.²² So events like this will not satisfy our definition of luck.

The rising of the sun thus provides a case where the probability account of luck overgenerates. It also undergenerates. Lottery cases provide examples of such undergeneration. Suppose I buy a ticket to a fair lottery, and before hearing anything about the results come to believe that my ticket does not win. Among epistemologists it is common to deny these beliefs the status of knowledge. One explanation for this fact is that the truth of my belief, even though overwhelmingly probable, is still a matter of luck, given that I formed it on the basis of statistical considerations alone. This indicates that high probability by itself is not enough to exclude luck altogether. As long as there are nearby possibilities for the event to fail to occur, the event will be lucky to some extent.

We discussed the way our modal account treats such cases above. There will be close possibilities where I form a false belief on the same basis. Yet, there are relatively few of such possibilities, meaning that our account can explain why we would consider such cases to be genuine cases of luck, even if the degree of luck is not very high. Again, it seems the modal account is better able to deal with the cases than a purely probabilistic account of luck.

²² While I have said that luck is a gradual notion, and that our account can accommodate this by saying that the more easily possible the failure of the even, the less luck, generally speaking, there will always be distinction between things that are easily possible and things that are not. Thus, while the boundary between what is easily possible and what is not is relatively flexible, it is not *endlessly* flexible. To say that it is easily possible that the Sun would not rise tomorrow would be to give up the distinction between things that are easily possible and things that are not altogether.

Finally, the empirical literature, in particular with the work of Karl Halvor Teigen, clearly indicates that not just probabilities, but also modal distance influences people's ascriptions of luck. Teigen found that certain judgements of luck do not depend on probability but rather on counterfactual closeness of alternatives: subjects were found to judge events as dissimilar in the amount of luck that was involved, even while explicitly admitting that the probability for the events was the same (Teigen, 2005). Teigen's examples included judgements of luck concerning roulette-wheel outcomes. If the outcome was physically close to the number on which subjects had placed bets, the subjects judged the event to be a case of (bad) luck. The further the outcome was physically removed from this number, the less willing subjects were to regard the event as a case of luck, even if they agreed that the probabilities of the relevant events were the same, indicating that our judgements of luck here are responsive to modal distance rather than to probability.

Thus, Rescher's account is incomplete due to its neglect of the particularly modal dimension on luck. This does not mean that probability is irrelevant for luck. As we have seen in Section 3.2, it is plausible that degree of luck depends on relative proportions of possibilities, and one way to capture this dependence is by including a (inverse) probability condition on luck. We have also seen that significance is important for luck. But what I take the above considerations show, is that there is more to luck than just probability and significance. What is missing from Rescher's account is a specifically modal factor, a factor that relates degree of luck to the modal *distance* of the nearest possibility where the event fails to occur to the actual world. The only account of luck that recognizes all factors relevant for luck, or so I maintain, is the modal account of luck developed in this chapter.

3.4 OBJECTIONS TO THE MODAL ACCOUNT

In this chapter, I have so far provided an overview of my preferred modal account of luck. I have argued that it is to be preferred over its rivals. I will now complete its defence by discussing the way it handles some counterexamples that have been raised against modal accounts of luck generally.

Consider first a counterexample due to Hiller and Neta:

ENHANCED NOGOT: Naira believes that Jones owns a Ford, and so she uses disjunction-introduction to infer, and so comes to believe, that *either Jones owns a Ford or Brown will not win a Grammy award*. As it turns out, Brown is not even a musician and has nowhere near the skills needed to win a Grammy, so the second disjunct is true. Indeed, we may suppose that Brown is so lacking in even the most rudimentary musical ability that the second disjunct is true in nearly all, if not all, nearby possible worlds in which Naira forms her belief in the same way. (Hiller & Neta, 2006, p. 308)

The thought here is that Naira's belief in the disjunction is still luckily true, even though it does not satisfy our modal condition: for given the way Naira forms her belief, the belief could *not* have easily been false due to the utter lack of musical ability on Brown's side.

As Goldberg points out, however, ENHANCED NOGOT fails as a counterexample to the modal account of luck (Goldberg, 2015, p. 277). Take as the relevant event the formation of the true belief about Jones and Brown. This event can fail to occur in two ways: first, the same belief is formed, but it is false; second, a *different* belief may be formed. Hiller and Neta are right that the first way is not

easily possible, but they gloss over the fact that presumably, the second way *does* represent an easy possibility. Given the way Naira formed her belief, she might have easily formed a *different* belief than the one she actually did. In line with our intuitions, the formation of Naira's true belief about Jones and Brown satisfies our conditions on luck.

Jennifer Lackey provides another well-known counterexample to Pritchard's modal condition on luck:

BURIED TREASURE: Sophie, knowing that she had very little time left to live, wanted to bury a chest filled with all of her earthly treasures on the island she inhabited. As she walked around trying to determine the best site for proper burial, her central criteria were, first, that a suitable location must be on the Northwest corner of the island—where she had spent many of her fondest moments in life—and, second, that it had to be a spot where rose bushes could flourish—since these were her favourite flowers. As it happens, there was only one particular patch of land on the Northwest corner of the island where the soil was rich enough for roses to thrive. Sophie, being excellent at detecting such soil, immediately located this patch of land and buried her treasure, along with seeds for future roses to bloom, in the one and only spot that fulfilled her two criteria.

One month later, Vincent, a distant neighbour of Sophie's, was driving in the Northwest corner of the island—which was also his most beloved place to visit—and was looking for a place to plant a rose bush in memory of his mother who had died ten years earlier—since these were her favourite flowers. Being excellent at detecting the proper soil for rose bushes

to thrive, he immediately located the same patch of land that Sophie had found one month earlier. As he began digging a hole for the bush, he was astonished to discover a buried treasure in the ground. (Lackey, 2008, p. 261)

Lackey argues that Vincent's discovery fails Pritchard's modal condition, and yet is a clear case of luck. Since the example does not seem to turn on the difference between Pritchard's conditions and our own, the case would pose a problem for our account as well. As I shall argue, however, Lackey's example fails to provide a counterexample to our modal account of luck.²³ To this end, I will disambiguate between various interpretations of the case, and show that where the events are lucky, they satisfy our condition (AMAL), and where they are not, they fail to satisfy it.

What is the relevant event in BURIED TREASURE? As Lackey presents the case, it seems that what is lucky is Vincent's discovery of the treasure. Lackey contends the event does not satisfy a modal condition on luck. Why so? It seems that given the way the case is set up, it is *not* an easy possibility for Vincent to fail to discover the treasure. That is, given the way the case is set up, the relevant event could not have easily failed to obtain, and thus, by our light, should not be counted as lucky.

As I have been stressing throughout this chapter, however, luck depends on how we describe the case, in particular on how we specify the relevant initial conditions for an event. Looking at BURIED TREASURE, AMAL fails to be satisfied only if we include in the relevant initial conditions for Vincent's discovery all the information provided (about both Sophie's and Vincent's mother's preferences, their skill at detecting proper soil for rose bushes, the ter-

²³ Nor do I think it succeeds as a counterexample to Pritchard's account. Here I will restrict attention to my own account.

rain on the Island, etc.). For it is *only* given all this information that Vincent could not have easily failed to find the treasure. Had Sophie's mother loved a different kind of flower, or were the island made up slightly differently, for example, then it would be an easy possibility for Vincent to fail to find the treasure.

Now that it is clear that we are considering the question whether Vincent's discovery is a case of luck, *given* that everything about the set-up of the case is kept fixed, I myself do not have such a strong intuition that his find is clearly a case of luck. Sure, it seems to be a case of luck that circumstances combined in such a way to make the discovery possible, but as we have seen above, we cannot validly infer from the fact that it is a matter of luck that preconditions for an event combined in the way they did that it is a matter of luck that the event occurred, given that these conditions combined as they did.

So even if Vincent's discovery is not itself lucky, there is still a lot of luck present in Lackey's case. Given some reasonable description of the case, it is a matter of luck that both Vincent's and Sophie's mother share the same favourite flower, or that both Sophie and Vincent are excellent at determining the proper soil for rose bushes to grow on. We can acknowledge this while denying that—given that these events occurred—Vincent's discovery specifically is lucky. So we can admit to Lackey's claim that that BURIED TREASURE is riddled with luck while resisting the claim that it is a counterexample out account of luck.

One of the reasons why Lackey thinks that Vincent's find is itself a matter of luck is that "not only does he[Vincent] have no reason to think that the treasure has been buried in the particular location in which he was digging, he also has no reason to think that a treasure has been buried *anywhere* on the island" (Lackey, 2008, p. 262, *italics in original*) But this just shows that the discovery will *seem* to be lucky from Vincent's perspective. The modal account can accom-

modate this fact without problems. For example, it is plausible that Vincent will not consider the fact that Sophie's mother favourite flowers were roses—a fact of which he is ignorant—to be a relevant initial condition for his discovery. From Vincent's perspective, therefore, his find *could* have easily failed to occur. And so the modal account correctly predicts Vincent to regard his discovery as a case of luck.

In sum, it seems that when the details of Lackey's case are properly spelled out, it does not present a counterexample to our modal account of luck. Doing so once again brings out the fact that whether an event is a case of luck depends on how we describe the event in question, on what we take to be its relevant conditions.

A third and final putative counterexample to modal accounts of luck is due to Sandy Goldberg:

TRICK-OF-LIGHT: Through an exceedingly rare trick-of-light, the light rays reflecting off of one particular bush will generate the appearance of a rock formation to a person located in one particular spot 20 feet away from the bush. The trick of light is a rarity among rarities: it happens only once a millennium, and the event then only has its effect when there is a a person who happens to occupy that particular spot 20 feet away, and who happens to be looking in that right direction at the time. Now it happens that there *is* a rock formation on the hill. However, given its location behind the bush, this rock formation is not visible to anyone standing at the spot in question. Imagine finally that it just so happens that there is a subject, *S*, standing at just that point at just the right time looking in just the right direction, and so who has a perceptual experience

as of a rock formation located on the hill. (S has no other access to the rock formation; she does not go closer, or move to get a different orientation.) What is more, the rock formation in question has always been at that spot: it is part of a largely buried rock mass that has been at this site for millions of years, and its presence there reflects fundamental laws of geology. What is more, owing to the laws of nature, the content of the trick-of-light illusion is invariant: whenever (on those every rare occasions) such an illusion obtains, these laws ensure that it comes in the form of an experience *as of a rock formation on the hill*. (No other illusory experience would be compatible with the actual laws of nature.) Now it so happens that S, standing at just the right spot, at just the time when the trick-of-light obtains, and who, looking at the bush in just the right way at just the right time, has a perceptual experience as of a rock formation on that hill. On this basis S comes to believe that there is a rock formation on that hill. (Goldberg, 2015, pp. 277-278)

Goldberg's argument is directed against a modal account of *epistemic* luck, about which I will have much more to say in the next chapter. However, the argument can be directed at our general modal account of luck as well.

Why would TRICK-OF-LIGHT present a counterexample to the modal account of luck? First, the relevant event is the event that S forms a true belief about the rock formation. Second, as Goldberg argues, it seems that given how S formed her belief, it is *not* the case that this event could have easily failed to occur. After all, S forms her belief on the basis of a hallucination with invariant content 'as of a rock formation'. Moreover, the rock formation is present as

a matter of geological necessity (the implication being that it could not easily have failed to be there). Thus, formed in this way, *S*'s belief could not have easily been false. It seems thus that we are forced to judge that the formation of the true belief is non-lucky. Nevertheless, Goldberg argues, we do have the strong intuition that it is a matter of luck that *S* forms the true belief she does. Since the modal account cannot account for this, it must be wrong.

Why does Goldberg think *S*'s belief is lucky? He says that her belief is lucky because "her belief is in the relevant respects precisely like the subject in the sheep-on-the-hill case, which is a paradigmatic example of epistemic luck" (Goldberg, 2015, p. 278). The sheep-on-the-hill case is Chisholm's famous Gettier case, where we are to imagine a person looking from a distance at a hill, and mistakenly identifying a creature that he sees on the hill to be a sheep (while it is actually a dog). Nevertheless, there is a sheep on the side of the hill that is out of sight. Thus, the man's belief that there is a sheep on the hill is true only as a matter of luck.²⁴ Goldberg's claim is that TRICK-OF-LIGHT is in the relevant respects like Chisholm's case, and thus that since the latter is uncontroversially regarded as a case of epistemic luck, we should be willing to similarly judge that the former is a case of luck.

I do not think, however, that the cases are as alike as Goldberg suggests. For example, it seems that our verdict in Chisholm's case depends on certain tacit background assumptions, one of them crucially being that the hill is not overly populated with sheep. After all, if there were always many sheep on the hill, it seems that it is not particularly lucky that your belief is true, when you form it on the basis of seeing something that looks like a sheep from afar. If there are always many sheep on the hill, this method will not often lead you astray. Our intuitions about Chisholm's

²⁴ See (Chisholm, 1977, p. 105).

case are thus strongest when there are seldom any sheep around.

Goldberg's case is much closer to the first version of Chisholm's case than to the second. For the stable hallucination and the stable presence of the rock formation make it such that your method of belief-formation would rarely lead you astray, just as in the first version of Chisholm's case. In that case, however, the intuition that the agent formed her true belief by luck is much weaker than in the second case. Therefore, Goldberg's case should not be regarded as a 'paradigmatic' case of luck.

Even if Goldberg's case is not a 'paradigmatic' instance of luck, there is still some luck present. In particular, it is a matter of luck that circumstances combined in the way they did, for presumably, S could have easily arrived a little later than she did, or could have stopped at a slightly different location, thus preventing the relevant belief to be formed. It is not a matter of luck however, that *given* that circumstances combined in the way they did, S formed her true belief. Goldberg judges that the latter event is lucky *because* the former is. This is a kind of reasoning about luck that we also saw with Jennifer Lackey. But as we have seen before, luck does not transfer in this way. Thus, by clearly separating the events that we are regarding as lucky, we can explain our intuition that there is luck present in TRICK-OF-LIGHT as well as maintain a modal account of luck.

This concludes the defence of AMAL. We have considered some putative counterexamples, and argued that they fail.²⁵ The modal account of luck developed and defended in this chapter seems to be the best account of luck around.

²⁵ For some other objections to modal accounts of luck, see (Coffman, 2010; Lackey, 2006; Hales, 2016).

3.5 CONCLUDING REMARKS

In this chapter, I defended and developed a particular modal account of luck (AMAL), based on Pritchard's Modal Account of Luck (MAL). In Section 3.1, I provided an overview of Pritchard's Modal Account of Luck. In Section 3.2, I developed and improved the account in various ways. In particular, I argued we should incorporate a significance condition on luck, and recognize the fact that luck depends on proportions of possibilities as well. I also argued that there are various ways in which degree of luck depends on non-objective factors. This gives rise to a problem for accounts of luck parallel to the reference class problem for accounts of probability. In Section 3.3, I discussed the two main rivals to a modal account of luck, the lack-of-control account and the probabilistic account of luck, and argued that we should prefer the modal account of luck to either of these alternatives. In Section 3.4, finally, I defended my modal account of luck against the various objections raised to it in the literature. In the next chapter I will put my present efforts to use by distinguishing various kinds of *epistemic* luck.

EPISTEMIC LUCK

In the previous chapter, I argued in favour of a modal account of luck. In this chapter, I apply the account in discussing various forms of luck that are relevant for epistemology. My main focus will be on two forms of luck identified in the literature: veritic and reflective luck. Whereas it is nearly universally agreed that veritic luck is incompatible with knowledge, it is less clear whether the same holds for reflective luck. I distinguish two kinds of knowledge, and argue that the one, but not the other is compatible with reflective luck.

Why think luck is relevant for the analysis of knowledge? Arguably, one of the first philosophers to notice a (problematic) relation between luck and knowledge is Plato. In the *Theaetetus*, Socrates suggests that jurymen who are persuaded by savvy lawyers that a particular defendant is guilty do not thereby come to *know* that proposition, even if it turns out to be true. What seems to be problematic about the jurors' belief is that, if true, the belief would be true only by *luck*.¹ Here it looks like luck and knowledge are incompatible. However, we will see that not all forms of luck are incompatible with knowledge.

The relation between luck and knowledge has received considerable attention in recent literature.² No consensus has been reached, however, about the exact relation between various forms of luck and knowledge. In this chapter, I argue for my own view on the matter, starting from

¹ See, for example (Chappell, 2013; Williams, 1992).

² Cf. (Baumann, 2014; Hiller & Neta, 2006; Hall, 1993; Pritchard, 2008b; Heller, 1999; Lackey, 2006; Engel, 1992a; Riggs, 2007; Harper, 1996; Pritchard, 2007; Goldberg, 2015).

the modal account of luck developed and defended in the previous chapter.

The structure of the chapter is as follows. In Section 4.1, I provide an overview of the forms of luck that are compatible with knowledge. In Section 4.2, I discuss two varieties of luck that one may think are incompatible with knowledge: veritic luck and reflective luck. While it is almost universally agreed that veritic luck is incompatible with knowledge, the same does not hold true for reflective luck. In Section 4.3, I defend the claim that we can distinguish between kinds, or ‘grades’ of knowledge that are compatible with reflective luck, and kinds that are not. I close with some concluding remarks in Section 4.4.

4.1 BENIGN FORMS OF EPISTEMIC LUCK

Let us call the forms of luck that are relevant to the acquisition of knowledge *epistemic* forms of luck. Not all epistemic luck is *incompatible* with knowledge. I shall call the kinds that are not ‘benign’ forms of epistemic luck. The first person to call our attention to such benign forms of epistemic luck is Peter Unger (1968).³

4.1.1 *Content Luck*

Our first benign form of epistemic luck is called *content* luck. Content luck can be defined in the following way:

CONTENT LUCK: *S*’s belief that *p* is subject to content luck if and only if it is a matter of luck that *p* is true.

³ Unger does not clearly distinguish the different varieties from each other, however. For reasons of clarity, I shall therefore use the labels that have now become standard (Engel, 1992a; Pritchard, 2005a).

Thus a belief is subject to content luck insofar as it is a matter of luck that the proposition believed is true. On the modal account, this means that *S*'s belief that *p* is content lucky, given a set of initial conditions, if and only if *p* is actually true, but could have easily been false, given these conditions.

Also, to count as a case of luck at all, the truth of the belief must be somewhat significant to the agent. Here and beyond, I will assume this latter condition is automatically met when the truth of one's beliefs is concerned. This significance is easily explained for the beliefs that have practical import: we will be able to achieve our desired ends more efficiently when we base our action on true belief. Not all our beliefs have practical import, however. How are we to account for the significance of beliefs that lack such practical import? Here a slight retreat is perhaps in order. Significance, and value more generally, come in different kinds. Things may be morally valuable/significant, or pragmatically, for example. In a similar vein, things may be *epistemically* valuable. One way to look at epistemology is that it investigates rules for engaging in the particularly *epistemic* project of trying to get it right about the world. Such rules will be normative, of course, only for agents actually engaged in this project. So at least *qua* epistemic agents, we may assume that even the truth of beliefs that lack practical import has epistemic significance.

One reason why this is important for an anti-luck epistemology is brought out by the exchange between Nathan Ballantyne and Lee Whittington (Ballantyne, 2014; Whittington, 2016). Ballantyne argues that the requirement that lucky events be significant to the agent concerned spells trouble for anti-luck epistemology because it leads to an unacceptable form of pragmatic encroachment. After all, he argues, we would be able to turn Gettiered beliefs (where there is a problematic amount of luck present) into

ungettiered, knowledgeable belief just by decreasing the significance of the truth of the belief in question.

Whittington responds by drawing a distinction between different *kinds* of value. The kind of significance important for epistemic luck is epistemic value. Whittington argues that we cannot decrease the *epistemic* value of true belief enough to exclude knowledge without changing its truth-value. Since false beliefs by definition cannot constitute knowledge, Ballantyne's strategy fails to establish any problematic encroachment.⁴ The upshot is that all true beliefs have some epistemic value and are therefore positively significant to agents, at least as they are engaged in the epistemic project.

With this small caveat out of the way, let us consider an example of content luck:

POKER PLAYER: Jenny is an avid poker player. One day she participates in a high-stake tournament. As luck would have it, she is being dealt pocket aces for five consecutive hands. During the sixth hand, she looks at her cards. Aces again!⁵

Suppose Jenny believes that her hand consists of two aces for the sixth time in a row on the basis of both perceptual and memory faculties. Her belief is a case of content luck because the proposition believed (my hand consists of two aces for the sixth time in a row) is true but could have easily been false. Jenny could have easily been dealt a different hand. Also, in this case, the truth of this belief has clear practical significance for Jenny, so her belief satisfies our conditions for content luck.

This example makes it clear that content luck is a benign form of luck. For it seems clear that it is at least possible

⁴ This cursory glance leaves out some of the details of the exchange. It will suffice for our present purposes, however.

⁵ Unless otherwise stated, the examples in this chapter are my own.

for Jenny to come to know that she has been dealt two aces for the sixth time in a row by appropriately relying on memory and perception. That it is a matter of luck that the proposition she believes is true, thus does not rule out that her belief amounts to knowledge. Content luck is compatible with the possession of knowledge.

4.1.2 *Capacity Luck*

A second kind of epistemic luck that is compatible with knowledge is *capacity* luck. Sometimes, it is a matter of luck that we possess the capacities to acquire knowledge. In the literature on luck, it is widely acknowledged that this form of luck does not undermine knowledge. For example, Unger says about a turtle-watcher that while it may be “quite an accident [sic] that the turtle watcher is alive at the time he sees the turtle crawling on the ground before him ... at that time, as we have supposed, the turtle watcher knows that there is a turtle crawling there upon the ground” (Unger, 1968, p. 160). Similarly, drawing on the same example, Pritchard says that “genuine knowledge possession is not undermined merely by the fact that it is a matter of luck that the agent is in a position to know anything at all at that moment” (Pritchard, 2005a, p. 135).

Why do Unger and Pritchard think that capacity luck does not undermine knowledge? Take the following example:

LSD: Jimmy is dead-set on playing a devilish joke on one of his friends, Johnny. The joke consists in Jimmy slipping some LSD into Johnny’s drink while they are out in the cinema. Even though Jimmy prepared carefully for this joke, just before slipping the LSD into Johnny’s drink, Jimmy drops it into some spilled beer on the cinema-floor, rendering the drugs useless.

Jimmy and Johnny watch the film, which ends with a subtle Hitchcock reference.

Suppose Johnny comes to know that the film ended with a subtle Hitchcock reference through the usual perceptual and memory faculties. In this case, it is a matter of luck that Johnny possesses the capacity to know that the film ended with a subtle Hitchcock reference. For it is clear that Jimmy could have easily succeeded in drugging him, and it is equally clear that, if Johnny had not been able to rely on his perceptual and mnemonic faculties in the same way, the subtle reference would presumably have escaped him. We may also suppose that Johnny is an avid Hitchcock fan, such that possessing the capacity to notice that the film ended with a reference to his favourite director is of some significance to him. Thus, Johnny possessing the capacity to notice the reference is an event that satisfies our conditions for luck.

However, it is equally clear that this does not prevent Johnny from *knowing* that the film ended with a Hitchcock reference. Actually, Johnny *is* able to rely on his perceptual and memory faculties. These faculties are functioning as they should, and it seems that they should allow Johnny to acquire knowledge of the Hitchcock reference. The fact that Johnny could have easily been deprived of the reliable operation of these faculties does not detract from their ability to produce knowledge when they *are* operating reliably.

Some of the most extreme cases of capacity luck occur when people are almost killed. If I am almost hit by a car, it is a matter of luck that I have the capacity to then go on and form the memory belief that I was nearly hit by a car a moment ago. While it will be a matter of luck that I possess the relevant capacities, it is plausible that I can come to know this proposition on this basis. I conclude that capacity luck is a second benign form of epistemic luck.

4.1.3 *Evidential Luck*

Evidential luck is another benign form of luck (Engel, 1992a). Evidential luck can be defined in the following way:

EVIDENTIAL LUCK: *S*'s belief that *p* is evidentially lucky if and only if it is a matter of luck that *S* is in possession of the evidence she has for *p*.

This means that on our account of luck, a belief is evidentially lucky if and only if *S* could have easily failed to be in possession of the evidence she has for *p* (and the fact that *S* has the evidence she has for *p* is at least somewhat significant to her). Mylan Engel provides the following example of evidential luck:

One afternoon while working on a rather depressing novel in her study, Nadine looks out the window and sees rain pouring down, leaves whipping around, and intermittent flashes of lighting. Moreover, she hears rain slapping against the window, not to mention some rather loud thunderclaps. Nadine comes to believe that it is storming outside, and she is right. (Engel, 1992a, p. 69)

Does Nadine know that it is storming outside? As Engel rightly points out, Nadine's case seems to be a standard case of perceptual knowledge. Now suppose, however, that in the above case Nadine is only luckily at work in her study. Suppose that normally, Nadine works at a windowless, soundproof office at the other side of town, but because that building got struck by lightning yesterday, she decided to stay home. This means Nadine could have easily failed to see and hear the thunderstorm today. Thus, Nadine is only luckily in possession of the evidence

she has for her belief that it is storming outside. Nevertheless, Nadine acquires knowledge that it is storming outside on the basis of her perceptual faculties. If this is so, then knowledge is compatible with evidential luck.

Why is evidential luck compatible with knowledge? One suggestion would be that whether one knows that p does not depend how one acquired one's evidence for p . If this is so, then it follows that whether or not one's evidence is acquired in a lucky fashion is irrelevant to the question whether one possesses knowledge. This suggestion may seem radical, so it is worthwhile to pause on it a little. Consider the following example:

ENCYCLOPAEDIA: Jona is a superstitious person. One day in the park, Jona sees a broken twig that points downwards. Jona instantly starts digging, believing that this is a clear sign from the gods that there is a treasure buried underneath. As luck would have it, he finds an encyclopaedia buried underneath the broken twig. Curious, Jona starts reading. Two months later, he has finished the encyclopaedia, and it has made him a different man. As a result of reading the encyclopaedia, he has revised many of his superstitious beliefs.

Take any of Jona's revised beliefs. Suppose the evidence Jona has for this belief exclusively consists of the material provided by the encyclopaedia. Were it not for all his superstitious beliefs, Jona would have never started digging underneath the broken twig, so this evidence came into Jona's possession as the result of many epistemically dubious beliefs. But it seems we would still grant that Jona acquires knowledge when he revises many of his beliefs on the basis of his new evidence. The case thus serves to make plausible the claim that whether one possesses knowledge

that p does not depend on how one acquired one's evidence for p .⁶

Obviously, this does not imply that knowledge does not depend on the *quality* of one's evidence. In the above case, for example, we would deny Jona knowledge if the encyclopaedia he found were written by a madman, containing many falsehoods.⁷ But the claim that whether one knows that p depends on *what* evidence one has for p is of course compatible with the claim that one's knowledge does not depend on *how* one acquired that evidence.

Our claim here should also be distinguished from the following one. When Alvin Goldman distinguishes knowing that broccoli is healthy on the basis of having read it in the *New York Times* from merely believing that broccoli is healthy on the basis of reading it in the *National Inquirer*, he takes this to indicate that whether one knows that p depends on the causal history of one's belief (Goldman, 1999). However, Goldman's example features causal histories that are distinguished by the *quality* of evidence that they contain. A piece from the *New York Times* is simply *better* evidence than a piece from the *National Inquirer*. So the difference between knowing and not knowing in Goldman's example depends on a difference in the quality of evidence for one's belief. And that is compatible with our claim that it does not depend on how one came into possession of that evidence. It would be fully compatible with Goldman's story to say that one can come to know that broccoli is healthy on the basis of reading this in the *New*

6 A similar point is made in (Feldman & Conee, 1985, pp. 21–22).

7 What about if the encyclopaedia were written by a madman, but accidentally contained only truths? Arguably, this would constitute a mixture of evidential and content luck, say *evidential content luck*, for what is now a matter of luck is not necessarily that one is in possession of the evidence, but that the evidence has the content it has. I thank Jeanne Peijnenburg for bringing this to my attention.

York Times, even if this newspaper came into one's possession only luckily.

We can thus provide an *explanation* of the fact that evidential luck is compatible with knowledge by pointing to the fact that the way one's evidence comes into one's possession is irrelevant to the question whether one possesses knowledge.

In this section, we saw that content, capacity and evidential luck are benign forms of luck in the sense that they are compatible with the possession of knowledge. This concludes my discussion of the benign forms of luck. Let us continue to the forms of epistemic luck that are more problematic.

4.2 MALIGNANT FORMS OF EPISTEMIC LUCK

The forms of luck discussed in this section are called 'malignant' forms of epistemic luck because such luck seemingly precludes the agent from possessing knowledge. I will discuss two malignant forms of epistemic luck: veritic luck and reflective luck. While it is relatively uncontroversial that knowledge is incompatible with veritic luck, it is far more controversial to claim that knowledge is not compatible with reflective luck. I shall therefore devote Section 4.3 to arguing for this latter claim.

4.2.1 *Veritic Luck*

The first malignant kind of epistemic luck that we will discuss is called *veritic luck*. Roughly, a belief is veritically lucky if it is a matter of luck that the agent formed that true belief. Let us consider first some paradigmatic cases of veritic luck:

CASE 1: Janine forms the belief that the continuum hypothesis is true on the basis of a simple guess. Her belief is true.

CASE 2: Jill is driving through barn-façade county. She forms the belief that there is a barn in front of her. Her belief is true.⁸

CASE 3: Jeremy is a brain in a vat. On the basis of the stimuli evil scientists apply to his brain, he forms the belief that the sun is shining. His belief is true.

In all three cases, it is a matter of luck that the relevant agent forms the true belief. Formulated in this way, it may seem that veritic luck is quite close to content luck. Veritic luck is different from content luck, however, as Case 1 shows. For no belief in a necessary proposition can be content lucky (for necessary propositions cannot be false at all). As Engel notes, veritic luck is also distinct from evidential luck:

there is ... an epistemologically relevant difference between a person who is epistemically lucky in virtue of the fact that, given her evidential situation, it is simply a matter of luck that her belief turns out to be true, and a person who is epistemically lucky in virtue of the fact that she is lucky to be in the evidential situation she is in but that, given her evidential situation, it is not a matter of luck that her belief is true. I call the kind of epistemic luck had by the former "veritic luck" because it is just a matter of luck that her belief is true and the kind had by the latter "evidential luck" because it is just

⁸ See, for more on such cases, Section 4.2.2.

a matter of luck that she has the evidence she does. (Engel, 1992a, p. 67)

From this it seems that Engel adopts a definition of veritic luck along the following lines:

VERITIC LUCK₁: A belief is veritically lucky if and only if it is a matter of luck that one forms a true belief, given the evidential situation one is in.

On this interpretation of veritic luck, what is central to veritic luck is that it is easily possible that one is in the same evidential situation but forms a false belief. It is clear that the three cases provided above satisfy this definition of veritic luck. In all three cases, the agent could have easily formed a false belief, given their respective evidential situations.⁹

Because of the possible danger of confusing content luck with veritic luck, is very important to be clear on one's interpretation of veritic luck. Unfortunately, the formulations of veritic luck in the literature are often less clear than they could be in this respect. Consider, for example, one of Duncan Pritchard's accounts of veritic luck:

VERITIC LUCK₂: It is a matter of luck that the agent's belief is true. (Pritchard, 2005a, p. 146)

This specific formulation is equivalent to content luck. Moreover, it is not what Pritchard intended, since he does not want veritic luck to reduce to content luck. Suppose I form the belief that your lottery ticket has won by reading about it in the newspaper. This is a case of content luck, since surely, it is a matter of luck that you won. However, it is not a case of veritic luck, since not easily could I have

⁹ In case 2, this requires that one takes the evidential situation to be broader than just concerning the one real barn Jill is actually looking at.

failed to form this true belief *in this way*. What seems to be important for veritic luck is thus not whether it is a matter of luck that the proposition believed is true, but rather whether it is a matter of luck that the belief one formed is true, *given* the way it was formed.

This is a substantially different concept of veritic luck than in Pritchard's earlier work, however:

VERITIC LUCK₃: It is a matter of luck, given that the agent's belief meets all the relevant epistemic conditions, that the belief is true. (Pritchard, 2004, p. 204)

On this formulation, it must be a matter of luck that the belief is true, *given that all relevant epistemic conditions are satisfied*. This formulation is distinct from both Veritic Luck₁ and Veritic Luck₂. This is so because the relevant set of possibilities for Veritic Luck₃ includes only possibilities where the relevant epistemic conditions are met. While it is unclear what these relevant epistemic conditions would be in each case, it is clear that this could be a different set of possibilities than the set of possibilities relevant for either Veritic Luck₁ or Veritic Luck₂, which make no such restriction. For example, what is relevant for Veritic Luck₂ is how one formed one's belief, not whether one meets 'all the relevant epistemic conditions' in doing so. Because of this, the notions are not equivalent.

Yet another notion of veritic luck can be extracted from the work of Baumann, who rephrases Engel's characterization as follows:

VERITIC LUCK₄: It is a matter of luck that a given belief is true, given the evidence for it (Baumann, 2014, p. 526)

This notion is distinct from all foregoing notions, since the set of possibilities relevant for determining whether

a belief is veritically lucky now seems to consist of all and only those easy possibilities where the same belief is formed and where there *is* the same evidence for that belief. It is unclear how exactly Baumann intends to interpret this criterion, but he seems to think his formulation is equivalent to Veritic Luck₁. On the face of it, however, the existence of evidence for a belief does not depend on the believing subject's relation to that evidence. In this sense there can 'be' the same evidence for a belief, while the subject's evidential situation differs nonetheless. Again, this would mean that the two notions come apart.

In sum, contemporary literature on epistemic luck features many distinct notions of veritic luck. This deserves to be stressed, especially since on some of these accounts veritic luck reduces to content luck, and is benign rather than malignant. My own account of veritic luck is similar to what I take to be the intended meaning of Pritchard's (2005a) account:

VERITIC LUCK₅: It is a matter of luck that the agent's belief is true, given the way the agent formed her belief.

However, my account will not be identical to this formulation. For Veritic Luck₅ runs into a similar difficulty to the one discussed above: no belief in a necessary truth can ever be veritically lucky, according to this definition. For beliefs in necessary truths are such that they cannot be false, no matter what one's evidential situation is, or no matter how one's belief is formed. A belief in a necessary truth will therefore never satisfy Veritic Luck₅. Because of this, Veritic Luck₅ gives the wrong result in the above Case 1. There still is a sense in which it is a matter of luck that one forms a true belief if one forms a true belief in a necessary truth on the basis of random guessing. This luck is not different in kind from the luck in Case 2 and 3, and the

kind of luck that Pritchard and Engel try to target when they talk about veritic luck. In all these cases what seems problematic is the relation between one's evidential situation and the fact that one forms a true belief. That is the (problematic) relation captured by the concept of veritic luck.

The above considered definitions of veritic luck fail to adequately capture this problematic relation. What is common between luckily guessing a necessary truth, and luckily guessing a contingent truth, is that it is a matter of luck that one ends up with a true belief *at all*, given the way one's belief is formed. If we want to classify both cases as cases of veritic luck, the following definition thus suggests itself:

VERITIC LUCK₆: A belief is veritically lucky if and only if it is a matter of luck that the method one used to form one's belief produced a true belief.

The above formulation of veritic luck brings to the fore the importance of one's method of belief-formation for veritic luck. This introduces some complexity in our account. As Conee and Feldman have long argued, it may be hard to spell out in general how these belief-forming methods should be individuated (Conee & Feldman, 1998). This is a problem that I set aside here. It is a problem for many accounts of justification, since many such accounts refer either explicitly or implicitly to belief-forming methods in their definition of justification.¹⁰ It is not on me to solve this problem. I operate under the same assumption as do the proponents of these accounts of justification, namely that at least on a case by case basis, we can identify which belief-forming method is of epistemic relevance.

¹⁰ See Chapters 5, 6.

Veritic Luck₆ differs from Veritic Luck₅ in that lucky guesses of necessary truths will now count as veritically lucky, since in these cases one may have easily formed a *different* (false) belief instead. From now on, when I talk about veritic luck, I will mean Veritic Luck₆.

Why think that veritic luck is incompatible with knowledge? As Engel admits, arguing for the incompatibility of veritic luck and knowledge on deductive grounds will be difficult, since “any thing short of an a priori proof of the incompatibility of veritic luck and knowledge will fail to be conclusive, and the prospect of constructing such an a priori proof seems rather dismal.” (Engel, 1992a, pp. 67-68). We can, however, following Engel, find paradigmatic instances of knowledge-failure in which there is present a substantial amount of veritic luck. While falling short of a conclusive proof, such cases would support the claim that knowledge and veritic luck are incompatible.

As it turns out, perhaps *the* most paradigmatic instances of knowledge failure in the epistemological literature—*Gettier*-cases—are cases in which a large amount of veritic luck is present. Consider one of Gettier’s own cases (some-what abbreviated for ease of use):

DISJUNCTION: Smith has excellent evidence for the proposition that Jones owns a Ford, and forms the corresponding belief. From this proposition, Smith competently deduces the further proposition that either Jones owns a Ford, or Brown is in Barcelona, and again forms the corresponding belief. Smith has no evidence whatsoever that indicates that Brown is in fact in Barcelona, and so formulates the second disjunct quite at random. Now suppose that through some elaborate deception, all Smith’s evidence for believing that Jones owns a Ford is misleading, and Jones in fact does

not own a Ford at all. Suppose further, however, that Brown is in Barcelona at the moment Smith forms his belief in the disjunction. His belief thus ends up being true.

It is widely accepted that in the above case Smith does not *know* that either Jones owns a Ford or Brown is in Barcelona. However, note that Smith's belief is subject to veritic luck. For given the way Smith formed his belief (he formulated his second disjunct quite at random) it seems that this method could have easily produced a false belief. For example, Smith could have easily formed the false belief that either Jones owns a Ford or Brown is in *London* in the very same way he formed his actual belief. And so for many equivalent beliefs. This Gettier case thus clearly involves a substantial amount of veritic luck.

There are, of course, many more Gettier cases. To go through all of them would be both impossible and unnecessary, all the more so because Linda Zagzebski has provided a general formula for generating Gettier cases (Zagzebski, 1994). If we can show that, following this formula, one will be guaranteed to end up with a belief that is veritically lucky, this will suffice to show that all Gettier cases (at least of the standard sort covered by Zagzebski's formula) involve veritic luck. Since, as Gettier cases, none of these will involve knowledge, this will provide further support for the claim that knowledge and veritic luck are incompatible.

Linda Zagzebski argues that Gettier cases can be constructed for any theory of knowledge that consists exclusively of a non-factive epistemic condition in addition to the truth and belief condition (1994). That is, Zagzebski argues that any account of knowledge that consists of a truth condition, a belief condition, and a further (epistemic) condition that does not entail that the belief in question is true,

will be susceptible to Gettier cases.¹¹ Moreover, she argues that such cases can be constructed in the following way: take any non-factive epistemic condition you like and construct a case such that a given subject's true belief satisfies it. Then, modify the case such that accidentally, satisfying the epistemic condition does nothing to make it more likely that you end up with a true belief. Finally, make it so that as a second case of luck, you end up with a true belief nonetheless. In these cases, the subject will, according to Zagzebski, end up with a belief that satisfies the preferred conditions for knowledge, but will still fail to qualify as such. In short, the subject will end up with a Gettiered belief.

It is immediately apparent that in Zagzebski's recipe for constructing Gettier cases, luck plays an important role. As she herself comments, in Gettier cases "an accident of bad luck is cancelled out by an accident of good luck. The right goal is reached, but only by chance" (1994, p. 66). Here, the bad luck concerns the fact that while the evidence on which one bases one's belief is misleading, and the good luck the fact that the belief is true nonetheless. The question before us is then, whether the luck that features in the construction of Gettier cases is *veritic* luck. As we said above, on our conception of veritic luck, a belief is veritically lucky if one's belief-forming method actually produced a true belief, but could have easily produced a false belief instead. Now, given that one's evidence is misleading, it seems that one's method of belief-formation could have easily produced a false belief. The second kind of luck that Zagzebski talks about makes it such that while one could have easily formed a false belief, one ends up with a true belief nonetheless. In sum, cases that satisfy Zagzeb-

¹¹ We will argue against this claim in Chapter 6. Some non-factive accounts *can* evade Gettier cases. This point does not matter for the argument presented here. What matters for our present purposes is Zagzebski's 'recipe' for Gettier cases.

ski's criteria for being a Gettier-case are such that while the subject could have easily formed a false belief, they luckily end up with a true belief nonetheless. Thus, it seems that all Gettier cases, or at least the ones that can be constructed using Zagzebski's method, will feature veritic luck. Since Gettier cases are paradigmatic examples of the absence of knowledge, this provides support for the claim that knowledge and veritic luck are incompatible.

4.2.2 *Intervening and Environmental Luck*

We argued, based on Zagzebski's recipe for constructing such cases, that all Gettier cases feature veritic luck. But that is not quite right. As Pritchard argues, we can distinguish between Gettier cases that feature *intervening* luck and those that feature *environmental* luck (Pritchard, 2010; Carter & Pritchard, 2015). The former kind of luck, according to Carter and Pritchard, is "[t]he standard type of Gettier-style epistemic luck ..., and concerns epistemic luck which 'intervenes' between the agent's cognitive performance and her cognitive success (2015, p. 444). As an example they provide a Gettier case similar to Chisholm's famous sheep case: you form a belief that there is a sheep in the field on the basis of an object that looks like a sheep. This object is actually a stone, but behind the stone there is a real sheep. Thus, your belief is true, but only by luck, it seems. Particularly, Carter and Pritchard argue that in this case

we have cognitive success, in that the agent truly believes that there is a sheep in the field, and we also have cognitive performance, in that the agent is skilfully forming her belief that there is a sheep in the field. But the intervening epistemic luck in play means that the cognitive success is disconnected from the cognitive

performance that led to the agent forming this belief. (Carter & Pritchard, 2015, p. 444)

Intervening luck is a subspecies of veritic luck. For it is an essential feature of intervening luck that the cognitive success—the *truth* of the belief—is only luckily connected to the belief-forming method of the agent. Here I take the agent's 'cognitive performance' to be constituted by her belief-forming method. This means that in each case of intervening luck, we have a belief-forming method that could have easily produced a false belief, which will make each case of intervening luck a case of veritic luck.

Note that while each cognitive performance is constituted by a belief-forming method, it is not the case that every belief-forming method constitutes a cognitive performance. Only those belief-forming methods that constitute cognitive 'skills' on the part of the agent will be able to do so. While these cognitive performances will normally produce true beliefs reliably enough, they do not in the specific circumstances of these Gettier cases (Zagzebski's first case of bad luck). In fact, were it not for some extraneous factors (Zagzebski's second case of compensating good luck), they would reliably produce *false* beliefs in these cases.

Not all cases of veritic luck are cases of intervening luck, however. Consider a simple guess of mine that tomorrow it will rain. Suppose this method produces a true belief: I start believing that it will rain tomorrow, and this is true. However, it is clear that my method could have easily produced a false belief instead: if it is a truly uneducated guess, I could have just as easily believed that it would *not* rain tomorrow. So guesses like this one are cases of veritic luck. Yet they are not cases of intervening luck because here we cannot speak of any cognitive performance at all. It is not the case that in normal circumstances my method of belief-formation would reliably enough produce true be-

lief. Guessing is unreliable *tout court*. Because of this, my belief will be subject to veritic luck but not intervening luck.

All cases of intervening luck are thus cases of veritic luck, but not the other way around. Intervening luck is a subspecies of veritic luck. As a species of veritic luck, it is incompatible with knowledge.

Not all Gettier cases feature intervening luck, however. Some feature a kind of epistemic luck that is neither clearly malignant nor clearly benign. This kind of luck is called *environmental* luck. Roughly, environmental luck is the kind of luck that is in play in Gettier cases like the Goldman/Ginet *fake barn* cases. Consider Goldman's description of such a case:

Henry is driving in the countryside with his son. For the boy's edification Henry identifies various objects on the landscape as they come into view. "That's a cow," says Henry, "That's a tractor," "That's a silo," "That's a barn," etc. ... Suppose we are told that, unknown to Henry, the district he has just entered is full of papier-maché facsimiles of barns. These facsimiles look from the road exactly like barns, but are really just façades, without back walls or interiors, quite incapable of being used as barns. They are so cleverly constructed that travellers invariably mistake them for barns. Having just entered the district, Henry has not encountered any facsimiles; the object he sees is a genuine barn. (Goldman, 1976, pp. 772–773)

In this case, there clearly seems to be something lucky about Henry's belief that the object he is looking at is a barn. For one, he might have easily looked at one of the façades and formed the same belief, in which case that be-

lief would be false. On the other hand, as the case is constructed, Henry is looking at a real barn and forms a true belief about it on the basis of using his reliable perceptual faculties. If we strip away the environment, this seems like a case of normal perception. The kind of luck in play thus seems to concern the fact that the subject is lucky to be able to be in the epistemic situation she is in, given how the environment is structured. As Carter and Pritchard argue, in cases of environmental luck, the cognitive success *is* properly ascribable to the agent's cognitive performance. Rather, the luck in such cases concerns the fact that it is only a matter of luck that this is so, that the agent is able to achieve success through performance. Fake barn cases are lucky in this sense because in such cases it is only a matter of luck that the agent is able to form her true belief that there is a nice barn ahead on the basis of her cognitive skills.

The fact that fake barn cases are traditionally conceived of as Gettier cases indicates that environmental luck should be regarded as a malignant form of luck, incompatible with knowledge. Indeed, this is how Carter and Pritchard conceive of the case (2015, p. 445). But it is hard to find in the literature an argument for this claim other than an appeal to the intuition that in such cases no knowledge is present, an intuition which I must admit I do not share wholeheartedly.¹²

One reason for doubting that environmental luck is malignant is that the fake barn cases seem to share some relevant features with the cases of evidential luck described above. For in cases of environmental luck, subjects are quite lucky to be in the evidential situation they are in: quite easily could they have been in a different, misleading evidential situation. On the other hand, as is the case with

¹² This is not to say that I have an intuition that these cases *are* cases of knowledge; I simply do not feel intuitively inclined to judge one way or another.

evidential luck, it is *not* very lucky that *given* that they possess this evidence, they form a true belief. In Carter and Pritchard's terms, the cognitive success in these cases genuinely derives from the agent's cognitive performance. So one diagnosis of Henry's case above is that it is *not* very lucky that Henry forms a true belief, *given* that he is looking at a real barn. His belief-forming method would not easily have him believe otherwise than that there is a barn in front of him. These similarities between evidential and environmental luck seem to motivate a conception of environmental luck as a benign form of luck.

The fact that environmental luck is truth attained through competence also means that it is not a subspecies of veritic luck. For not easily could a false belief be attained through competence. If Henry looks at the real barn, and consequently forms the true belief that there is a real barn before him, he could not easily have formed a false belief *in this way*. Not easily would he have looked at that real barn but formed the false belief that there was no barn, or just a barn façade, for example. So given the way Henry's belief is *actually* formed, not easily would this method produce a false belief, even if his environment is such that he could have easily formed his belief in a *different* way, by looking at a barn façade, in which case his belief would not be success attained through competence. So where intervening luck is a subspecies of veritic luck, environmental luck is not. Thus we cannot derive the epistemically problematic nature of the latter from the problematic nature of veritic luck.

The above argument assumes that misleading evidence will constitute a different evidential situation than non-misleading evidence. One may deny this assumption. On this interpretation, Henry's evidential situation is the same if he is looking at a real barn and if he is looking at one of the façades. In that case, it *would* be a matter of luck that he formed a true belief, given the evidential situation he is in,

for given this situation, he could have easily formed a false belief instead. Similarly, what we said above assumed that Henry's belief-forming method is different when he looks at a real barn than when he looks at a barn façade. This means that we individuate belief-forming methods externalistically, with reference to features of the environment.

One may object to this way of individuating belief-forming methods. However, veritic luck depends on the belief-forming method the subject actually used, rather than on the method she believes she used. Even if Henry himself can make no distinction between how he forms his belief in the actual case from how he would form his belief were he to look at one of the façades, we, having access to facts that Henry does not have access to, can make a distinction, and the distinction is relevant for whether his belief is veritically lucky or not.

In any case, it is not clear whether environmental luck is compatible with *all* forms of knowledge. For example, while Carter and Pritchard argue that this kind of luck is incompatible with knowledge-*that* (propositional knowledge), they argue it is compatible with knowledge-*how*. Sosa, as we will see, also distinguishes different kinds of knowledge and argues that some of them are compatible with environmental luck but others are not (Sosa, 2007, p. 31).¹³ For now, we will remain impartial on the compatibility between environmental luck and knowledge.

To recap, all veritic luck, including intervening luck, is incompatible with knowledge. It is unclear whether environmental luck is incompatible with (all forms of) knowledge. In the remainder of this chapter, we will focus mainly on the distinction between veritic luck and the kind of luck to which we will now turn: *reflective luck*.

¹³ See Chapter 6.

4.2.3 *Reflective Luck*

The above forms of luck are not the only kinds of luck that are epistemically problematic. To see this, consider the following example from Robert Brandom:

Industrial chicken-sexers can, I am told, reliably sort hatchlings into males and females by inspecting them, without having the least idea how they do it. With enough training, they just catch on. In fact, as I hear the story, it has been established that although these experts uniformly believe that they make the discrimination visually, research has shown that the cues their discriminations actually depend upon are olfactory. (Brandom, 1998, p. 375)

When one of Brandom's chicken-sexers forms the belief that a particular chick is male, what is the epistemic status of this belief? On the one hand, it seems that the belief will not be veritically lucky. Given the reliability of the way their belief is produced, they could not have easily formed a false belief.

However, on the other hand, there still seems to be a kind of luck in play that is epistemically problematic. To bring this out, suppose these chicken-sexers have no idea how they form their beliefs, nor whether that method is reliable. That is, we are supposing here that the sexers do not have available to them any track record information.¹⁴ On this assumption, then, there is an important sense in which, *from their perspective*, they could have easily formed a false belief instead of a true one. Note that the issue is not that these chicken-sexers doubt the belief they actually

¹⁴ The chicken-sexer case is used here as a thought-experiment. As such, the situation is highly idealized. I do not want to commit myself to the claim that this is how chicken-sexing actually works.

formed, but rather whether, from their perspective, they could have easily formed a *different*, false belief instead. So these chicken-sexers need not consider their beliefs subject to content luck. For example, they may be strongly convinced that a particular chick is male. The point is that even if they are somehow deeply convinced that their belief is true, the fact that they have no idea whether this belief is produced in a reliable way means that they have to consider it an easy possibility that they would have formed a different, false belief instead. From their subjective perspective, it is thus a matter of luck that their method of belief-formation produced a true belief.

This kind of luck is known in the literature under the heading of *reflective luck*, a term introduced by Duncan Pritchard (Pritchard, 2005a). Pritchard characterizes reflective luck in the following way:

REFLECTIVE LUCK₁: Given only what the agent is able to know by reflection alone, it is a matter of luck that her belief is true. (Pritchard, 2005a, p. 175)

This definition, however, should be rejected for the same reason as Pritchard's definition of veritic luck. For it is possible that some agent would have the reliable ability to tell mathematical truths from mathematical falsehoods, but who, like our naive chicken-sexers above, had no idea how she did it. Her belief would still be lucky in the same way as that of our chicken-sexer. What matters is not that the particular belief *itself* could easily have been false, but whether, from the subject's perspective, the method used to produce the belief could have easily produced a false belief instead. We thus arrive at the following definition of reflective luck:

REFLECTIVE LUCK₂: S's belief that *p* is reflectively lucky if and only if, given the informa-

tion reflectively accessible to S, it is a matter of luck that the method S used to form her belief that p produced a true belief.

Note that Reflective Luck₂ is essentially an agent-relative version of Veritic Luck₆; the only difference between the two is that the luckiness of the formation of a true belief is made relative to the information reflectively accessible to the believing subject in Reflective Luck₂, whereas it is not made relative to any restricted set of information in Veritic Luck₆.

What does it mean to determine the luckiness of the production of a true belief with respect to the information reflectively accessible to a believing subject, as is required by Reflective Luck₂? Plugging in our modal account of luck, we need to determine whether from the subject's perspective, her belief-forming method could have easily produced a false belief. If this is so, then her belief is subject to reflective luck. As Pritchard notes, this has two notable consequences (Pritchard, 2005a, sec. 6.5).

First, the fact that we are judging luckiness from a subjective point of view means that what counts as an easy possibility is determined by what, from the subjective viewpoint, could have easily occurred. This may result in a different ordering of possibilities, since what is an easy possibility from the perspective of a believing subject may not be so objectively speaking. For example, if the subject believes that airplanes are very dangerous, then it may well be that from her perspective, airplanes could easily crash, even if in fact this possibility is quite remote. In this case, if we want to determine whether, from *the subject's* perspective, it is a matter of luck that the plane did not crash, we need to see whether *she* considers this to be easily possible. In this case, she does, and so it will be a matter of luck, from her perspective, if the plane does not crash. In a similar fashion, when determining whether a belief is

reflectively lucky, we need to determine whether from the subject's perspective, her method of belief-formation could have easily produced a false belief.

Second, for the same reason we need to identify the relevant belief-forming method as the method that the subject believes she used, rather than the one she actually used. Given that our chicken-sexer has no idea how she formed her belief, it is a matter of luck, from her perspective, that she formed a true belief. But that is only so because she does not know what the actual method was that produced her belief. If she took into account the actual method that produced her belief, involving olfactory stimuli, she would probably *not* consider it an easy possibility that she would have formed a false belief. If we thus want to capture the sense in which the belief is lucky from the subject's point of view, we need to take into account the method the subject believed she used rather than the method she actually used.

It is clear from the above that veritic and reflective luck are two substantially different forms of epistemic luck. It is therefore perhaps unsurprising that the extensions of the two notions will diverge: there will be cases of veritic luck that are not cases of reflective luck and vice versa. We will provide an example of each. Our first case is BonJour's Norman, whom we have already encountered in Chapter 2:

NORMAN: Norman has the mysterious condition that from time to time, he forms beliefs about the whereabouts of the president of the United States. These ideas seem to be formed spontaneously; they just 'pop' into Norman's head. As it happens, Norman has never been able to check whether his beliefs are correct. As a matter of fact, however, Norman is the one and only human being with a reliable faculty

of clairvoyance concerning the whereabouts of the president of the United States.¹⁵

Since Norman's beliefs are the result of a reliable faculty of clairvoyance, they will not be veritically lucky.¹⁶ For not easily would his faculty of clairvoyance produce false beliefs. However, Norman's beliefs are subject to a substantial degree of reflective luck: given that Norman has no idea that he is in possession of such a reliable faculty of clairvoyance, and given that he has no further evidence that supports his beliefs about the president's whereabouts, it seems that *from his perspective*, he could have very easily formed a false belief instead. So Norman presents a case of reflective luck without veritic luck.¹⁷ Now consider the following case:

NOLA: Suppose Nola has excellent reasons for most of her beliefs.¹⁸ She behaves in an epistemically conscious way, always double-checking her evidence, and making sure she proportions her beliefs to the evidence as much as possible. Unfortunately, Nola is plugged into the Matrix, a very advanced computer simulation, meaning that she rarely forms a true belief at all.

Suppose Nola comes to believe that she is looking at an elm tree, and infers from this that there are elms nearby. Suppose further that her belief is true: as a matter of fact, there are elms nearby. It seems that, from Nola's perspective, the method she used to form this belief could not have

¹⁵ This case is a slightly modified version of BonJour's famous case (BonJour, 1980).

¹⁶ We will defend this claim in more detail in Chapter 6.

¹⁷ Note that Norman's case is structurally similar to the case of our naive chicken-sexers.

¹⁸ This case will presuppose that one's reasons can be excellent even while false. If one does not agree, modify the case such that Nola *thinks* she has excellent reasons for her beliefs.

easily produced a false belief. This means that Nola's belief is not reflectively lucky. However, since the way she actually formed her belief (on the basis of artificial stimuli, say) *could* have easily produced false belief, her belief is veritically lucky nonetheless. So Nola's case presents a case of veritic luck without reflective luck.¹⁹

Now that we have seen that reflective and veritic luck have different extensions, the question becomes relevant whether reflective luck is compatible with knowledge. That is, whether we are right to classify reflective luck as an epistemically *malignant* form of luck. For while there is clearly something epistemically problematic about the beliefs of our chicken-sexers, or Norman the clairvoyant (at the very least, we would rather be in *our* epistemic position than in theirs), it is not thereby said that what is problematic is that these people do not possess *knowledge*. The question whether reflectively lucky beliefs can constitute knowledge is therefore taken up in the next section.

4.3 REFLECTIVELY LUCKY KNOWLEDGE

The question before us is whether knowledge is compatible with reflective luck. In this section I borrow Sosa's well-known distinction between animal and reflective knowledge, and argue that only one of them is compatible with reflective luck.²⁰

¹⁹ Nola's case is a special case of what are known as 'New Evil Demon-arguments' See (Lehrer & Cohen, 1983; Dutant & Dorsch, n.d.). Such cases, like Norman's case, are often put forward as an argument for internalism. See also the next chapter.

²⁰ Cf. (Sosa, 2009a, 2010, 2015). Note that Sosa makes the further distinction between reflective knowledge and knowledge 'full well'. Everything we say in this section about the relation between reflective luck and reflective knowledge can be said about the relation between reflective luck and knowledge full well as well, since knowledge full well entails reflective knowledge, according to Sosa (2015, p. 74).

What is required for animal knowledge? According to Sosa, animal knowledge is apt belief, where apt belief is understood as “belief whose correctness is attained sufficiently through the believer’s epistemic competence”, where an epistemic competence is a competence to believe truly (Sosa, 2015, p. 9). Thus, I may possess animal knowledge whenever I form a belief that is true, and whose truth is attained through a competence of mine to believe truly. These competences come in many different forms, and may reside at the personal level, as well as at the sub-personal level. For example, I may possess the epistemic competence to form true beliefs about medium-sized objects at close distances using my visual system. Another epistemic competence that I may possess is the competence to form true beliefs about the answers to crosswords-puzzles using a combination of my memory and reasoning faculties.

The above examples bring out that the term ‘animal’ knowledge is perhaps unfortunate, for such knowledge may be acquired in ways that are unavailable to most animals, as the crossword-puzzle case brings out. Nevertheless, it seems that we share at least *some* animal knowledge with other animals, as the perceptual case brings out.

Animal knowledge is distinguished from reflective knowledge in the following way:

[a]nimal knowledge does not require that the knower have an epistemic perspective on his belief, a perspective from which he endorses the source of that belief, from which he can see that source as reliably truth conducive. Reflective knowledge does by contrast require such a perspective. (Sosa, 2009c, p. 135)

The idea seems to be this: we can distinguish the kind of knowledge that we may achieve through (extensive) reflection from the kind of knowledge that more automatic

and that non-human animals may perhaps also achieve. Whereas it is a general property of our reflective knowledge that it is embedded in an 'epistemological perspective', a perspective from which the sources of our beliefs are deemed reliable, the same cannot be said for the kind of knowledge that we could perhaps be said to share with some non-human animals. The distinction is such that there is a necessary condition for reflective knowledge (being embedded within a particular epistemic perspective) which animal knowledge need not satisfy.

For Sosa, reflective knowledge is a higher grade of knowledge than animal knowledge, both more valuable and more difficult to achieve, because it requires animal knowledge *plus* an epistemic perspective from which the believer endorses her belief.²¹ However, it seems possible to endorse Sosa's distinction between animal and reflective knowledge without committing to his specific framework for analysing the two notions. The main point of the distinction is that there seems to be a grade of knowledge that does, whereas there also seems to be a kind of knowledge that does not, require a perspective from which the sources of the relevant beliefs are reliable. In Chapter 6, we will discuss Sosa's virtue theoretical framework for analysing the notions of animal and reflective knowledge. Here we take a more neutral approach, and argue that irrespective of how one fleshes out these notions, it is plausible that animal knowledge is compatible with reflective luck but reflective knowledge is not.

Let us first take a closer look at the extent of animal knowledge. First, obviously, the notion targets the kind of

²¹ The term could be read as pejorative. This is certainly not what is intended. Humans are animals too, and a lot, if not most of our knowledge is of the animal kind. We simply rarely have the time and resources available to achieve reflective knowledge. So often animal knowledge will have more pragmatic value. But from the epistemic point of view, reflective knowledge does seem to be more valuable.

knowledge that we attribute when we say things like “the dog knows that his owner is at the other side of the door”, or “the elephant knows that there is water nearby”. But animal knowledge can also be had by human beings. For example, we sometimes say that “the baby knows that his mother smiles at him”. It is clear that babies do not have an epistemic perspective on the reliability of the sources of their beliefs. Therefore, the knowledge that we attribute here seems to be of the animal, rather than the reflective, kind.

Our chicken-sexers may also be said to fall within the extension of animal knowledge. For they too lack an epistemic perspective on the reliability of the sources of their beliefs. As a contrast case, we may imagine our ‘naive’ chicken sexers have more enlightened counterpart chicken-sexers, who have read about the relevance and reliability of olfactory cues in making their judgements. Only the latter, it seems, could plausibly be said to have an epistemic perspective from which the sources of their beliefs are reliable. If there are circumstances in which even our naive chicken-sexers can be said to know at all, the knowledge attributed will thus be of the animal kind.

More inclusive still, perhaps some (or much!) of our own, ordinary knowledge is of the animal kind. For there are many occasions, it seems, where one trusts one’s sources of belief without considering their reliability, and just takes the beliefs produced at face value. Both the above visual belief case and crossword-puzzle case can serve as examples here. Without the presence of an epistemic perspective from which our sources are reliable, such beliefs could at most constitute animal knowledge.

Thus, perhaps only a small subset of our knowledge constitutes reflective knowledge. Perhaps there is only a select subset of my beliefs for which I can claim that my epistemic perspective supports the reliability of their sources. In fact, when faced with sceptical scenarios, we may won-

der whether *any* of our beliefs satisfy the criteria for reflective knowledge. That is, we may wonder whether any of our beliefs are such that our epistemic perspective *adequately* supports the reliability of their sources.²²

Prima facie, however, we all take ourselves to have at least some such reflective knowledge. And clearly, we value this kind of knowledge over 'mere' animal knowledge, at least epistemically speaking. I take it that most of us would place greater epistemic value in the position of the enlightened chicken-sexer than in that of her naive counterpart, who, in an epistemic sense, is no better than a reliable thermometer.²³ It is not the case that we just think that the former knows differently than the latter, we think the former knows *better*.²⁴

With this distinction between human and animal knowledge in place, we are in a position to investigate the relation between these grades of knowledge and reflective luck. First, we will argue that animal knowledge is compatible with reflective luck. All we need to do is provide a case of animal knowledge that is subject to reflective luck. We have already identified two of such cases.

The first case is the paradigm case of animal knowledge: knowledge that some non-human animals have. Take the belief of my dog Jupiter that it is me, his owner, on the other side of the door. Suppose this belief constitutes animal knowledge. Since all knowledge entails truth, the fact that Jupiter knows entails that Jupiter's belief is true. But presumably, Jupiter does not have any beliefs about the reliability of the sources of his belief. He may even be incapable of having such beliefs. This means that there is

²² We will come back to this issue in Chapter 7.

²³ The 'thermometer' model of knowledge was developed by Armstrong as an account of the knowledge constituted by our most basic perceptual beliefs (Armstrong, 1973).

²⁴ Of course, one may want to deny that the latter knows at all. That is not the route taken in this chapter.

nothing in Jupiter's perspective that rules out the possibility that the source of his beliefs is unreliable. If this source is in fact unreliable, then it could have easily produced a false belief. From Jupiter's perspective, there is nothing that rules out that he could have easily formed a false belief. Thus, while Jupiter's belief may constitute animal knowledge, it will necessarily be reflectively lucky.

One may object that the above reads too much into the epistemic perspective of Jupiter. Indeed, it may be contested that animals possess knowledge at all, even of the 'animal' variety. In this case the terminology would be very unfortunate, but there may still be other cases of 'animal' knowledge. Perhaps a better example of such animal knowledge would then be provided by our naive chicken-sexer. She too, as we saw, could perhaps possess animal knowledge. Suppose one of them possesses such knowledge of the proposition that the hatchling before her is male. Since, again, this means that her epistemic perspective does not favour her belief being reliably produced over it being *unreliably* produced, this means that from her epistemic perspective, it is an open possibility that the source of her belief could have easily produced a false belief. Therefore, from her epistemic point of view it is a matter of luck that her method of belief-formation turned out a true belief: her belief is reflectively lucky.

The above two cases indicate that animal knowledge and reflective luck are compatible. The above two cases suggest, moreover, that we can even go further and show that the presence of animal knowledge *entails* that the relevant belief is reflectively lucky. This is because the fact that both Jupiter's belief and our naive chicken-sexer's belief were found to be reflectively lucky seemed to *derive* from the fact that both of them lacked an epistemic perspective from which the source of their beliefs was reliable. Because the lack of such a perspective is essential to animal knowledge,

we can conclude that all animal knowledge will be reflectively lucky.

Thus not only is animal knowledge compatible with reflective luck; it entails it. The converse does not hold, of course, since there may be conditions for animal knowledge that are not satisfied by every reflectively lucky belief.

While animal knowledge is compatible with reflective luck, the same does not hold for *veritic luck* and animal knowledge. If our arguments from the previous section are sound, *no* knowledge is compatible with veritic luck, and that, of course, includes animal knowledge. For the same reason reflective knowledge will not be compatible with veritic luck. In the remaining part of this section, I will argue that in contrast to animal knowledge, however, reflective knowledge is incompatible with reflective luck as well.

First, we focus on paradigm cases of reflective knowledge and show that in these cases, the relevant beliefs are not reflectively lucky. I take our scientific knowledge to constitute a paradigm case of reflective knowledge, as well as the case of the enlightened chicken-sexer. Let us take these cases in order. First, when we consider our collective scientific beliefs, they do seem to be supported by an epistemic perspective from which the sources of these beliefs are reliable. Scientists typically have good reasons for their beliefs, reasons which are scrutinised by the scientific community. Scientific methods, one may argue, are among the most reliable methods we have at our disposal for forming true beliefs. Even if these methods are too involved to use in most everyday contexts, when we *really* want to know whether something is likely to be true, we often turn to science. Moreover, the reliability of scientific method is itself something that is intensively studied, within the sciences themselves.²⁵ All in all, it thus seems that if *any* of

²⁵ And also from an 'outside' perspective in philosophy of science.

our knowledge would constitute reflective knowledge, it would be our scientific knowledge. And it is equally clear that given our collective epistemic perspective on the reliability of our scientific methods, it is *not* the case that given this perspective, our scientific methods could have *easily* produced false beliefs. Thus, our scientific beliefs will not be reflectively lucky.

To a lesser extent, neither will the beliefs of the enlightened chicken-sexer. She has read about the fact that chicken-sexers like herself tend to discriminate the chicks on the basis of smell, and after some training are able to do so quite reliably. She knows that she has received training, and thus, that she should be able to sort the chicks reliably. Thus, from her perspective, it seems, she could not easily form false beliefs. Her beliefs will thus not be reflectively lucky.

Reflective luck is thus absent in paradigm cases of reflective knowledge. We can provide a more general argument for the conclusion that reflective knowledge is incompatible with reflective luck when we consider Sosa's specific interpretation of reflective knowledge.

As Sosa argues, reflective knowledge is apt belief (animal knowledge), *aptly noted*. Reflective knowledge thus requires a second-order belief that aptly indicates the first-order belief is apt. Above we have seen some examples. Another example of reflective knowledge may be the following. I believe there is a dog in my neighbour's yard. My belief is true because it results from the competent exercise of my auditory faculties, and it is accompanied by a second-order belief that this is so, a second-order belief that itself manifests one of my epistemic competences.

It is relatively easy to see that such knowledge requires the elimination of reflective luck. For contrary to animal knowledge, reflective knowledge of p requires that the subject believes that she formed an apt belief that p . That is, it requires that the subject has a second-order belief q , that

her first-order belief p constitutes animal knowledge. And as we have seen, that entails that her belief that p is not veritically lucky, i.e. that she could not have easily formed a false belief instead of p . Since what a subject's second-order beliefs are is part of what is reflectively accessible to her this means that given the information reflectively accessible to her, she could not have easily formed a false belief instead of p , and thus, that her belief that p is not reflectively lucky.

The above suffices to at least make plausible the claim that reflective knowledge is incompatible with reflective luck. The primary result of this section is thus that while animal knowledge is compatible with reflective luck, reflective knowledge is not.²⁶ By distinguishing animal from reflective knowledge, we thus see that reflective luck is a *malignant* kind of epistemic luck only insofar as the higher-order kind of knowledge is concerned.

4.4 CONCLUDING REMARKS

In this chapter, I discussed various 'benign' forms of epistemic luck: content luck, capacity luck and evidential luck. I discussed two potentially *malignant* kinds of luck: veritic luck and reflective luck. Only the first of those is incompatible with *any* kind of knowledge, for we saw in the last section of this chapter that animal knowledge is compatible with reflective luck. We also briefly discussed the two forms of luck present in Gettier cases: intervening and environmental luck. We argued that only the first of these is a subspecies of veritic luck and thus clearly problematic.

The upshot of this chapter is that different forms of luck are compatible with different grades of knowledge. This will prove important later on. For we will see in the next

²⁶ I argue for the same claim in my paper "Epistemic justification and epistemic luck" (de Grefte, 2017).

chapters that various theories of justification are incompatible with different of these forms of luck. Indeed, we will see that the line between *internalist* and *externalist* theories of justification can be drawn in terms of the kinds of epistemic luck these theories are incompatible with. This will open up the way for arguing that these distinct concepts of justification are relevant for different grades of knowledge.

INTERNALISM AND LUCK

In the previous chapter I argued that reflective knowledge requires the elimination of both reflective and veritic luck. I argued that animal knowledge requires only the elimination of veritic luck. It seems thus, that different kinds of knowledge can be individuated in terms of their anti-luck conditions.

In this chapter, we will focus our attention on justification rather than knowledge. As we saw in Chapter 2, one of the reasons Plato proposed a necessary justification condition on knowledge was to exclude (certain forms of) luck. In this sense, one can think of epistemic justification as an anti-luck condition on knowledge. And this is indeed how many people still think of justification:

“... one role of the justification is to rule out lucky guesses as cases of knowledge.”
(Ichikawa & Steup, 2017)

“... in order to rule-out lucky ‘knowledge’, we think that we need a further condition: the justification condition. (Booth, 2011, p. 38)

“Fallibilistic justification is thought to rule out epistemic luck by making one’s belief extremely probable.” (Engel, 2015)

One problem with this way of thinking about justification, however, is that at least according to the standard story, Gettier showed that justification is not up to this task: it fails to rule out certain problematic forms of epistemic luck:

... justification and knowledge must somehow not depend on coincidence or luck. This was just the point of the Gettier counter-examples; nothing in the tripartite definition excluded knowledge by luck.¹" (Dancy, 1985, p. 134)

There are many different concepts of epistemic justification, however, and, as we have seen in the previous chapter, many different forms of epistemic luck. The question is whether all concepts of justification are compatible or incompatible with luck in the same way. It is this question that we will take up in the present chapter and the next. As it will turn out, we can draw a line between internalist and externalist concepts of justification in terms of different kinds of epistemic luck these concepts rule out. Specifically, I argue in this chapter that internalist conceptions of justification are incompatible with reflective luck, but compatible with veritic luck. In the next chapter I shall argue that the converse holds for externalist conceptions of justification. Viewed in this light, part of the disagreement between internalists and externalists is thus about the kinds of luck that justification is meant to rule out. This will lead to a novel interpretation of the internalism/externalism-debate.

The upshot of this investigation will be a picture quite different from the one sketched above. What Gettier showed is that *internalist* justification does not eliminate veritic luck. But that should have been clear from the start: as I argue in this chapter, none of the prominent internalist forms of justification are directed at eliminating veritic luck. Rather, the kind of luck that is problematic for the internalist is *reflective* luck. And this kind of luck is indeed absent in Gettier cases. On this picture, the upshot of Gettier's paper is *not* that justification was not up to the task it was meant to do. Rather, Gettier showed that there is

¹ See also (Booth, 2011)

another task that needs to be done if we want to achieve knowledge—eliminating *veritic* luck. In the next chapter I argue that externalist theories of justification are particularly well-suited for this task.

Before we continue, a few remarks are called for. First, I will speak of concepts of justification being compatible or incompatible with certain kinds of luck. By this I just mean that when a belief satisfies the conditions for that particular concept of justification, it cannot be subject to the kind of luck in question. When I say that a particular concept of justification *eliminates* a certain kind of luck, I mean the same thing.

Second, as we have seen, both luck and justification are gradual notions: one can be more or less lucky, more or less justified. When I say in this chapter and the next that a particular concept of justification is incompatible with a specific kind of luck, this does not mean that every degree of that kind of justification is incompatible with every degree of that kind of luck, no matter how low. Rather, the claims defended in this chapter and the next ones should be read as (implicit) claims about relative proportions: the more of a particular kind of justification, the less of the relevant kind of luck. At the end-point of the scale, the *maximal* degree of justification of a particular type may require the complete absence of a particular kind of luck. Less than complete justification will allow for more of the relevant luck.

That said, most theories of justification are formulated as if justification were a discrete and absolute notion: either one is justified or one is not justified. Thus, such theories include necessary and sufficient conditions for justification *in general*. Usually, however, these theories allow for their conditions to be satisfied to a fuller or lesser degree, thus allowing for different degrees of justification.

Third, as before, I will restrict the discussion in this chapter to the two internalist approaches to justification that we

already came across in Chapter 2, the deontological/accessibilist approach, and the evidentialist/mentalist approach. These two approaches of course do not exhaust the possible internalist views of justification. Nevertheless, they represent by far the most dominant internalist theories of justification in contemporary literature, and we may thus hope that our findings can be generalized. The cautious reader, however, should read my claims concerning internalism as qualified to these two—important—internalist theories only.

With these preliminaries out of the way, I present the structure of this chapter. In Section 5.1, I argue that a deontological/accessibilist theory of epistemic justification entails that justified beliefs cannot be reflectively lucky, but that it allows for justified beliefs that are subject to a substantial degree of veritic luck. In Section 5.2 I do the same for the mentalist/evidentialist account of epistemic justification. Section 5.3 contains some concluding remarks.

5.1 ACCESSIBILISM

In this section I argue that historically one of the most prominent forms of internalism about justification—accessibilism—eliminates reflective luck but not veritic luck.² Accessibilism can be captured, as we saw in Chapter 2, by the following thesis:

ACCESSIBILISM: S is justified in believing *p* only
if *p*'s justifiers are reflectively accessible to S.

² As we saw in Chapter 2, forms of accessibilism have been endorsed by such eminent epistemologists as Bonjour (1980), Chisholm (1977) and Ginet (1975), among many others. Nowadays, its popularity seems to be in decline, either in favour of externalism or in favour of the other kind of internalism to be discussed in this chapter, mentalism (see Sections 2.3 and 5.2).

Accessibilism at heart is thus a supervenience thesis: it entails that justification supervenes on what is reflectively accessible to the relevant subject. As we saw in the previous chapter, reflective luck, too, supervenes on the information reflectively accessible to an agent. Why is this so? Consider Jim and Jane, two reflective duplicates. That is, Jim and Jane have reflective access to exactly the same information; everything that is reflectively accessible to Jim is similarly accessible to Jane and *vice versa*. Jim and Jane both believe that *p*. Now suppose, for reductio, that Jane's belief is reflectively lucky, but Jim's is not. In the previous chapter we said that a belief is reflectively lucky if and only if it is an easy possibility, from the perspective of the agent, for her method of belief-formation to produce a false belief (where the perspective of the agent is constituted by the information reflectively accessible to her). According to this definition, this means that given what is reflectively accessible to Jane, she could have easily formed a false belief, whereas given the information reflectively accessible to Jim, he could *not* have easily formed a false belief.

However, as we also said in the previous chapter, what is easily possible from one's perspective is determined by the information that one has reflective access to. If one believes that flying is dangerous, then a plane crash will generally be an easy possibility from one's own perspective. If reflectively accessible information determines one's subjective possibility-ordering, however, then it is impossible that Jim and Jane have different such orderings. Their subjective possibility-orderings will be identical because they are determined by the same set of reflectively accessible information. What this means is that it cannot be the case that it is an easy possibility that her method of belief-formation would have produced a false belief *for Jane*, but not *for Jim*. What is easily possible from Jane's perspective must be so from Jim's perspective as well, since their perspectives are identical. Hence, it is impossible that Jane's belief is sub-

ject to a different degree of reflective luck than Jim's belief is. Reflective luck, like accessibilist justification, supervenes on what is reflectively accessible.

A shared supervenience base does not entail, however, that accessibilist justification cannot be reflectively lucky. As it stands, accessibilism is nothing *but* a claim about the supervenience basis of justification: it states that justification depends on what is reflectively accessible, but not *how* it depends on this information. In principle it would thus be possible for the proponent of accessibilism to adopt a concept of justification that is compatible with reflective luck.

An example of such an accessibilist theory would be to say that one's belief is justified if and only if it is supported by at least one other belief. Under the assumption that all beliefs are reflectively accessible, this kind of justification supervenes on what is reflectively accessible. But one's belief may be justified in this sense yet reflectively lucky. For it is possible that one's belief is supported by another belief, even if it seems from one's perspective that one's method of belief-formation could have easily produced a false belief. An example would be to guess a correct answer, an answer that one fails to see is supported by some other beliefs one has.

Historically, however, accessibilists have tended to endorse a concept of justification that *is not* compatible with reflective luck. Or so I will argue. In Chapter 2, we saw that one of the main historical motivations for accessibilism draws on the idea that epistemic justification is deontological in nature. I shall argue in what follows that deontological justification entails an absence of reflective luck. This will show that accessibilist justification, at least as it is historically motivated, is incompatible with reflective luck as well.³

³ Even if not all possible forms of accessibilism entail the absence of reflective luck, the fact that (versions of) this form of accessibilism has been

Why think that accessibilism combined with a deontological concept of justification eliminates reflective luck? Let us first briefly rehearse the deontological concept of justification. As we saw in Chapter 2, it can be formulated in the following way:

DJ: S is justified in believing p if and only if S does not violate any epistemic norms in believing p .

If justification is a matter of not violating any epistemic norms, as proponents of the deontological concept presuppose, then to evaluate whether someone is justified, we need to know what these epistemic norms are.

Presumably, there are many epistemic norms. For example, it seems plausible that, epistemically speaking, we should do such things as proportion our belief to our evidence, believe what is true, strive for understanding, refrain from believing what is false, strive for coherent belief sets, trust our experiences, and so on. Reductionists about these norms think that all these norms can be reduced to one ultimate norm. Anti-reductionists deny this.⁴ For our purposes, the distinction does not matter: if justification requires not violating *any* norm, then any plausible norm incompatible with reflective luck will suffice to show that justification, on the present conception, is incompatible with reflective luck.

As we have already seen in Chapter 2, a common way to formulate such an ultimate epistemic norm is the following: "One's primary epistemic duty ... includes both

endorsed by textbook internalists like Bonjour, Chisholm and Ginet, means that the claim that it is incompatible with reflective luck should thus have some interest in its own right. See (Bonjour, 1980; Chisholm, 1977; Ginet, 1975).

⁴ Marian David is an example of the former, Jonathan Kvanvig of the latter (cf. David, 2013; Kvanvig, 2013).

seeking the truth and avoiding error" (BonJour, 2002, p. 236).⁵

At first sight, it may seem that BonJour specifies an *externalist* condition on justification. After all, whether one believes in a way that maximizes true belief and minimizes false belief depends on factors, most notably *truth*, that supervene neither on what is reflectively accessible, nor on one's mental states. Of course, insofar as the deontological concept of justification is to be used in support of *accessibilism*, it cannot be interpreted in such an externalist way. That is why proponents of accessibilism tend to interpret the epistemic norm in a more 'subjective' way: such that it does not require the subject to *actually* believe in a way that maximizes truth and minimizes falsity, but requires the subject does so to the best of her knowledge.⁶

Thus William Alston—himself no supporter of the deontological concept of justification—writes:

[I]f our basic intellectual obligation is to maximize truth and minimize falsity, one cannot be deontologically justified in a belief unless one is believing in such a way that *so far as one can*

5 For similar formulations, see (Alston, 1989, p. 116; BonJour, 1985, p.7-8; Chisholm, 1977, pp. 12–15; David, 2001).

6 Such an interpretation seems to fit with the distinction between *evaluative* and *prescriptive* norms (McHugh, 2012). Roughly, where evaluative norms specify what is good or bad, prescriptive norms specify what one ought to *do*. These norms do not always imply each other: "Something may be bad without its badness being a matter of anyone's having done anything they ought not have done, and without its being the case that there is anyone who ought to change it; some prospective state of affairs or object may be good without its being the case that there is anyone who ought to produce it or bring it about. Evaluations do not presuppose accountability or blameworthiness" (McHugh, 2012, p. 10). The externalist norm would be the evaluative epistemic norm, and the 'subjectivized' norm the derivative prescriptive norm. Simion, Kelp, and Ghijzen argue that satisfying a prescriptive rather than an evaluative norm may be sufficient for epistemic justification (Simion et al., 2016).

tell, is well calculated to reach the truth. (Alston, 1989, p. 201, my emphasis)

For Alston's deontologist, then, justification requires not that one *actually* believes in such a way as to maximize true belief and minimize false belief, but rather that one does so *as far as one can tell*. For example, one could be deontologically justified if one has done one's best to acquire as much evidence as possible, and to weigh the evidence accordingly, even if the evidence is *in fact* highly misleading, such that one does not in fact believe in a way that maximizes true belief and minimizes false belief. And conversely, one fails to be deontologically justified if, for all one can tell, one does not believe in a truth-conducive way, even if one in fact does.⁷

'So far as one can tell' is a notion that is less than clear, of course. Here, I will assume the standard reading that interprets it as 'so far as one can *reflectively* tell', which in turn can be glossed as 'according to the information reflectively accessible to one'. On this assumed reading, whether one is deontologically justified thus supervenes on the information reflectively accessible to the subject. This is one reason to think the deontological concept of justification implies accessibilism.

It follows from the above that S is deontologically justified only if, given the information reflectively accessible to her, she forms her beliefs in such a way that she maximizes truth and minimizes falsity. Now it is notoriously difficult to say what exactly is required for believing in a way that maximizes true belief and minimizes false belief. *Prima facie*, however, it seems reasonable to suppose that whatever is required, it excludes believing in a way that for all one can tell could have easily produced a false belief instead. Thus, under the assumption that forming beliefs in a way

⁷ Examples of this latter kind include the 'chicken-sexer' cases, as well as BonJour's clairvoyant Norman, both discussed in the previous chapter.

that from one's perspective could have easily produced a false belief is not a way that, so far as one can tell, maximizes true belief and minimizes false belief, it follows that one's belief cannot be deontologically justified and reflectively lucky at the same time.

What can be said in favour of this assumption? First, it seems that when we are considering which methods of belief-formation will maximize true belief and minimize false belief, it is clear that methods that could easily produce false beliefs are always trumped by methods that not easily could easily produce false belief. Using the first kind of method will result, in the long run, in more false beliefs than the second. Since maximization requires using the best available method, it seems that the epistemic norm under consideration would prohibit using methods that, for all we can tell, could have easily produced false belief. Second, the assumption that using such methods will result in a violation of the epistemic norms is backed-up by paradigmatic cases from the literature. For example, both Norman the clairvoyant and Brandom's chicken-sexers are usually considered as paradigm examples of subjects who are not deontologically justified. For neither of them believes in a way that, for all she can tell, is 'well-calculated' to reach the truth. As we have seen in the previous chapter, however, the beliefs of these subjects are also paradigmatic cases of reflective luck. This provides further support for the claim that forming your belief in a way that could have easily produced false belief will result in a violation of the epistemic norm of maximizing true belief and minimizing false belief.

Our argument does not depend on the specific epistemic norm under consideration. It is plausible that the argument can be generalized to work with any norm that will vindicate the accessibilist criterion on justification. Take, for example, an epistemic norm that says that you should adopt a degree of belief in p equal to the strength of the evidence

for p available to you. If this norm is to support the accessibilist framework, then the evidence available to you must be reflectively accessible. Otherwise, since on the deontological account justification just is matter of not violating any epistemic norms, justification would not supervene on what is reflectively accessible to you. It seems clear, however, that if one's belief is reflectively lucky, one does not meet the evidentialist norm thus understood. For it seems that complying with this norm requires you not to form beliefs that, according to that same evidence, are formed in a way that could have easily produced false belief.

If the above is right, then deontological justification, at least as appealed to by accessibilists, is incompatible with reflective luck. But such justification is quite compatible with the presence of high degrees of veritic luck. To see this, we need only consider some famous cases of deontological justification provided by the literature, such as the beliefs of victims of the New Evil Demon.⁸ In this case, victims of the New Evil Demon have reflective access to exactly the same information that we do. And just like us (we may suppose), these people are doing the best they can, from their perspective, to maximize true belief and minimize false belief. According to the deontologist, their beliefs are therefore justified. However, due to the Demon's interventions, their beliefs are formed in such way that they could very easily produce false belief (indeed, many of their beliefs *are* false). This means that insofar as they do form true beliefs, their beliefs will be subject to a high degree of veritic luck. While accessibilist justification thus eliminates reflective luck, it is compatible with veritic luck.

⁸ See Chapter 2.

5.2 MENTALISM

I argued above that one traditionally dominant form of internalism is incompatible with reflective luck but not with veritic luck. In this section, I will do the same for a form of internalism that is particularly dominant in contemporary literature: mentalism. As we saw in Chapter 2, mentalism can be captured by the following thesis:

MENTALISM: The justificatory status of a person's doxastic attitudes strongly supervenes on the persons occurrent and dispositional mental states, conditions and events. (Conee & Feldman, 2001, p. 2)

In Chapter 2, I argued that while mentalism and accessibilism are logically independent, it is not clear that understood as bare claims about the supervenience base of justification, they represent substantially different viewpoints. We saw that under some plausible assumption, their supervenience claims would amount to the same thing: that every justificatory difference implies a difference in the reflectively accessible mental states of a subject.

However, we also saw that mentalism's prime defenders adopt a concept of justification that is different from the concept usually endorsed by accessibilists. Therefore, we cannot straightforwardly infer from the fact that that accessibilist justification as supplemented with a deontological concept of justification is incompatible with reflective luck, that the same holds for mentalism when proponents of the latter understand justification in a different, evidentialist way.

Thus, I must argue for the incompatibility between mentalist justification and reflective luck independently. I will use the following strategy. First, I will revisit some of the cases used by Feldman and Conee to support their mentalist theory of epistemic justification. I will argue that insofar

as these cases support the claim that justification supervenes on the mental, they also support the claim that justification and reflective luck are incompatible. This serves to show that at least in these paradigm cases, more mentalist justification means less reflective luck.

Second, I will focus on the fact that, as we have seen in Chapter 2, mentalism is commonly supplemented with a more substantial view on the nature of epistemic justification. Whereas accessibilists have tended to opt for a deontological concept of justification, mentalists primarily draw on an *evidential* concept. I will argue that under some reasonable assumptions, we can show that mentalism combined with an evidentialist concept of justification eliminates reflective luck, but not veritic luck.

We thus look at some of the classic cases used to support mentalism:

CASE 1: Bob and Ray are sitting in an air-conditioned hotel lobby reading yesterday's newspaper. Each has read that it will be very warm today and, on that basis, each believes that it is very warm today. Then Bob goes outside and feels the heat. They both continue to believe that it is very warm today. But at this point Bob's belief is better justified. (Conee & Feldman, 2001, p. 3).

Bob's perception of the heat is a mental state that raises the justification of his belief that it is very warm today. Bob's perceptual state is not only reflectively accessible, it is reflectively *accessed*.⁹ A similar perceptual state is not reflectively accessible to Ray, however, as he stays put in the hotel lobby. Because of this, the mental state that Conee and Feldman identify as responsible for a difference in justification between Bob and Ray constitutes a difference in

9 Cf. Section 2.2.

the information reflectively accessible to Bob and Ray as well. Moreover, given the information reflectively accessible to Bob, he could have less easily formed a false belief than Ray could have. Both Bob and Ray believe the actual world to be very warm. But only the facts reflectively accessible to Bob include the fact that the world *feels* very warm. For Ray's method of belief-formation to produce a false belief, all that is required is a misspelling in the newspaper. For Bob's method to produce a false belief, not only the newspaper needs to be mistaken, but there needs to be some elaborate deception going on outside the hotel as well. Given their respective bodies of reflectively accessible information, Bob's method of belief-formation could have less easily produced a false belief, and his belief is thus subject to a lesser degree of reflective luck than Ray's belief.

The second case that we will discuss is the following:

CASE 2: A logic Teaching Assistant and a beginning logic student are looking over a homework assignment. One question displays a sentence that they both know to express a truth and asks whether certain other sentences are true as well. The TA can easily tell through simple reflection that some of the other sentences express logical consequences of the original sentence and thus she is justified in believing that they are true as well. The student is clueless. (Conee & Feldman, 2001, p. 4).

As this case is set up, it is explicitly stated that the TA has better reflectively accessible reasons than the student has for supposing that her beliefs about the other logical sentences are true. Since the student is clueless, we may suppose she forms her beliefs on the basis of simple guessing. Since simple guessing could easily produce a false belief,

the student's belief will be subject to a high degree of reflective luck. Not so for the TA. It is precisely the fact that she has the ability to easily tell through reflection that some of the sentences express logical consequences, whereas others do not, that means that she could not have easily formed false beliefs about the truth values of these sentences in the same way. Because of this, the belief of the student is subject to a higher degree of reflective luck than the belief of the TA. We thus again find that the mental state that according to Conee and Feldman makes the justificatory difference also eliminates reflective luck.

Our third and final case is the following:

CASE 3: Initially, Smith has excellent reasons to believe that Jones, who works in his office, owns a Ford. Smith deduces that someone in the office owns a Ford. The latter belief is true, but the former is false. Smith's reasons derive from Jones pretending to own a Ford. Someone else in the office, unknown to Smith, does own a Ford. The fact that Jones is merely simulating Ford ownership keeps Smith from knowing that someone in his office is a Ford owner, but it does not prevent Smith from being justified or diminish his justification. At a later time Smith gains ample reason to believe that Jones is pretending. At that point Smith is not justified in believing either that Jones owns a Ford or that someone in his office owns a Ford. (Conee & Feldman, 2001, p. 4).

In this case, Smith loses the justification he has for believing that someone in the office owns a Ford as a consequence of gaining 'ample reasons' to believe Jones is merely pretending to own a Ford. It seems these reasons should be reflectively accessible if they are to explain why

it would be reasonable for Smith to change his belief. As in the above cases, a change in justification is thus accompanied by a change in the information reflectively accessible to the agent. Does this change in reflectively accessible information lead to an increase of reflective luck? In the first case, Smith forms his belief that Jones owns a Ford on the basis of 'excellent reasons'. Given that he bases his belief on these excellent reasons, it seems that from *his* perspective, Smith could not easily have formed a false belief in same way.¹⁰ However, continuing to believe that Jones owned a Ford after receiving ample evidence to the contrary is a method that even from Smith's own perspective could easily produce false belief. So were Smith to disregard the evidence contradicting his belief in the second case, this would increase the degree of reflective luck to which his belief is subject. Again we find that the difference in justification is mirrored by a difference in reflective luck.

The above three cases each feature differences in mentalist justification that are mirrored by differences in reflective luck. While this is of course not logically conclusive, it does suggest a correlation between mentalist justification and the absence of reflective luck.

We can further support the connection between mentalist justification and the absence of reflective luck by considering the concept of justification that mentalism's chief defenders rely upon.

As we saw in Chapter 2, Earl Conee and Richard Feldman support a theory of justification that they call *evidentialism* (1985). Evidentialism essentially consists of the following thesis:

¹⁰ Of course, in reality, Smith's method of belief formation *could* have easily produced a false belief. That means his belief is veritically lucky. Crucial for the present point, however, is that from *Smith's* perspective his reasons would not easily lead him astray.

EJ: Doxastic attitude D toward proposition p is epistemically justified for S at t if and only if having D toward p fits the evidence S has at t . (Feldman & Conee, 1985, p. 15).

What is one's evidence? As we saw, Conee and Feldman's view is that one's evidence consists of one's mental states. If one's evidence consists of one's mental states, we may ask whether it consists of all of them, or only a subset thereof. Feldman defines one's *total possible evidence* as "all and only the information the person has stored in his mind at the time" (Feldman, 2004, p. 226). It includes not only the information actively considered by the subject at the time, but also the information that can only be recalled with great difficulty.

Interestingly, however, Feldman argues that only a subset of one's total possible evidence is *possessed* by the subject in the sense required for evidentialist justification. For that to be the case, two further conditions need to be met: a psychological accessibility condition, and an epistemic acceptability condition. With regard to the first condition, Feldman notes that one's total possible evidence will include mental states that can only be recalled with great difficulty. For example, there may be certain psychological primings that will allow me to remember the colour of the alarm-clock at my parents' first house. Not being primed in this way, I cannot for the life of me remember what the relevant colour was. Since there are conditions, however, under which I will recall the colour, it is part of my total possible evidence, but it not plausible that this deeply buried memory belief is part of the evidence currently in my possession, or so Feldman argues. As regards the second condition, my total possible evidence includes all mental states that I could possible recall, including presumably some mental states that are formed in epistemically deficient ways. Suppose I believe for no good reason that the

number of stars is even. From this belief I infer that the number of stars cannot be prime. But, as Feldman argues, we would not in this case say that I am currently in possession of a piece of evidence for the thesis 'the number of stars cannot be prime'.

Controversially, Feldman then goes on to argue that the evidence possessed by a subject consists exclusively of those propositions the subject is currently thinking of. Obviously, this implies, together with EJ, that not many of the beliefs we take to be justified will count as justified. For example, I take my belief that Amsterdam is the capital of the Netherlands to be justified, even if I am not currently thinking of anything that supports it. If so, it will not meet EJ, and so Feldman will be forced to say it is not justified. Feldman replies to this rather obvious objection by distinguishing two senses of justification: occurrent justification and dispositional justification. He then argues that EJ should be taken as an account of occurrent justification rather than of dispositional justification.

The distinction between occurrent and dispositional justification is controversial, as is Feldman's claim that the evidence we possess is exhausted by what we are currently thinking of. Is it really the case that the beliefs we have but are not currently thinking of are justified in a different way from the ones that are currently before our mind? Here, I will remain neutral on this topic. The points made will not depend on it.

We should say something about what it means for a belief to 'fit' one's evidence. Unfortunately, Conee and Feldman do not specify what it means for a belief to 'fit' a body of evidence in their intended sense, noting instead that while "there are difficult questions concerning the concept of fit ... [t]he application of EJ is clear enough to do the work that we intend here—a defense of the evidentialist position" (Feldman & Conee, 1985, fn. 2). The idea seems to be that

[b]elieving is the justified attitude when the person's evidence on balance supports a proposition, disbelieving is the justified attitude when the person's evidence on balance supports the negation of a proposition, and suspension of judgment is the justified attitude when the person's evidence on balance supports neither a proposition nor its negation. (Conee & Feldman, 2004, p. 102).

This formulation is still not entirely clear, but it will have to do for our purposes. In any case, my point will not depend on the specific definition of evidential 'fit'.

Perhaps surprisingly, I will now argue that EJ *can* be satisfied for reflectively lucky beliefs. Remember that we defined a reflectively lucky belief as a belief that, given the information reflectively accessible to the subject, is produced in such a way that she could have easily formed a false belief instead. Now, EJ says that a belief is justified if and only if it fits the evidence possessed by the subject. Crucially, EJ says nothing about *how* the subject has acquired her belief, nor about how she believes she acquired her belief. Suppose, for example, that I believe that my friend is innocent of a crime she is accused of. Suppose further that I believe I formed this belief out of loyalty to my friend, rather than on the basis of any evidence. Suppose finally that I believe many things that together imply that my friend is innocent. I simply fail to see the connection between my evidence and my belief. In this case, the proposition that my friend is innocent will satisfy EJ, but is nevertheless subject to a substantial degree of reflective luck. For given the way I *believe* I formed my belief, I could have easily formed a false belief instead, even if this belief fits the evidence I have very well.

Thus, it is possible for a belief to fit one's evidence perfectly, even if one believes one could have easily formed

a false belief instead. Moreover, it may fit one's evidence perfectly even though one *knows* one has formed the belief on the basis of simple guessing. For in the above example, it may be *true* that I formed my belief on the basis of loyalty, and we may suppose any further conditions for knowledge to be met as well. The belief would still satisfy EJ, and would still be subject to a very high degree of reflective luck.

It would thus seem that, contrary to what we set out to argue for, the view under consideration is compatible with reflective luck. However, as we have seen in Chapter 2, Conee and Feldman use the phrase 'epistemic justification' in a non-standard way, referring to *propositional* justification, rather than doxastic justification.

We have also seen that our main concern is doxastic rather than propositional justification. That is, we are concerned with a property of held beliefs, rather than a property of propositions. This is also the kind of justification that our other internalists target. We saw above that the accessibilist justification of a belief depends on whether the subject *in so believing* satisfied the epistemic norms. For deontologists about epistemic justification, the method of belief formation is thus of crucial importance for epistemic justification. This means that they are talking about doxastic justification, since only held beliefs have a method of belief-formation. Propositions do not. And a similar argument can be given for why the notion of epistemic justification that is given central importance in externalist theories like reliabilism is doxastic justification rather than propositional justification. For reliabilism, what matters for justification is whether one's belief is formed through a belief-forming method that is *reliable*. Again, only actual beliefs have a method of belief-formation, propositions do not.

Thus, beliefs that satisfy EJ but are reflectively lucky do not provide counterexamples to our claim that internalist theories of *doxastic* justification are incompatible with

reflective luck. For that, we have to consider the mentalist picture of doxastic justification. As we saw in Chapter 2, Conee and Feldman call doxastic justification 'well-foundedness', and provide the following definition:

WELL-FOUNDEDNESS: *S*'s doxastic attitude *D* at *t* toward proposition *p* is well-founded if and only if

1. having *D* toward *p* is justified for *S* at *t*;
and
2. *S* has *D* toward *p* on the basis of some body of evidence *e*, such that
 - a) *S* has *e* as evidence at *t*;
 - b) having *D* toward *p* fits *e*; and
 - c) there is no more inclusive body of evidence *e'* had by *S* at *t* such that having *D* toward *p* does not fit *e'*. (Feldman & Conee, 1985, p. 24)

As is apparent from its definition, well-foundedness is a notion that applies to actually held doxastic attitudes, rather than merely possibly believed propositions. The notion is clearly meant as an epistemic evaluation of the doxastic attitude. It is thus closer to our notion of doxastic justification than Conee and Feldman's notion of epistemic justification.

Can a belief be well-founded yet reflectively lucky? As the above definition makes clear, a belief is well-founded only if it is based on a set of evidence *e* that is available to the subject and that fits the belief, and there is no more inclusive set of evidence available to the subject that does not fit her belief.

Whatever the precise notion of fit, it seems, *prima facie*, that forming beliefs on the basis of evidence that fits them is a method that, from one's perspective, would not easily

produce false beliefs. As a paradigmatic case, consider my belief that the grass outside is green on the basis of having a lush green perceptual experience. It seems that this belief satisfies the criteria for well-foundedness, in that it fits the body of evidence on the basis of which it is formed. Similarly, it seems that from my perspective, I could not easily have formed a false belief in the same way. Some elaborate deception would have to be going on for my method to have produced a false belief. If this *would* have been an easy possibility from my perspective, it is hard to see how my belief still fits my evidence.

It is also clear that my belief in the innocence of my friend discussed above does not provide a counterexample to this claim. For in that case, while I possess evidence for believing that my friend was innocent, I do not form my belief *on the basis* of this evidence. So even though the case satisfied EJ, it will not satisfy WF (at least assuming that I formed my belief purely out of loyalty, and this itself would not constitute evidence that fits my belief).

A possible objection to this line of argument would be that the basing relation may be interpreted externalistically; that it may not be clear from one's perspective on what basis one's belief is formed. On this conception of the basis relation it would be possible that I believe that I form my belief about my friend's innocence out of loyalty, while in fact I form the belief on the basis of evidence that fits it. In this case, my belief would be well-founded but still reflectively lucky.

The following passage makes clear, however, that Conee and Feldman intend a more internalistically friendly interpretation of the basing relation, and that consequently, well-foundedness excludes reflective luck:

If a belief is justified and otherwise apt for knowledge, but not well-founded, then the belief is accidentally correct in a way that some-

what resembles what occurs in the classic Gettier cases. But, unlike in classic Gettier cases, there is nothing epistemically defective about the connection between the person's justification for P and a fact making P true. The belief itself is not held *in light of* the justification, though. It is held dogmatically, or from wishful thinking, or on some other epistemically faulty basis. (Conee & Feldman, 2004, p. 105, *my italics*)

This passage fits well with our story in two ways. First, well-foundedness requires that the belief is held *in light of* its justification. The phrase 'in light of' suggests that the subject must be aware of the connection between her belief and its justification. This seems to exclude an externalist reading of the required basing relation. In the example above, for instance, I clearly do not believe that my friend is innocent *in light of* the fact that this belief fits my evidence. So this belief would not be well-founded according to the passage above.

Second, Conee and Feldman argue that justified but not well-founded beliefs are 'accidentally correct'. Using the notion of reflective luck, we can spell out the exact sense in which they are accidentally correct: her beliefs are accidentally correct in the sense that given the information available to our subject, her method of belief formation could have easily produced a false belief instead. Her beliefs are *reflectively* lucky.

We thus see that there are good reasons to suppose that a belief that is well-founded in Conee and Feldman's sense cannot be subject to a substantial amount of reflective luck. Well-foundedness does not exclude veritic luck, however. Consider again the case where I form the belief that the grass outside is green on the basis of a lush green experience. My belief may be well-founded in the sense that it fits

the evidence on which it is based very well and there are no defeaters present. However, now suppose that my evidence is highly deceptive. While my experiences tell me I am an embodied philosopher currently looking outside his window, I am *actually* a brain in a vat stimulated by evil scientists. Note that in this case my belief will still be well-founded. It fits the evidence in light of which it is held, and the evidence available to me does not include any defeaters for the belief. But my belief, if true at all, will of course be subject to a rather large degree of veritic luck. For the way I *actually* formed my belief could have very easily produced a false belief. Even if in the world of the evil scientists the grass outside is actually green, trusting the experiences these scientists feed me could have very easily produced a false belief.

5.3 CONCLUDING REMARKS

In this chapter we investigated the relation between internalism about epistemic justification and luck. In particular, we argued that the two main internalist theories of justification in contemporary epistemological literature are incompatible with reflective luck but not with veritic luck.

The first internalist theory of justification we investigated combined a deontological account of justification with an accessibilist condition on the supervenience base of justification. According to this form of internalism, a belief is justified if and only if it is formed in a way that does not violate any epistemic norms, and the factors that determine whether this is the case should all be reflectively accessible. We saw that on a common interpretation of the norm, reflectively lucky beliefs will violate it. Reflectively lucky beliefs will then not be justified according to this conception of justification. This form of internalism was

found to allow for justified beliefs that were subject to high degrees of *veritic* luck, however.

The second internalist theory of justification we identified was the theory of justification recently put forward by Earl Conee and Richard Feldman. It consists of a combination of views that have been labeled evidentialism and mentalism. Where mentalism claims that justification supervenes on one's mental states, evidentialism concerns the claim that epistemic justification is a matter of fitting the evidence. While our analysis was complicated by the fact that what Conee and Feldman call 'epistemic justification' is something different from what we—and as we argued most other internalists and externalists too—call epistemic justification, we found that, properly understood, Conee and Feldman's views on doxastic justification are incompatible with reflective luck but not with veritic luck.

We thus found that the two main internalist theories of justification in contemporary literature are incompatible with reflective luck, but not with veritic luck. Of course, the forms of internalism investigated in this chapter are not the only possible ones. They do constitute by far the most prominent forms of internalism in contemporary literature, however. Showing that these two forms share a commitment to the absence of reflective luck, but not veritic luck, constitutes the first step towards showing that the internalism/externalism-debate about epistemic justification can be captured in terms of veritic and reflective luck.

The findings of this chapter support a non-standard interpretation of the import of Gettier cases. As we said above, many epistemologists think that Gettier showed that justification could not deliver on the promise of excluding (problematic forms of) luck from knowledge. This chapter shows that we can provide a more nuanced picture. For internalist justification eliminates reflective luck but not veritic luck. We have seen in the previous chapter that reflective luck is incompatible with the possession of

reflective knowledge but compatible with the possession of animal knowledge. This raises questions about the function of internalist justification. For it seems that there is a kind of knowledge that is incompatible with the kind of luck that internalist justification eliminates. Perhaps the kind of knowledge that philosophers like Plato have been interested in has always been reflective knowledge. For that kind of knowledge, the elimination of reflective luck is necessary, and the results of this chapter indicate that an internalist justification condition is well-suited for that task. What Gettier draws our attention to, on this picture, is precisely the fact that there are different forms of malignant epistemic luck, forms of luck that require different forms of justification to eliminate.

In the next chapter, we will see that externalist justification eliminates veritic luck but not reflective luck. This will open up the way for a new interpretation of the internalism/externalism-debate about epistemic justification.

EXTERNALISM AND LUCK

In the previous chapter I argued that the most prominent forms of internalism are incompatible with reflective luck, but compatible with veritic luck. In this chapter, I will argue that the situation is exactly the opposite with regards to *externalist* justification. That is, I will argue that most prominent externalist concepts of justification are compatible with reflective luck, but *incompatible* with veritic luck. I will also discuss some implications of these findings.

We closed the last chapter with the suggestion that the main import of Gettier's paper is that there is a kind of luck that is compatible with internalist justification—veritic luck—but that nevertheless precludes knowledge. One *prima facie* reason for thinking that externalist theories in general are incompatible with veritic luck stems from the historical development of such theories in response to Gettier's famous problem. Generally, Gettier-subjects are taken to satisfy internalist criteria for justification, but lack knowledge because their beliefs are veritically lucky: given the way these subjects formed their beliefs, they generally could have very easily formed a false belief instead. One way to see the development of externalist criteria for justification in response to Gettier cases is as attempts to provide justification conditions that *would* eliminate veritic luck.¹ One of the main aims of this chapter is to see whether and how prominent externalist concepts of justification are able to do so.

The chapter consists of two parts. In the first part I show that prominent externalist theories of justification are in-

¹ For similar interpretations, see Poston (2016) and Pritchard (2005a, p. 152).

deed incompatible with veritic luck, but compatible with reflective luck. In Section 6.1 I show that reliabilist justification is incompatible with veritic luck. In Section 6.2, I discuss Sosa's virtue epistemology, and argue that it, too, is incompatible with veritically lucky justified belief. In Section 6.3, I argue that both reliabilism and virtue epistemology allow for justified beliefs that are reflectively lucky.

In the second part of the chapter I consider some consequences of the results of the first part. In Section 6.4, I argue against Zagzebski's (1994) claim that only *factive* accounts of justification—accounts that hold that justification entails truth—evade Gettier cases. In Section 6.5, I discuss some implications of our findings concerning the relation between justification and luck and the recently popular *safety* conditions on knowledge. In Section 6.6, I use the findings of this chapter and the previous to argue for a re-conceptualization of the internalism/externalism-debate about epistemic justification. I close in Section 6.7 with some concluding remarks.

6.1 RELIABILISM

In the first part of this chapter, I argue that the externalist theories of justification we discussed in Chapter 2 are incompatible with veritic luck but compatible with reflective luck. The first of these theories is simple process reliabilism as proposed in (Goldman, 1979). In this section I argue that, properly understood, and some recent claims to the contrary notwithstanding, this kind of justification eliminates veritic luck.

In Chapter 2, we saw the core of simple process reliabilism consists of the following thesis:

RELIABILISM: If [and only if] *S*'s believing *p* at *t* results from a reliable cognitive belief-forming

process (or set of processes), then S 's belief in p at t is justified.² (Goldman, 1979, p. 96)

What does it mean for a belief to be reliably produced? According to Goldman, the cognitive process that caused the belief must have a 'tendency' to produce true beliefs (Goldman, 1979, p. 96). As we saw, the required tendency should be understood as a tendency to produce true belief *across nearby possibilities*:

[A] cognitive mechanism or process is reliable if it not only produces true beliefs in actual situations, but would produce true beliefs, or at least inhibit false beliefs, in relevant counterfactual situations. (Goldman, 1976, p. 771)

It is doubtful that a belief can be reliable in this sense yet be veritically lucky. Recall our definition of veritic luck:

VERITIC LUCK₆: A belief is veritically lucky if and only if it is a matter of luck that the method one used to form one's belief produced a true belief.

On our definition of luck, this means that a belief is veritically lucky if and only if it is true and produced in a way that could have easily produced false belief instead. If the counterfactual situations relevant for reliability include the nearby possibilities where one forms the belief in the same way, then reliably produced belief will not be veritically lucky. For reliability is now simply defined in terms of whether the relevant belief-forming process could have easily produced false belief instead. Thus understood, reliabilism excludes veritic luck.

² As we saw, Goldman formulates his base clause as a sufficient condition, even though it is clear from the rest of his work that he intends it to be necessary as well. See Section 2.4.

The claim that reliability excludes veritic luck is controversial, however. Here I will discuss two objections. The first one derives from William Harper. Harper asks us to suppose that

Smith forms a reliably formed belief, by normal methods, that Jones owns a Ford, but, unbeknownst to Smith, Jones' Pinto is blown to dust by a terrorist's bomb and simultaneously Jones wins a Falcon in the State lottery. Smith has a reliably formed true belief that Jones owns a Ford, but his belief is not knowledge. (Harper, 1996, p. 277)

If Harper is right, reliability does *not* eliminate veritic luck, contrary to what we said above. My response starts by noting that it is unclear in Harper's description of the case whether the terrorist attack is supposed to be an easy possibility or not. Let us suppose first that it is. In this case Smith's method of belief-formation could have easily produced false belief, namely in the case of the terrorist attack, and so his belief will not be modally reliable.

If the terrorist attack is not easily possible, Smith's method of belief-formation will be reliable across the relevant counterfactual situations, and his belief will thus be modally reliable. But if the terrorist attack is not easily possible, then Smith's belief is not veritically lucky either. For now his method could not have easily produced a false belief instead. Jones will keep his Pinto across the relevant easy possibilities where Smith forms his belief in the same way. In some of them, he will also have won the Falcon. But that does not make Smith's belief that Jones owns a Ford false.

Once we remove the ambiguity concerning the modal distance of the terrorist attack, we see that neither of the possible interpretations provides us with reason to revise our claim that reliability is incompatible with veritic luck.

Another alleged counterexample to our claim comes from Linda Zagzebski (1994). As Zagzebski argues:

a breakdown in the connection between a reliable belief-forming process and the truth is possible. When that happens, even if you manage to hit on the truth anyway, you do not have knowledge. (Zagzebski, 1994, p. 66)

The implication is that you do not have knowledge because your beliefs could still be true merely by luck. As an example, she provides following case:

The well-known fake barn case can be described as an example of this sort. Here we are to imagine that you are driving through a region in which, unknown to you, the inhabitants have erected three barn facades for each real barn in an effort to make themselves look more prosperous. Your eyesight is normal and reliable enough in ordinary circumstances to spot a barn from the road. But in this case the fake barns are indistinguishable from the real barns at such a distance. As you look at a real barn you form the belief 'That's a fine barn'. The belief is true and justified, but is not knowledge. (Zagzebski, 1994, p. 66)

Let us grant that the belief does not constitute knowledge.³ Still, there is a way for the reliabilist to deny that the relevant belief is reliably produced. For reliability is context relative. What may be a reliable method of belief-formation for us may not be so for a Brain In a Vat (BIV).

³ We have seen, in Chapter 4, that fake barn cases involve environmental, not veritic luck, and that it is not clear that they undermine all forms of knowledge. More on this below.

Whether a method is reliable thus depends on the environment in which it is used. Relative to the environment described by Zagzebski, the subject's method of forming beliefs about barns by trusting one's eyesight is not reliable. True, as Zagzebski notes, in *our*, normal, context this method is (presumably) reliable, even *modally* reliable, but this need not be so in fake barn county. Our actual world, and the worlds surrounding it are not the actual world and the worlds surrounding the actual world in fake barn scenario's. In these scenario's the method used could have easily produced a false belief instead, and would thus not be modally reliable. On our interpretation of veritic luck and reliability the considered counterexamples thus fail to present problems for our claim that reliability eliminates veritic luck.

We can further support our claim by considering another important motivation for reliabilism. As we saw in Chapter 2, reliabilism can be supported by the truth-conducive account of justification. The truth-conducive account holds that a belief is justified only if it is produced in such a way as to make it 'at least very likely that the belief is true'. Such a conception of justification motivates reliabilism, because reliably produced beliefs are guaranteed to satisfy this criterion.

I will now argue that the truth-conducive concept of justification itself is incompatible with the presence of veritic luck. This would mean that insofar as reliabilists intend their conditions of justification to guarantee truth-conducivity, they must eliminate veritic luck. Truth-conducive beliefs are beliefs that are very likely to be true. We can interpret this claim in different ways.

First, we may distinguish first between 'objective' and 'subjective' likelihood of truth. A belief is objectively likely to be true if and only if it is likely to be true, given the objective facts of the situation. It is subjectively likely to be true if and only if it is likely to be true, given the subjective per-

spective of the believing subject. The most plausible way to understand the truth-conducive concept, at least insofar as it is used to support externalist theories of justification like reliabilism, is in terms of objective, rather than subjective, likelihood of truth. Externalists usually count the beliefs of victims of the New Evil Demon as unjustified. In particular, they will want to say that this is so because the relevant beliefs are not truth-conducive. Yet these beliefs may have a high subjective likelihood of truth. So insofar as reliabilists draw on a concept of truth-conducivity in which these beliefs will fail to be truth-conducive, they are committed to an objective understanding of the relevant likelihood of truth.

Second, truth conducivity can be given a modal and a probabilistic reading. On the modal interpretation, a belief satisfies the truth-conducive concept if and only if it is produced in a way that could not *easily* have resulted in a false belief instead. On the probabilistic reading, the truth-conducive concept is satisfied if and only if the belief is produced in a way that has a low probability of producing a false belief. I argued in Chapter 3 that how easily something is possible depends on more than just probability alone. Also modal distance is important. For this reason, the modal and the probabilistic reading do not reduce to each other. For if I form my belief that my lottery ticket will lose purely on the basis of the relevant odds, my method may have a low probability of producing false belief, but since the nearest world where that belief is false may be very similar to the actual world, there is still a sense in which my method could have easily produced a false belief, and the resulting belief will thus still be subject to a certain degree of veritic luck.

On the modal interpretation, the truth-conducive concept is nothing over and above the requirement that justified beliefs are not veritically lucky. For to be veritically lucky the belief would have to be produced in a way that

could have easily produced false belief instead, which is precisely what is the truth-conducive concept says justified beliefs cannot be.

As lottery beliefs show, belief-forming methods that have a high probability of producing true belief, and are truth-conducive in the purely probabilistic sense, will not suffice to rule out veritic luck.⁴

The fact that we commonly do not consider lottery beliefs to be cases of knowledge suggests that the truth-conducivity required for knowledge is stronger than the purely probabilistic kind.⁵ If we want our concept of justification to be truth-conducive in the sense required for knowledge, we should thus interpret it in modal terms.

If justified beliefs are truth-conducive in the modal sense, then lottery beliefs are not justified. Yet, there seems to be some intuitive pull towards saying that these beliefs *are* justified. Here, one person's *modus ponens* is another's *modus tollens*. I feel inclined to bite this particular bullet and say that lottery beliefs are not justified. The advantage of this is that it allows us to provide a straightforward explanation of *why* we are disinclined to ascribe knowledge in these cases: the relevant beliefs are not justified and justification is necessary for knowledge. This also allows me to explain why reliabilists like Goldman would provide modal interpretations of their reliability conditions: they think justification is truth-conducive in a modal sense, which requires a modal form of reliability. Such a concept of justification will be incompatible with veritic luck.

⁴ The epistemic relevance of this is recognised in (Williamson, 2009), where it is argued that knowledge is subject to a modal safety condition, rather than a probabilistic one.

⁵ Martin Smith uses similar arguments to argue for a modal interpretation of justification directly (Smith, 2010, 2016). The present account differs most significantly from Smith's account in that I maintain, while Smith denies, that many of the usual accounts of justification should be understood in modal terms. For Smith, justification has classically been understood in probabilistic terms.

This concludes my defence of the claim that reliabilism eliminates veritic luck. In the next section, I consider the other prominent externalist account of justification that we discussed in Chapter 2, and argue that it, too, is incompatible with veritic luck.

6.2 VIRTUE EPISTEMOLOGY

In Chapter 2, we saw that virtue epistemology provides another prominent externalist theory of justification. As before, I will focus on Sosa's version.

To recap, Sosa distinguishes between animal and reflective knowledge.⁶ Here we focus on the former, since it is this kind of knowledge which brings out the externalist character of Sosa's concept of justification most clearly.⁷

About justification, Sosa says the following:

we can then evaluate it [a belief] as "epistemically justified", in one or another sense: "competently adroit" perhaps (or reliably based, or counterfactually safe, etc.), or perhaps "rationally justified" (coherently fitting, and held in part on that basis). (Sosa, 2009b, p. 114)

It is clear from the context of this remark that the latter kind of justification is the kind required for reflective knowledge, whereas the former is the kind required for animal knowledge.⁸ Adroitness is required for either sense of

6 And in his most recent work between knowledge full well too. But for our present purposes the distinction between animal and reflective knowledge is sufficient.

7 Although, as we have seen, reflective knowledge requires animal knowledge, and so reflective knowledge will require externalist justification as well. More on reflective knowledge below.

8 See Section 4.3.

epistemic justification.⁹ What is required for a belief to be *adroit*? As we saw in Chapter 2, Sosa thinks that a belief is *adroit* just in case it manifests an epistemic competence to believe truly. In general, to possess a competence for X requires that one could not easily fail to X if one tried, under the appropriate circumstances. Thus, to possess an epistemic competence requires that one does not easily fail to believe truly, were one to try to do so, under the appropriate circumstances.

As Sosa says, manifesting a competence requires that circumstances *are* indeed appropriate (Sosa, 2015, p. 103). Thus, a shot *manifests* a competence for archery only if the conditions are in fact appropriate in the sense that reliable success in *these* conditions is required to possess the competence of archery at all. Similarly, manifesting an epistemic competence requires that the conditions are such that if one under those conditions could have easily failed to believe truly, one would not count as possessing the epistemic competence at all.

A belief is thus justified only if it is *adroit*, which in turn requires that in those conditions one could not easily have failed to form a true belief. This rules out veritic luck. Thus, Sosa's account of justification eliminates veritic luck as well.

It is instructive to consider what happens in cases where the conditions are not appropriate. In such cases, a belief cannot *manifest* the epistemic competence, and thus fails to be justified according to Sosa's criteria. This is one point where the externalist character of Sosa's virtue theory clearly shines through. As is typical for externalist theories of justification, Sosa is committed to the claim that victims of the new evil demon and recently envatted brains will not be justified. Presumably, in such situations condi-

⁹ Interestingly, *adroitness* is spelled out in counterfactual terms: on Sosa's view, a belief is *adroit* just in case it is 'counterfactually safe'. Below I discuss whether this kind of safety suffices to rule out veritic luck.

tions are not appropriate for the competences possessed by the believing subjects, and as a result, their beliefs will not manifest those competences, and will thus not be justified.

However, perhaps surprisingly, given the above, Sosa does allow for justification, even knowledge, in cases of substantial *environmental* luck. Consider the following case:

You see a surface that looks red in ostensibly normal conditions. But it is a kaleidoscope surface controlled by a jokester who also controls the ambient light, and might as easily have presented you with a red-light+white-surface combination as with the actual white-light+red-surface combination. Do you then know the surface you see to be red when he presents you with that good combination, despite the fact that, even more easily, he might have presented you with the bad combination?

Arguably, your belief that the surface is red is an apt belief, in which case it amounts to knowledge, or so it does according to our account. (Sosa, 2007, p. 31)

This makes clear that Sosa thinks environmental luck is compatible with knowledge.¹⁰ However, he recognizes that this is somewhat counterintuitive, given that fake barn cases and the like have been traditionally regarded as *Gettier cases*, that is, as cases where there is a distinctive *absence* of knowledge.

Sosa's solution lies in his insistence on the two distinct grades of knowledge we have already come across: animal and reflective knowledge. Where cases of environmental luck are compatible with animal knowledge, they exclude reflective knowledge. I want to mention this point rather

¹⁰ Sosa explicitly likens the kaleidoscope example to the Ginet/Goldman barns example (Sosa, 2007, fn. 5.2).

than discuss it in detail. What is important here is that on Sosa's conception of animal knowledge and justification, both of these are incompatible with veritic luck but compatible with environmental luck.

If Sosa thinks that justification is compatible with environmental luck, he must think that in such cases the conditions for the exercise of the epistemic competence in question are appropriate, and thus that the subjects in such cases will form true beliefs reliably enough to be justified. It should be noted that this is a different interpretation of the fake barn cases than the one we provided above when we discussed reliabilism. There we argued that reliabilism could deny knowledge in these cases on the basis that in these environments, the methods used for belief-formation are not reliable. What we see here is that Sosa maintains that we do have knowledge in such cases, and thus needs to show how beliefs in these cases can be competently produced.

To see how this can be done, consider two different versions of the fake barn case, both compatible with how such cases are usually described. Both cases start in the standard way. I drive through fake barn county and form the belief that there is a nice barn over there. In both cases, I am looking at the only real barn around. The relevant difference between the two cases is the amount of time that I spend in fake barn county. In version 1, we suppose that I live in fake barn county. In this version of the case, it is clear that fake barns are part of the 'appropriate environment' for forming such beliefs. Consequently, forming beliefs about these barns just by looking will not manifest an epistemic competence since, in what are for me normal circumstances, this method will produce many false beliefs. Consequently, in this version of the case, we can deny justification.

In version 2 of the case, we suppose that I drive past fake barn county, glimpsing across the border just long enough

to form the belief that there is a barn there. In this case, the relevant normal circumstance will be mostly *normal* circumstances, that is, circumstances in which there are only real barns and no fakes. In this case, my belief forming method may qualify as an epistemic competence. It may also qualify as a manifestation of such a competence, since I am using my normal perceptual faculties to arrive at the belief about *this particular* barn. Even if there are many fake barns around, this does not prevent me from manifesting my competence on *this specific* occasion. In this case, my belief will thus be justified, and qualify as knowledge. As we have seen, this implies that my belief will not be a case of veritic luck. The way I formed my belief could, in this version of the case, not easily have produced a false belief instead.

An important upshot of the present discussion is that we see how externalist accounts of justification can both accommodate and deny knowledge in fake barn cases. Which verdict one goes with depends on what one thinks is the relevant environment for evaluating belief-forming process in these cases. Relative to more global environments, one's belief-forming methods will be reliable and competent, which provides reason for thinking these cases involve knowledge. Relative to more local environments, they are not, and this provides reason for denying knowledge. What the relevant environments are for epistemically evaluating belief-forming processes is a question I will not try to answer here. What is important here is that we see *how* Sosa could argue to support his claim that environmental luck does not undermine knowledge.

6.3 EXTERNALISM AND REFLECTIVE LUCK

So far, the externalist accounts of justification considered are incompatible with veritic luck. As the cases of Norman

the Clairvoyant and our naive chicken-sexers show, however, neither of these theories rules out reflective luck.

Consider first Norman the Clairvoyant.¹¹ While Norman forms his belief about the current whereabouts of the president on the basis of a highly reliable faculty of clairvoyance, he has no reflectively accessible information to suppose that he is actually in possession of such a faculty. Since his belief is produced by a highly reliable belief-forming process, it is justified according to reliabilism. Since Norman does not have any idea how he forms his beliefs, however, it will seem like an easy possibility from *his* perspective that they are produced in an unreliable fashion. Consequently, as we saw in Chapter 4, his beliefs will be reflectively lucky.

Norman's case may or may not satisfy Sosa's criteria for justification. That depends on whether we want to regard Norman's clairvoyance powers as cognitive competences of his. When we consider just the fact that if Norman were to try to form beliefs about the president's whereabouts, he could not easily form false beliefs, it seems that his beliefs would qualify as justified.

Next we consider again the chicken-sexer case. The case is structurally very similar to Norman's case, with the exception that these cases are somewhat more straightforward because they do not rely on any implicit assumptions we may have about the possibility and/or reliability of clairvoyance. Beliefs formed by such chicken-sexers are reliably produced and thus satisfy the reliabilist criterion for justification. Yet they are, as we have seen in Chapter 4, subject to substantial reflective luck. Like Norman, these people do not have any idea how they form their beliefs. This means that from their perspective, their beliefs might as well have been produced by an unreliable method, and thus, from their perspective, their beliefs could have easily

¹¹ See Chapters 2, 4.

been false. Even if their beliefs are formed reliably, they are thus nevertheless reflectively lucky. In this case too, reliability fails to eliminate reflective luck.

Note that our chicken-sexers provide a natural analogue to how we form most of our beliefs ourselves. Often we do not reflect on how we form our beliefs, or whether this method is reliable. If the virtue epistemologist regards the latter as manifesting competence in so believing, she should similarly allow the chicken-sexers to do so. Thus, it seems, Sosa's requirements for justification will be fulfilled as well, even though the beliefs of these chicken-sexers are subject to substantial degrees of reflective luck.

Sosa's multi-level picture of knowledge and justification introduces some complexity into our story about the relation between externalist concepts of justification and luck. So far, the considered concepts are compatible with reflective luck. But this may not hold true for the kind of justification required for Sosa's reflective knowledge. Consider the following quote again:

we can then evaluate it [a belief] as "epistemically justified", in one or another sense: "competently adroit" perhaps (or reliably based, or counterfactually safe, etc.), or perhaps "rationally justified" (coherently fitting, and held in part on that basis). (Sosa, 2009b, p. 114)

Where for animal knowledge the first sense of justification suffices, reflective knowledge requires the second, stronger sense. Here it is not only necessary that one's belief is counterfactually safe, but also that one's belief is held in part on that basis. As we saw in Chapter 2, this requires that one has at least a second-order belief that one's first-order belief is adroit. If one believes that one's belief is adroit, one believes that one's belief is competently produced and thus produced in a way that could not eas-

ily have produced false belief. This entails that one's first-order belief will be immune from reflective luck, because now it seems that from one's perspective, one could *not* have easily formed a false belief instead. So the justification required for reflective knowledge eliminates reflective luck. Since reflective knowledge is required for knowledge-full well, this latter kind of knowledge will also be immune from reflective luck. Both these higher forms of knowledge require animal knowledge, so both forms will *also* be immune from veritic luck.

Note that the kind of justification required for reflective knowledge is weakly externalist, since justification still requires the absence of veritic luck, a property to which the believing subject does not have reflective access. This means that while *most* externalist concepts of justification are compatible with reflective luck, not all of them are. Some kinds of justification, such as the kind required for reflective knowledge, will exclude *both* veritic and reflective luck. In the next section we will see another example of such a 'strong' form of justification. These strong forms of justification can be regarded as hybrid concepts of justification: concepts that share both internalist and externalist elements. I will then continue in the next chapter to discuss the prospects of such strong accounts of justification in more detail. For now, let us rest content by saying that at least at the most basic level, Sosa allows for justified belief that is nevertheless reflectively lucky. In the next section I discuss some immediate implications of our findings so far.

6.4 FACTIVE ACCOUNTS OF JUSTIFICATION

I will start the second part of this chapter by considering the implication of our findings for Zagzebski's well-known

claim that only factive accounts of justification are able to evade Gettier cases (1994).

As we saw in Chapter 4, Zagzebski argues that no non-factive account of justification will be immune from Gettier cases. By this she means that for every non-factive account of justification, we will be able to construct a case where this condition, together with the truth and belief condition, is satisfied, but in which the subject nevertheless fails to have knowledge.

What is problematic about Gettier cases, according to Zagzebski, is the following:

What generates the problem for JTB, then, is that an accident of bad luck is cancelled out by an accident of good luck. The right goal is reached, but only by chance. (Zagzebski, 1994, p. 66)

'The right goal' here is *truth*. So Gettier cases in general are cases where a true belief is formed but only by chance. In Chapter 4 I argued this means that Gettier cases necessarily involve veritic luck. Importantly, Zagzebski thinks this means that only accounts of justification that are *factive*—in the sense that having a justified belief entails that the belief is true—are immune to Gettier cases:

As long as the property that putatively converts true belief into knowledge is analysed in such a way that it is strongly linked with the truth, but does not guarantee it, it will always be possible to devise cases in which the link between such a property and the truth is broken but regained by accident. Such is the nature of Gettier cases. (Zagzebski, 1994, p. 69)

One of the reasons why Zagzebski's point is particularly pressing is that most philosophers would like to accommodate the possibility of false but nonetheless justified beliefs.

Most philosopher's would thus endorse a non-factive interpretation of justification. If Zagzebski is right, such justification conditions will never suffice (together with true belief) for knowledge.

Our findings suggest, however, that Zagzebski is mistaken. It is not the case that only factive accounts of justification can evade Gettier cases. What is problematic about Gettier cases is the presence of veritic luck. Every justification condition that eliminates this kind of luck will thus be immune to them. Justification that eliminates veritic luck need not entail truth. The elimination of veritic luck only requires that one could not have *easily* formed a false belief in the same way instead, not that one could not *possibly* have formed a false belief. Indeed, we have argued that properly understood, most prominent non-factive externalist accounts of justification *are* able to eliminate veritic luck, and thus are immune to Gettier cases.

It seems thus that Zagzebski is mistaken in her claim that factive justification is *necessary* to escape Gettier cases.¹² We may still wonder whether she is right in her claim that they would be sufficient. To test this claim, we consider two factive concepts of justification. These concepts have not been considered so far because, as we said above, the view that justification is factive is a minority view. Nevertheless, they are interesting for our present purposes.

¹² Zagzebski responds to a similar argument in her (1999, p. 103). She argues that while it will exclude Gettier-cases, an anti-luck condition on justification will be both uninformative and *ad hoc*. (Howard-snyder, Howard-snyder, & Feit, 2003) provide a response, which I take to be largely adequate. Such a condition is informative since it tells us something substantial about the general constraints on epistemic justification. It is not *ad hoc*, because it would be plausible even if no one thought of Gettier cases. Indeed, as we saw, the condition is implied by most prominent externalist concepts of justification, concepts that can be motivated without an appeal to Gettier cases.

The first approach we consider is Clayton Littlejohn's concept of justification (2012). According to Littlejohn, justification is factive:

FACTIVITY_J: You cannot justifiably believe p unless p is true. (2012, p. 122)

The second factive account of justification that we consider is Jonathan Sutton's account of justification:

KNOWLEDGE_J: A subject's belief that p is justified if and only if he knows that p .¹³ (2005)

That Littlejohn's account of justification is factive is clear. That Sutton's account is factive is clear if we make the assumption that knowledge is factive explicit. You cannot know what is false, and therefore, if you cannot justifiably believe p unless you know it, you cannot justifiably believe it unless it is true. Let us see if these factive accounts of justification eliminate veritic luck.

The argument is easier to make for Sutton's account of justification than for Littlejohn's. For, as we have been arguing throughout this dissertation, the main lesson to be drawn from Gettier's paper is that knowledge is incompatible with veritic luck. If knowledge is incompatible with veritic luck, and a justified belief is a belief that constitutes knowledge, as Sutton maintains, then a justified belief cannot be veritically lucky. So on Sutton's account, justification straightforwardly, by definition, eliminates veritic luck.

On the other hand, Factivity_J does not rule out Gettier cases. By definition, all Gettiered beliefs are true. So in this sense, the factivity of justification does not rule out the kind of luck present in Gettier cases. For that, justification needs to exclude veritic luck. Does that mean that Zagzebski is wrong when she claims that factivity is sufficient

¹³ See also (Haddock, 2010; Sutton, 2007; Williamson, 2016).

to evade Gettier-cases? I do not think it does, because the notion of factivity that Zagzebski uses is of the stronger kind that requires not just that every justified belief is true, but rather that one's justifiers *guarantee* that one's belief is true. What seems to be required by Zagzebski to evade Gettier cases is that it is impossible for one to form a false belief in the same, justified way. If this is the case, then of course one's belief will not be a matter of veritic luck. So, contrary to Littlejohn's weaker kind of factivity, this stronger notion of factive justification does seem to rule out veritic luck. Nevertheless, this means that in full generality, the claim that factive accounts of justification evade Gettier cases is wrong. As Factivity_J brings out, there can be factive accounts of justification that are compatible with veritic luck.

It should be noted, however, that this does not mean that Littlejohn's own account of justification is compatible with veritic luck. For Littlejohn does not consider Factivity_J to be *sufficient* for justification. For Littlejohn,

[w]hether you meet all of the requirements on doxastic justification, ... depends upon whether the mental states by virtue of which you meet the requirements on personal justification direct your attention to the facts that show that your relevant beliefs are correct. (Littlejohn, 2012, p. 240)

Personal justification, for Littlejohn, requires that you do not have any available defeaters for your belief and that it seems to you that the belief is correct. This is an internalist element of Littlejohn's account of justification, since whether it seems to you that your belief is correct, and whether you have any available defeaters "turns on what your mental life is like" (Littlejohn, 2012, p. 240).

However, what is primarily relevant for our present purposes is that besides personal justification, doxastic justi-

fication requires that the facts under consideration ‘show that your relevant beliefs are correct’. What does it take for something to show to someone that she is correct? Littlejohn notes that there has to be a subject, and something that is correct. Furthermore, there must be “something that puts the person and the facts in the right relation” (2012, p. 241). Interestingly for our purposes, the required relation is spelled in terms of luck:

Whether you are related in the appropriate way depends upon whether the connection between you and the facts is sufficiently non-accidental.
(Littlejohn, 2012, p. 241)

Thus, for Littlejohn, to be justified, you must stand in a non-accidental relation to the facts that make your belief true. What kind of luck must be excluded? Here Littlejohn draws on the distinction between intervening and environmental luck. He argues that only intervening luck undermines justification, although he admits his arguments are not entirely conclusive.¹⁴ In Chapter 4 we saw that intervening luck is a subspecies of veritic luck. At the very least we can thus say that Littlejohn’s account is designed to rule out an important variety of veritic luck.

Further, while Littlejohn does not go into the issue, it seems plausible to suppose that his account is incompatible with other forms of epistemically problematic luck as well. For one thing, he requires that justified beliefs be based on reasons that show that you are right. If your belief is formed in this way, it is hard to imagine how you

¹⁴ “If our intuitions about these [fake barn] cases are not to be trusted ... I have failed to show that justification is an externalist notion distinct from knowledge” (Littlejohn, 2012, p. 253). Since, at least for Littlejohn, it is clear that *knowledge* is incompatible with the environmental luck present in fake barn cases, this would entail that justification is incompatible with environmental luck as well. Cf. our discussion of Sosa’s account of justification above.

could easily have formed a false belief on the same basis. In any case, paradigmatic cases of veritically lucky beliefs such as lucky guesses and the beliefs of victims of the new evil demon will not be justified. Even if not conclusive, this provides support to the claim that Littlejohn's account of justification rules out other kinds of veritic luck besides the intervening luck present in Gettier cases.

Finally, Littlejohn's account of justification eliminates reflective luck because it requires that you are in possession of reasons that you take to show you are right. If you are in possession of such reasons, and you believe that you formed your belief on the basis of those reasons, it seems that from your perspective, you could not have easily formed a false belief in the same way. Littlejohn's account of justification is thus another example of a 'strong' account of justification that not only eliminates veritic but also reflective luck.

Let us recap. In this section we argued that factivity is neither necessary nor sufficient to eliminate veritic luck. This means that Zagzebski is wrong to say that only factive accounts of justification are able to evade Gettier cases. Nevertheless, we found that the factive accounts of Littlejohn and Sutton *are* able to evade Gettier cases. But that is because they are incompatible with veritic luck, not because of their factivity.

6.5 SAFETY AND LUCK

With the exception of a bare-bone factivity account of justification, all externalist accounts of justification considered eliminate veritic luck. This is striking, because it is usually thought that we need a condition on knowledge additional to justification in order to eliminate veritic luck. In this section, we will briefly switch our attention from justification to knowledge, and discuss one such condition.

According to Pritchard, Sosa and Williamson, knowledge should be *safe* (Pritchard, 2007; Sosa, 1999; Williamson, 2000, Ch. 5). Different as their overall views on knowledge are, they all seem to be committed to something like the following thesis:

SAFETY I: S knows that p only if not easily would S believe that p without it being the case that p .¹⁵

That knowledge is subject to such a safety condition is brought out by cases like the following. I form the belief that Donald Trump is going to save America on the basis of looking at his campaign posters. However, as campaign posters work, they would tell me this even if it was not in fact true. Suppose I am very gullible in this regard. Thus, I would easily form the same belief in the same way even it were false. My belief is thus unsafe, and this meshes with our intuition that, even if it turns out that my belief is true, I cannot on this basis alone claim to know that Donald Trump is going to save America.

Thus formulated, Safety I specifies a necessary condition on knowledge. Does it eliminate veritic luck? Even though safety is usually put forward specifically to eliminate veritic luck, as formulated, it will not do so. This is because what is at issue in the formulation of safety above is the modal profile of the belief that p *itself*. This means that beliefs in necessary truths, which are modally speaking, maximally 'robust' (you cannot possibly falsely believe a necessary proposition to be true), will be safe by default. But beliefs in necessary truths can be veritically lucky. On the present conception of veritic luck, what is important for veritic luck is not whether the belief *itself* could have easily been false, but rather that *the way the belief is produced* could have easily produced a false belief instead, whether

¹⁵ The formulation is based on Sosa's (1999, p. 142).

that belief is the same belief or a different belief produced in the same way. As formulated, safe belief might thus still be veritically lucky. Luckily, it is easy to provide a modified safety condition on knowledge:

SAFETY II: S knows that p only if not easily would S form a false belief instead.

This condition rules out veritic luck, since beliefs that are veritically lucky are produced in a way that could have easily produced a false belief instead, and this is just what is precluded by Safety II.

However, note that, as formulated, safety eliminates more luck than just *veritic* luck. Consider a case of evidential luck:

I luckily run into you on my way to my exam. As we discuss the exam, it turns out that I have been calculating a specific class of statistical problems using an unsuitable method. As the result of our talk, my handling of these cases is greatly improved.

Here, my belief will not satisfy Safety II because I could have easily still be using my faulty method, believing wrongly that certain answers to the exam questions are correct. But this seems to be the wrong result, in particular if safety is a necessary condition on knowledge. For I may still, if I actually used the right method, come to know the correct answers to my questions. So Safety II cannot be the kind of condition that is necessary for knowledge.

Not all formulations of safety fall prey to this problem, however. Consider Prichard's formulation:

SAFETY III: For all agents, ϕ , if an agent knows a contingent proposition ϕ , then, in nearly all (if not all) nearby possible worlds in which she

forms her belief about ϕ in the same way as she forms her belief in the actual world, that agent only believes that ϕ when ϕ is true. (Pritchard, 2005a, p. 163)

Here, nearby possible worlds represent easy possibilities. We note two things about this definition. First, in order to avoid the necessary safety of necessary truths, Pritchard's formulation needs to be modified.¹⁶ Second, however, Safety III explicitly references the belief-forming method used to produce the safe belief. Our example beliefs above would be safe, since in those cases, one could not have easily formed a false belief, *in the same way as one actually formed one's belief*.

On the basis of the points made above, it seems we need something like the following safety condition on knowledge:

SAFETY IV: S knows that p only if not easily would the method used to produce her belief result in a false belief instead.

This safety condition targets all and only veritic luck. If I am right, however, then externalists about justification do not need to pose any additional safety condition on knowledge. Externalist justification on its own would suffice to rule out veritic luck. As we have seen, the same does not hold for internalist concepts of justification. If the above authors are right in supposing that knowledge should be safe, then internalists about justification will need to endorse something like Safety IV as an additional safety condition on knowledge besides their preferred justification condition.

¹⁶ This fact is recognized by Prichard in his (Pritchard, 2012a).

6.6 THE INTERNALISM/EXTERNALISM-DEBATE REVISITED

In the previous chapter, I argued that the most prominent internalist accounts of justification eliminate reflective luck but not veritic luck. In this chapter, I argued that the converse is true for the most prominent externalist accounts of justification: externalist accounts eliminate veritic luck but not reflective luck. In this section I take a step back and review the consequences of these findings for the internalism/externalism-debate about epistemic justification.

Our findings indicate that internalist and externalist justification can be distinguished in terms of the kinds of luck they are compatible and incompatible with. Internalist justification, as we have seen, is incompatible with reflective luck but not with veritic luck, whereas externalist justification is incompatible with veritic luck but in principle compatible with reflective luck.¹⁷ My proposal as regards our view of the internalism/externalism-debate boils down to taking the incompatibility between justification and reflective and veritic luck as *essential* characteristics of internalism and externalism, respectively.¹⁸ That is, what unites internalists is a commitment to the claim that reflective luck should be eliminated in order for a belief to be justified, and I propose to take this joint commitment as the defining characteristic of internalist conceptions of justification. I characterize externalism about justification as any position that entails that veritic luck should be absent in order for a belief to be justified.

This proposal is not entirely new. In particular, Duncan Pritchard has claimed that we should understand externalist theories of *knowledge* as directed mainly at the elimina-

¹⁷ Even though we have seen externalist accounts that are incompatible with both kinds of luck.

¹⁸ I do not mean to commit myself to any metaphysical claim here. 'essential' should be understood loosely as 'unique' or 'defining'

tion of veritic luck (from knowledge), whereas the elimination of reflective luck seems to be more of an internalist preoccupation (Pritchard, 2005a). I take the present findings to provide further support for this claim.

While our findings thus support Pritchard's claim, they also go beyond it because his claim concerns internalist and externalist theories of *knowledge*, not justification. Since Pritchard defines an internalist account of knowledge as an account that features a necessary internalist kind of justification (and consequently, an externalist account as an account that does not feature such an internalist justification condition), one may suspect that such internalist accounts of knowledge are incompatible with reflective luck *because* internalist kinds of justification are. However, this tells us nothing about why externalist accounts of knowledge are incompatible with luck, for these may either feature an externalist justification condition, or *no justification condition at all*. Pritchard's claim thus does not reduce to mine.

That being said, our findings motivate a novel conception of the internalism/externalism-debate about epistemic justification as essentially a debate about which kinds of luck is ruled out if a belief is epistemically justified. Such a re-conceptualization is supported by the following two points.

First, if we want to argue that internalism and externalism can be distinguished in terms of the luck they eliminate, it is necessary that these characteristics are possessed by all and only the internalist and externalist theories of justification respectively. That is, the theory has to be extensionally adequate. In this chapter and the previous one, we provided support for this claim by arguing that the two main internalist theories of justification are committed to the claim that justification eliminates reflective luck, whereas main externalist theories of justification are committed to the claim that justification eliminates veritic luck. Further, the considered internalist concepts of justification

are compatible with veritic luck, and, with the exception of some 'strong' accounts of justification, our main externalist concepts of justification are compatible with reflective luck. While we have not considered all possible internalist or externalist accounts of justification, it seems that our criterion would appropriately classify the main accounts of justification in the literature.

Second, this way of formulating the issue dividing the internalists and externalists makes it at once clear that internalism and externalism are compatible in principle. For as the 'strong' accounts of justification that we have come across in this chapter show, it is possible for an account of justification to eliminate *both* reflective and veritic luck. Rather than envisaging the internalist and the externalist as diametrically opposed to each other, we get a picture where internalism and externalism are quite compatible. I believe such a picture does better justice to the actual hybrid positions in the field and promotes a cooperation between different sides of the debate.

My approach to the justification debate is in line with Alston's pluralistic approach outlined in his (2005). According to Alston, internalists and externalists about justification are talking past another and are instead just stressing different epistemic *desiderata*. Contrary to Alston, however, I do not think that we best abandon all talk of 'justification'. As the present study brings out, internalist and externalist concepts of justification share the property of excluding *some* form of luck from true belief. Contrary to Alston's claim, there is thus a neutral way of specifying the property about which internalists and externalists are disagreeing: justification eliminates luck from belief. Internalists and externalists disagree about the *kind* of luck that is relevant here.

Second, this way of delineating internalism and externalism has some important virtues over the way this is usually done. Consider the following examples:

Epistemic internalism is the view that a thinker's epistemic status depends wholly on matters which are 'internal' to that thinker, rather than at least partially on matters which are 'external' to her, such as her relations to her environment. Let epistemic externalism be the denial of epistemic internalism. (Brown, 2007, pp. 13-14)

Internalism asserts that justification is internally determined, whether by evidence possessed, or by coherence among beliefs, or by some other internal condition. Externalism about justification is readily understood as the denial that internal factors are sufficient. Something external has an independent role in justifying beliefs. Justification does not supervene on the internal alone. (Conee, 2004, p. 48)

Note that both these accounts try to delineate internalism about justification from externalism in theory-neutral terms. For example, both Brown and Conee are trying to provide characterization of internalism that does not favour accessibilism over mentalism. But this has the unfortunate consequence of making their characterizations quite unspecific. Surely internalism is the view that justification depends on the internal in *some* way, but this observation alone is rather trivial. We would like to know in *what* way justification depends on the internal for internalism. Of course, this differs between different internalist conceptions. But what is nice about the proposal to define internalism and externalism in terms of luck is that it is able to maintain a neutral stance towards the different forms of internalism and externalism, while at the same time saying something informative about their respective commitments.

In what way is our proposal informative? First, on our proposal, it is clear that what is required by internalists is that a justified belief is a belief that from the perspective of the subject is formed in a way that could not easily have produced a false belief instead. For externalists, on the other hand, what matters for justification is whether one's method of belief-formation could not easily have produced a false belief instead, according to the *objective*, external description of the situation, independently of what is reflectively accessible to the agent. This brings out a shared commitment to eliminating luck from justified belief, while at the same time making clear in what sense the *kinds* of luck that are eliminated are different.

Contrast this with the above two accounts of Brown and Conee. On Brown's account, hearing that a belief is internalistically justified conveys no more information than that this epistemic status depends wholly on matters which are 'internal' to that thinker. Hearing that a belief is externalistically justified is even less informative: all we can infer from that information is that it is not the case that its epistemic status depends on such internal matters. A similar critique can be levelled against Conee's classification: on his account internalistically justified beliefs are beliefs whose justificatory status is determined by some internal conditions, and externalistically justified beliefs are beliefs for which this does not hold. Thus, our proposal has an important advantage over these attempts in that it provides more content to the internalism/externalism distinction than is usually done, while maintaining the neutrality with respect to different internalist and externalist conceptions of justification.

A second—related—advantage of our proposal is that it identifies a positive commitment of externalism. For both Brown and Conee, externalism is characterized as the denial of internalism, but not much more is said on what, if anything, unites externalist conceptions of justification

besides their incompatibility with internalist ones. For all their characterizations allow, there may be *nothing* which is common to externalist conceptions of justification apart from their incompatibility with internalist ones. On our characterization, however, this is not the case. In this chapter I argued that all externalist conceptions of justification are committed to the positive claim that justification requires the elimination of veritic luck, and so our characterization makes clear that there is such a thing as an 'externalist commitment', one that goes beyond the mere rejection of internalism.

We thus see that an application of different concepts of luck to the internalism/externalism-debate constitutes progress in the sense that it enables us to provide a more informative delineation of internalism and externalism, and in doing so uncovers a positive commitment of externalism, rather than the mere rejection of internalism. This alone, I think, justifies our efforts to relate various kinds of justification and luck. But the increase in clarity about the current state of the debate between internalists and externalists is not the only reward generated by this project. As I will argue in the next chapter, our investigations point us towards a multi-level view of justification, a view that recognizes both the importance of eliminating veritic luck from knowledge, as well as reflective luck.

6.7 CONCLUDING REMARKS

In the first part of this chapter, I argued that prominent externalist accounts of justification are incompatible with veritic luck but compatible with reflective luck. In the second part, I considered some implications.

First, I discussed at some length Zagzebski's claims about the inescapability of Gettier cases for non-factive accounts of justification. We found that factive accounts of

justification are neither necessary nor sufficient to escape Gettier problems.

Second, I argued that if externalist concepts of justification are incompatible with veritic luck, this means that externalists about justification can do without a separate safety condition on knowledge. Since internalist concepts of justification do little to eliminate veritic luck, this will not be true for those who endorse an internalist concept of justification.

Finally, I argued that the findings from this chapter and the previous one provide the basis for a novel characterization of the internalism/externalism-debate about epistemic justification in terms of luck. I argued that this conception of the internalism/externalism-debate has some advantages over existing ways of characterizing the debate, most notably that it is more informative about the commonalities among internalist and externalist accounts respectively.

JUSTIFICATION AND LUCK

In the last two chapters, I argued that internalist and externalist accounts of justification can be characterized in terms of various kinds of epistemic luck. In particular, I argued that externalist concepts of justification eliminate veritic luck but not reflective luck, whereas the reverse is true for internalist concepts. So far, I have thus tried to bring out connections between existing theories of justification and various forms of luck. In this chapter, I will—*informed by these findings*—develop my own account of the relation between justification and luck. Specifically, I will draw a distinction between two forms of justification, which I will call Type I and Type II justification. Whereas Type I justification requires only the elimination of veritic luck, Type II justification *additionally* requires the elimination of reflective luck. Type I and Type II justification can be linked to Sosa’s different levels of knowledge: Type I justification is required for animal knowledge, Type II justification for the other levels.

If we define internalism, as we have done in Chapter 6, as the claim that justification excludes reflective luck but not veritic luck, then neither Type I nor Type II justification is of the internalist kind, since neither is compatible with veritic luck. But Type II justification does have the internalist element of excluding reflective luck, and thereby opens the door to a kind of justification that can be endorsed and defended. But it also has externalist elements in that veritic luck needs to be excluded as well. In this regard, the concept can be seen as a synthesis between internalism and externalism. Type II justification has high epistemic value, but is difficult to achieve, and for some agents out of reach

in principle. Type I justification is purely externalist and more easily achievable.

This chapter will be somewhat more speculative than the previous ones. But I feel that anyone who spends so much time analysing the relation between *other* accounts of justification and luck owes an exposition of his own views on the matter. That is what I attempt to do in this chapter. In any case, what we will see is that we can use our findings regarding the various forms of malignant epistemic luck, and their relation to internalism, externalism and various grades of knowledge to say quite a bit about the relation between luck and justification. Crucially, by distinguishing different forms of justification we can provide a more nuanced picture of the relation between justification and luck than that given by either internalists or externalists.

I only claim that the conditions developed in this chapter are necessary for justification, not that they are sufficient as well. What I will say is thus compatible with other, non-luck related, conditions on justification. In particular, I will argue that the anti-luck condition on Type I justification is derivative from a particular relation between this kind of justification and what I call 'veritic risk'. We can explain why Type I justification is incompatible with veritic luck in terms of the incompatibility between this kind of justification and veritic risk.

The chapter is structured as follows. In Section 7.1, I present my account of the relation between justification and luck. As said, I distinguish two kinds of justification, and present two corresponding anti-luck conditions. In the second part of the chapter, I discuss two possible problems for this way of thinking about justification. The first is that every false belief will trivially satisfy my condition for Type I justification. I therefore discuss the relation between false belief and luck in more detail in Section 7.2. I introduce a notion related to the notion of veritic luck, the notion of veritic *risk*. It will turn out that this notion allows

us to formulate conditions on Type I and Type II justification that are more general than our conditions in terms of veritic luck, which are entailed by these more general conditions. Our discussion of the relation between veritic luck and justification will also lead us to distinguish between good and bad veritic luck in Section 7.3. Perhaps surprisingly, it turns out that while Type I justification is incompatible with good veritic luck, it is compatible with bad veritic luck. Indeed, I will argue that false beliefs require the *presence* of bad veritic luck if they are to be justified.

In Section 7.4, I discuss a second potential problem for my way of thinking about the relation between justification and luck. This problem is specifically related to my condition for Type II justification. It concerns the fact that some people have argued that reflective luck is *ineliminable*. This would have sceptical consequences on the assumption that Type II justification is required for certain kinds of knowledge. I present some ways to evade this objection. In Section 7.5, I close with some concluding remarks.

7.1 TYPE I AND TYPE II JUSTIFICATION

In this section I distinguish between two kinds of justification: Type I justification and Type II justification. Type I justification requires the elimination of veritic luck, Type II justification the elimination of both veritic and reflective luck. What can be said for drawing this distinction?

In Chapter 4, we saw that Sosa distinguishes three kinds of knowledge: animal knowledge, reflective knowledge, and knowledge full well. I argued that the first of these three requires the elimination of veritic luck, whereas the other two require the elimination of both reflective and veritic luck. The distinction between the kinds of justification identified in this chapter can be seen as an extension of Sosa's theory. Whereas Type I justification is naturally

interpreted as the kind of justification required for animal knowledge, Type II justification is the kind of justification that could be a necessary condition on Sosa's higher grades of knowledge.

If Sosa is right, and there are different kinds of knowledge, it stands to reason that there are different kinds of justification as well. This is especially plausible on the standard view that the justification condition is *the* epistemic condition on knowledge.¹ If, as Sosa maintains, the different kinds of knowledge have different epistemic statuses, we might expect they are governed by distinct epistemic conditions. Indeed, in the previous chapter we saw that Sosa does indeed distinguish between different grades of justification, one of which I argued requires just the absence of veritic luck other one requiring the absence of both veritic and reflective luck. On my picture, the former one would be a specific kind of Type I justification, the latter an example of Type II justification.

While the present view on justification thus fits naturally with Sosa's picture on the different kinds of knowledge, it does not depend on it. We have seen that reliabilism entails the absence of veritic luck just as Sosa's concept of justification does. So reliabilism may be considered a form of Type I justification as well. Similarly, we argued in the previous chapter that Littlejohn's concept of epistemic justification eliminates both veritic and reflective luck, so this concept may be considered a kind of Type II justification.

Furthermore, we may well distinguish different kinds of justification even if we allow for only one grade of knowl-

¹ The standard view is something like this. There are three conditions on knowledge: justification, truth, and belief. Of these three, justification is the epistemic, truth the metaphysical and belief the psychological condition. Perhaps an additional anti-Gettier condition is needed, though as we have seen some externalist conceptions of the justification condition are able to evade Gettier cases.

edge.² This would give rise to the question of which kind of justification is required for knowledge. On a liberal view, only Type I justification is required for knowledge. This would enable us to say that young children and non-human animals could have knowledge, even if they lack the capacity to achieve Type II justification. A more conservative view would say that Type II justification is required for knowledge, and would consequently entail that knowledge is more rare.

Independently of the above considerations, we can argue for the plausibility of the distinction between Type I and Type II justification by drawing on the notion of epistemic value. We have seen that both reflective and veritic luck are epistemically problematic or undesirable, in the sense that we accord a higher epistemic value to beliefs that are not veritically lucky than to beliefs that are, and more value still to beliefs that are neither reflectively nor veritically lucky. A 'strong' kind of justification that would eliminate both reflective and veritic luck would thus be highly valuable from the epistemic point of view.³ One can see the present chapter in part as an exploration of the prospects for such a strong concept of justification.

Be that as it may, it seems that most of our beliefs are held without the explicit reflective support required to eliminate reflective luck. If we want to allow for the justification of some of those unreflectively held beliefs, we thus need to distinguish a second, weaker, kind of epistemic justification as well. My dual account of justification

2 I use 'grades' and 'kinds' of knowledge interchangeably. No commitment to *natural* kinds is intended.

3 Note the present distinction between a strong and a weak type of justification is *not* the same as Goldman's distinction under the same heading (Goldman, 1988). Roughly, Goldman's strong concept of justification is a purely externalist one, whereas his weak concept of justification is a purely internalist one. The purely externalist concept of justification would be more akin to our *weak* type of justification, where our strong type of justification include both internalist and externalist elements.

will allow for these beliefs to be justified in one sense, but still acknowledge there is a different sense of justification in which they fall short.

Why not acknowledge a third kind of justification that only requires the elimination of reflective luck? That is, why think the elimination of veritic luck is necessary for *every* kind of justification? My reason for this is twofold. First, I believe that all forms of justification should be truth conducive, and only non-veritically lucky beliefs satisfy this criterion. Epistemic justification is something that is valuable from the epistemic point of view. What is this point of view? We saw in Chapter 2 that it is plausible that truth is among our ultimate epistemic aims, and in Chapter 6 that such a truth-conducive concept of justification requires the elimination of veritic luck. This means that only non-veritically lucky belief will be conducive to our ultimate epistemic aim, and thus, have the kind of value required for epistemic justification. On this conception, no form of epistemic justification is compatible with veritic luck.

Second, much of the appeal of extant theories of justification that *are* compatible with veritic luck can be explained away by drawing a distinction between primary and derivative epistemic norms, or between justifications and excuses, as Williamson and Littlejohn respectively have argued recently (Littlejohn, n.d.; Williamson, n.d.). When your belief is veritically lucky, the most you can hope for is an epistemic excuse, not a justification. Because both justification and excuse remove *blame*, it is easy to see how people can confuse veritically lucky but blameless beliefs with justified ones.

The distinction between (epistemic) justification and excuse can be supported by the claim above about the necessary truth-conduciveness of epistemic justification, but is not entailed by it. Justification could be truth-conducive even in the absence of a distinction between justifications

and excuses. In that case, the appeal of theories compatible with veritic luck would be harder to explain, but so much the worse, one might say, for such theories. Similarly, one might draw a distinction between justification and excuse without maintaining that justification is truth-conducive. Nevertheless, the claims fit together nicely: part of the explanation *why* internalist criteria for justification only generate epistemic excuses, not justifications, could be that justification must be truth-conducive.

Let us say a little more about Type I and Type II justification respectively. Type I justification, as I define it, is subject to the following anti-luck condition:

C₁: Belief B(p) of agent S is Type I justified only if the belief is not veritically lucky.

On the definition I adopted in Chapter 4, a belief is veritically lucky if and only if it is true, but the method that produced the belief could have easily produced a false belief instead. There are three ways in which a belief can fail to meet this criterion: it can be false, it can be produced in a way that could have easily produced false belief, or both. This means that all false beliefs will meet C₁. Lest we want to say that all false beliefs are justified, we had thus better not regard C₁ as a *sufficient* condition on justification. We will return to this issue in the next section.

Whether a belief-forming method could have easily produced a false belief is neither reflectively accessible nor a mental state, so C₁ is an externalist condition. This makes Type I justification an externalist concept of justification. In the previous chapter we saw that C₁ is implied by many important externalist theories of justification. In this sense, the claim that justification requires the elimination of veritic luck, although rarely noted, should be agreeable to externalists of different kinds.

I understand Type I justification to be compatible with reflective luck. Thus, Norman the clairvoyant, whose be-

iefs are reflectively, but not veritically lucky, could still be Type I justified in his belief. Whether he is so justified will depend on whether there are additional criteria on Type I justification besides C_1 , and whether Norman meets those criteria if so. But nothing in the present account excludes his belief as counting as Type I justified.

As I just pointed out, our condition on Type I justification should be agreeable to externalists. Somewhat more controversial is my claim that there is another kind of justification, Type II justification, that requires the elimination of *both* reflective *and* veritic luck. It is subject to the following necessary condition:

C_2 : Belief $B(p)$ of agent S is Type II justified only if it is neither veritically nor reflectively lucky.

What this means is that for a belief to be Type II justified, it must not only be the case that the belief is true and produced in a way that could not easily have produced a false belief instead, but also that this is so, *given the information accessible to the subject*. Thus, in order to evade reflective luck, the subject's reflective perspective must make it likely that her belief is produced in a way that could not easily have produced false belief instead. If she lacks such reasons, her belief will be subject to reflective luck.

In virtue of incorporating an anti-reflective luck condition, C_2 will satisfy the constraint we saw in Chapter 5 was implied by most prominent internalist theories of justification. In virtue of incorporating an anti-veritic luck condition, C_2 will satisfy the constraint we saw in Chapter 6 was implied by most prominent *externalist* theories of justification. Type II justification will thus satisfy some internalist, as well as some externalist constraints on justification. In this sense, it will be a hybrid view, a view of justification that has both internalist and externalist elements.

Since the absence of reflective and veritic luck implies the absence of veritic luck, a belief that meets our criterion for Type II justification will automatically meet our criterion for Type I justification. Not every kind of Type II justification will imply Type I justification, however, since our anti-luck conditions are only necessary for their respective kinds of justification, and the additional criteria for both Type I and Type II justification may be spelled out in different ways.

Both C₁ and C₂ treat justification as an absolute term. This is customary in the literature that deals with definitions of epistemic justification, but in fact it is highly implausible that justification is an all-or-nothing matter. One can be more or less justified. So I should say something about how my account is compatible with such a gradual picture of justification.

As we have seen in Chapter 3, the notion of luck is gradual too. A particular event can be more or less lucky. So while stated in absolute terms, both C₁ and C₂ can easily be modified to accommodate the above point. We can simply say that the *more* veritic luck to which a belief is subject, the less it will be Type I justified. The degree of Type II justification, on the other hand, will depend on both on the degree of reflective and on the degree of veritic luck. From a gradual perspective, we can read our conditions on extreme degrees of justification. One is *completely* Type I justified only if one's belief is subject to no veritic luck at all, and one is completely Type II justified only in the absence of any reflective or veritic luck. These extreme degrees of justification may be hard to achieve. But that would not be a problem, since it is plausible that for most of our purposes something less than complete justification will be sufficient.

This raises the interesting issue of how degrees of veritic and reflective luck are to be weighed against each other. Is it as bad, epistemically speaking, for a belief to be subject

to veritic luck as it is for the belief to be subject to an equal degree of reflective luck? I do not claim to have a precise answer to this question. Nor do I think that such a precise answer can be given, since, as we saw in Chapter 3, it is doubtful that we can put an exact measure on the 'easiness' of a particular possibility, and thus on the degree of luck (either of the veritic or of the reflective kind) to which it is subject.

Nevertheless, some general remarks are possible. A high degree of veritic luck will preclude a high degree of Type II justification, even if the belief is subject only to a low degree of reflective luck. This is in line with our comments earlier, that no belief is justified (either Type I or Type II) if it is not formed in a way that is truth-conducive. The fact that it *seems* to be produced in such a way will only generate an epistemic excuse, not a justification. So in order to achieve a high degree of Type II justification, a belief must be subject to a relatively low degree of veritic luck.

Similarly, a belief cannot achieve a high degree of Type II justification if it is subject to a high degree of reflective luck. If from your perspective, your belief is produced in a way that could easily have produced a false belief, then you will not be highly Type II justified, even if in fact your method could not easily have produced false belief at all. Type II justification requires a perspective from which your methods are safe from error. Of course, when your belief is subject to a high degree of reflective luck, but only to a low degree of veritic luck, it may still be Type I justified.

In order to achieve a high degree of Type II justification, your belief can thus neither be subject to a high degree of veritic luck, nor to a high degree of reflective luck. In this sense, Type II justification is quite demanding. It is precluded by either a substantial degree of veritic or of reflective luck.

This being said, while our conditions are thus compatible with a gradual account of justification, for reasons of

simplicity, we will focus in this chapter mainly on the absolute versions of our justification conditions.

With this sketch of Type I and Type II justification in mind, I continue to discuss two possible problems for my account. First, in the next section, I consider the objection that all false beliefs trivially satisfy our condition for Type I justification. This need not be a problem, since as we said, C_1 is only intended as a necessary condition on justification. However, the fact that false beliefs trivially satisfy our condition makes our anti-luck condition on justification somewhat uninformative. It is not able, for example, to explain the difference between false justified and false unjustified beliefs. To explain this difference we turn in the next section to the related notion of veritic *risk*.

7.2 FALSE BELIEFS AND VERITIC RISK

In this section I argue for a condition on justification in terms of the notion of *risk*. This notion, related to luck, is more informative than our anti-luck condition on Type I justification in that it specifies a criterion not only for distinguishing justified from unjustified *true* belief, but between justified and unjustified *false* belief as well.

First things first, however. Not all false beliefs are justified. This should not require much explanation, so let it suffice to name a few examples. If you have been given roughly the education I have been given, but you believe the earth is flat, this belief is presumably both false and unjustified. If you are bad at mathematics, and come to believe that 72 divided by 4 is 17, presumably this belief will be both false and unjustified. If you hastily generalize from a few incompetent metal bands to the belief that all metal bands are incompetent, presumably this belief is both false and unjustified. And we can of course think of many more such examples.

On the other hand, while this is sometimes denied, it seems there can be false justified beliefs.⁴ This possibility can be, and has been, accommodated by both internalists and externalists: even highly reliable methods may sometimes produce false belief, and even the best reflectively accessible evidence can point in the wrong direction.

What this means is that any plausible account of justification should allow for the possibility of false justified beliefs, as well as false unjustified beliefs. Ideally, such an account should furthermore be able to explain the difference.

Let us first see how our account of Type I justification fares in this regard. At the end of this section I will come back to Type II justification.

As we said, Type I justification requires the absence of veritic luck. Recall that a belief is veritically lucky if and only if it is true but produced in a way that could have easily resulted in a false belief. As we noted above, no false beliefs are veritically lucky according to this definition, since no false belief is true. This means that all false beliefs will satisfy our condition for Type I justification. If we regard C_1 as sufficient for Type I justification, this means that all false beliefs will be Type I justified. Since this is clearly wrong, we cannot take C_1 as a sufficient condition for justification.

It is thus of paramount importance to stress that the account of Type I justification sketched above is just a *partial* account. The anti-veritic luck condition is *necessary* for Type I justification only. No belief can be justified if it is veritically lucky. This will solve our problem, for now additional criteria for justification may be formulated that deal with the possibility of false unjustified beliefs. On this read-

⁴ Epistemologists who think false beliefs cannot be justified include Littlejohn (2012), Sutton (2007), and Williamson (2016).

ing, the account will not automatically license all false beliefs as Type I justified.⁵

Note that a similar problem arises for anti-luck theories of *knowledge*. Consider a theory of knowledge that states that a belief constitutes knowledge if and only if it is not veritically lucky.⁶ If the absence of veritic luck were *sufficient* for knowledge, all false beliefs would constitute knowledge. This of course is incompatible with the factivity of knowledge. In fact, *no* false beliefs constitute knowledge. So an anti-luck condition all by itself will do neither for an adequate account of justification nor for an adequate account of knowledge.

In the case of knowledge, however, the problem is easily fixed by adding a separate *truth*-condition. On this view, knowledge requires truth *and* the absence of veritic luck.⁷ Such a solution will not do in the case of justification, however, since we saw it is plausible that there are false but still justified beliefs. So we cannot just stipulate that justified beliefs must be true. What we need is a condition on justification that allows us to distinguish between true and false justified beliefs on the one hand, and true and false unjustified beliefs on the other in a uniform way. Preferably, such a condition would also explain *why* Type I justification is incompatible with veritic luck. As it turns out, we can give such a condition in terms of a notion closely related to the notion of veritic luck: the notion of veritic *risk*.

5 (Howard-snyder et al., 2003) identifies the same problem for anti-luck conditions on justification. The condition on justification proposed there differs from the one proposed below. No distinction is made between Type I and Type II justification, and no mention is made of veritic risk.

6 One may think of the theory of Pritchard (2005a) in this regard. Since then Pritchard has modified his account of knowledge, however. See (Pritchard, 2012a).

7 One could think of the safety-conditions endorsed by Pritchard (2005a, Ch. 6), Sosa (1999) and Williamson (2000, Ch. 5) as conjunctions of truth and anti-veritic luck conditions.

Let us define the following notion of veritic risk:

VERITIC RISK: Belief $B(p)$ of agent S is veritically risky iff $B(p)$ is produced in a way that could have easily produced a false belief.

The notions of veritic risk and veritic luck are very similar. The only difference between the two is that veritically lucky beliefs are required to be true, but veritically risky beliefs are not. This means that every true non-veritically risky belief will be non-veritically lucky, and every true veritically risky belief will be veritically lucky as well.⁸ The notions are not equivalent, however. No false belief is veritically lucky, but some will be veritically risky. We then specify the following anti-veritic risk condition for Type I justification:

C_3 : Belief $B(p)$ of agent S is Type I justified only if $B(p)$ is non-risky.

Adopting C_3 as a necessary condition on Type I justification solves our problem, since there clearly are risky false beliefs. For example, my belief that it will rain tomorrow on the basis of a simple guess may both be false and risky, since guessing can easily produce false belief. Moreover, C_3 still seems to capture the important externalist that justified beliefs are produced in a way that is truth-conducive. Simple process reliabilism, as a paradigm externalist conception of justification, would entail, for example, that justified beliefs are non-risky, at least on the modal interpretation of reliability. On this interpretation, reliably produced beliefs are produced in a way that not only results in a high ratio of true beliefs over false beliefs

⁸ See also Pritchard's discussion of the relation between luck and risk in general in (Pritchard, 2015). It is not clear, however, whether Pritchard's concept of risk is the same as ours, since he does not provide a definition (cf. Pritchard, 2015, pp. 149–152).

in this world, but in nearby worlds as well. On this interpretation, reliable belief-forming methods could not easily have produced false beliefs, and thus result in beliefs that are non-risky. So an account of Type I justification in terms of risk rather than luck would still seem to be compatible with prime externalist accounts of justification.

Relatedly, most prominent externalist accounts of justification are incompatible with veritic luck. At least for those true beliefs that satisfy these externalist criteria, this entails that veritic risk will be absent as well. If the satisfaction of externalist criteria is compatible with the absence of veritic risk in the case of true belief, it stands to reason that it is compatible with the absence of veritic risk in the case of false belief as well. Even the reliable methods may sometimes produce false belief, but the fact that these methods are reliable indicates that they could not *easily* have produced false belief, and so the beliefs formed on their basis will not be subject to high degrees of veritic risk.

Since C_3 entails C_1 , we can explain *why* Type I justified beliefs cannot be veritically lucky by referring to the fact that they cannot be veritically risky. Viewed in this light, the incompatibility between justification and luck would be derivative from the incompatibility between justification and risk. At first sight, thus, C_3 provides the condition on justification that we were looking for: it explains why Type I justification is incompatible with veritic luck and draws a distinction between justified and unjustified belief in a uniform manner.

As we saw in Chapter 3, how easily something is possible depends on two factors: i) the modal distance between the nearest world where the event occurs and the actual world, ii) the relative proportion of worlds where the event occurs compared to all (relevant) worlds. This means that how easily one's method of belief-formation could produce a false belief will also depend on two factors: i) the modal distance between the nearest world where one's methods

of belief-formation produces a false belief and the actual world, and ii) the relative proportion of nearby possibilities where one's method of belief-formation produces a false belief. This in turn means that the veritic risk to which a belief is subject will also be a function of these last two factors. In particular, the closer to the actual world the nearest world is where one's method of belief-formation produces a false belief, the more risky the belief. And similarly, the higher the proportion of relevant worlds where one's method of belief to produces a false belief, the more risky the resulting belief.

It should be noted, that this entails that false beliefs are always somewhat risky. For they are formed, by definition, by a method that actually produces false belief. Since there is no world more similar to the actual world than the actual world itself, this means in the modal distance sense of easy possibility, false beliefs are *always* produced in a way that could easily have produced false belief. They will thus always be *somewhat* risky. If any degree of risk is enough to prohibit justification, then no false belief will be justified. While, as I mentioned in the introduction, there are people who endorse this view, I would like to maintain the possibility of false justified belief.

We can accommodate the possibility of false justified belief in two ways. First, we can simply specify a threshold value for the degree of risk to which a justified belief may still be subject. Even if a method produces a false belief in the actual world, the relative proportion of nearby worlds where it does so may be limited, such that the degree of veritic risk to which the belief is subject may still be low. On the threshold view, such a low degree of risk may still be compatible with justification. A second strategy, which I think is the more plausible one of the two, stresses the fact, noted in the introduction of this chapter, that justification itself is plausibly regarded as a gradual notion. On this view, a false belief that is subject only to a low degree

of risk is simply Type I justified to a high, but non-maximal degree. As an example, take my belief, based purely on the long odds, that my lottery ticket will not win. Suppose this belief is false (I will in fact win the lottery). On our current understanding of veritic risk, my belief is somewhat risky. After all, in actuality the method that I used produced a false belief. But we can still say, on the present understanding of risk, that the risk is quite low, for there are proportionally *many more* easy possibilities for this method to produce a true belief. Therefore, my belief may still be highly justified, even if not *maximally* so.

We can thus draw on the earlier made distinction between maximal and non-maximal justification. Even if false beliefs will never be *maximally* Type I justified, since always subject to *some* risk, they may nevertheless be Type I justified to a high degree. So not all false beliefs will be unjustified. C3 fulfils our conditions for a plausible and informative necessary condition on epistemic justification. It draws the distinction between justified and unjustified belief in a principled way, and can explain why justified beliefs cannot be veritically lucky, for no belief that is produced in a way that could not easily produce false belief (and is therefore *not* subject to veritic risk) is true and produced in a way that *could* easily have produced a false belief (and thus subject to veritic luck).

I have amended my condition for Type I justification in order to explain the difference between justified and unjustified beliefs, irrespective of whether the belief in question is true or false. I similarly update my condition for Type II justification. But before doing this, we may wonder whether an adequate account of Type II justification would have to involve a notion of *reflective risk* instead of its present reference to reflective luck. As it turns out, however, where the notions of veritic luck and veritic risk are importantly distinct, this is not the case for the notions of reflective luck and reflective risk.

This is so because presumably, there are no beliefs, which from the reflective perspective of the subject herself are false. For it is hard to imagine how you could simultaneously believe something and believe that belief to be false. In any case, as Moore famously argued, expressing such a state of mind would be absurd (Moore 1993, pp. 207–212). Believing that the Eiffel tower is in Rome seems to be the same thing as believing that it is true that the Eiffel tower is in Rome. If you do believe that the Eiffel tower is in Rome, you think that it is true that the Eiffel tower is in Rome, and this implies that you do not think it is false that the Eiffel tower is in Rome.

Let us define reflective risk analogously to how we defined veritic risk:

REFLECTIVE RISK: Belief $B(p)$ of agent S is reflectively risky iff, given only the information reflectively accessible to the agent, $B(p)$ is produced in a way that could have easily produced a false belief.

If the only way that the agent can hold the belief $B(p)$ for it to be the case that from her perspective, $B(p)$ is true, then this definition is equivalent to our definition of reflective luck:

REFLECTIVE LUCK₂: S 's belief that p is reflectively lucky if and only if, given the information reflectively accessible to S , it is a matter of luck that the method S used to form her belief that p produced a true belief.

That this is so can be seen more clearly once we plug in our general account of luck:

REFLECTIVE LUCK₃: S 's belief that p is reflectively lucky if and only if, given the information reflectively accessible to S , her belief that p

is true but the method S used to form her belief that p could have easily produced a false belief instead.

No special account of the relation between reflective luck and Type II justified belief is needed. We thus arrive at the following definition of Type II justification:

C₄: The belief B(p) of agent S is Type II justified only if the following two conditions are satisfied:

- i B(p) is not veritically risky
- ii B(p) is not reflectively lucky.

I have argued that C₄ provides a plausible necessary condition on Type II justification. I have not claimed sufficiency. Even so, based on the consideration of the relation between luck and justification alone, it seems we can say quite a lot already about the requirements for epistemic justification.

7.3 GOOD AND BAD VERITIC LUCK

In the previous section I argued that Type I justification requires the elimination of veritic risk, and suggested that this explains why such justification is incompatible with veritic luck. We can say a bit more, however, about the relation between Type I justification and veritic luck. In this section, I argue for a nuanced picture of this relation by distinguishing between two kinds of veritic luck: *good* veritic luck and *bad* veritic luck.

To start, in Chapter 3, we argued that luck comes in good and bad varieties, and attributed the difference to the significance of the event in question: if an event is a case of good luck for an agent, then it has positive significance to

that agent, and if it is a matter of bad luck, then it has negative significance. Now we will apply this distinction to the case of veritic luck.

As we argued in Chapter 4, the significance important for epistemic luck is *epistemic* significance. We argued that from the epistemic point of view, true beliefs are always positively significant to epistemic agents. However, in a similar vein we may argue that false beliefs are always negatively significant to epistemic agents. From the epistemic point of view, false beliefs have negative value. On this assumption, then, we may distinguish between cases where it is a matter of luck that one's method of belief-formation produced a *true* belief, from cases where it is a matter of luck that it produced a *false* belief. So far, we have reserved the term 'veritic luck' only for instances of the first kind, but both kinds can be called forms of veritic luck in the sense that in both cases it is a matter of luck that a belief with a particular truth-value is produced. I will from now on call instances of the former kind instances of 'good' veritic luck, and instances of the latter kind instances of 'bad' veritic luck.

Interestingly, this distinction allows for a more nuanced picture regarding the relation between veritic luck and Type I justification in general. For while, as we have argued all along, this kind of justification is incompatible with good veritic luck, it is quite compatible with bad veritic luck.

Consider first what is required for bad veritic luck to occur. If a belief is a case of bad veritic luck, it is false but produced in a way that could have easily produced true belief instead. We can produce many such cases that intuitively seem to be justified. Consider first my lottery belief above. It is false, but given the way I formed it, I could have easily produced a true belief. Intuitively, my belief that I will win the lottery is justified. Further, suppose that I read in a reputable encyclopaedia that the atomic number of gold

is 76. To check whether this source is reliable, I look up its credentials. Various sources tell me that regarding atomic numbers, my source is the authority. I then form the corresponding belief about the atomic number of gold. Alas, the atomic number of gold in reality is 79, so my belief is false. The encyclopaedia simply contained an embarrassing misprint on one of its pages. Given the way I formed my belief (say, consulting a reputable source, and cross-checking its reliability), I could easily have formed a true belief instead. So my belief is a case of bad veritic luck. Still, it seems, my belief is justified. Finally, I may form a belief based on my reliably faculty of vision. Since this faculty is not perfectly reliably, it may still produce false belief. If it does, then my false belief will be a case of bad veritic luck: easily might my method have produced a true belief instead. Yet this does not seem to detract from my belief's justification.

Note that we are still considering Type I justification here, the kind of justification required for animal knowledge. We may have this kind of justification, even if our false beliefs are cases of bad veritic luck. Indeed, I will go even further and assert that if they are to be Type I justified, the presence of bad veritic luck is *required* for false beliefs. For consider what is the case if false beliefs are *not* subject to bad veritic luck. In that case, they are false but produced in a way that could *not* easily have produced true belief. I submit it is highly implausible that belief-forming methods for which the production of a true belief is only a remote possibility can generate justified beliefs, even if it is justification of the most basic kind.

It seems thus, that while Type I justification is incompatible with good veritic luck, it is compatible with bad veritic luck, and may even require it if the belief in question is false. This does not require an additional condition on Type I justification, since the requirement for the belief to be non-veritically risky suffices to guarantee bad veritic luck in case of false belief. For such a belief to be

non-*veritically* risky, it must be produced in a way that could not very easily produce false belief. Even if in actuality the method produced false belief, this suffices to ensure that there will be many close possibilities where the method produces true belief, and thus to ensure substantial degrees of bad veritic luck.

While of course *epistemic* justification is different from that of moral or pragmatic justification, it is interesting that those other kinds of justification seem to have similar relations to good and bad forms of luck as the one proposed here. For example, it seems doubtful that one's action could be pragmatically justified if it is just a matter of (good) luck that one's action had ultimately useful consequences. On the other hand, it seems that one can *only* be pragmatically justified in acting a certain way if it is a matter of bad luck that one's action did *not* have ultimately useful consequences.

In ethics too, there seems some plausibility (assuming a simple consequentialist view, at least) to the claim that one cannot be morally justified if it is just a matter of good luck that one's action has good consequences, but that one could still be morally justified if the fact that one's action had *bad* consequences is a matter of (bad) luck.⁹

More generally even, our ascriptions of *praise* and *blame* behave in the same way. We praise actions that have a positive outcome only if these outcomes are not due to luck, whereas we would praise actions that have a negative outcome only if this negative outcome *is* due to luck. Thus, we would not praise the CEO for preserving the rainforest if the method she followed could have easily destroyed it. Her actions might still be praised, on the other hand, if

⁹ See the classical debate on moral luck between Thomas Nagel and Bernard Williams (1976). Compare also the so-called 'Knobe effect', the effect that people's attributions of intentionality vary with the moral significance of the effect (Knobe, 2003; Nadelhoffer, 2004). See also (Hindriks, 2011).

the forest was destroyed, but not if her method could not have easily preserved it, thus, only on the condition that was a matter of bad luck that her method had the result of destroying the forest.

In general then, it seems to be the case that we positively evaluate methods or actions only when their positive results are not due to luck, and conversely, that the positive evaluation of actions that have negative consequences requires these consequences to be a matter of bad luck. The same, I submit, is the case when we evaluate beliefs. False beliefs can only be positively evaluated when their falsity is a matter of bad veritic luck.

By distinguishing between good and bad veritic luck, we thus arrive at a more nuanced picture of the relation between Type I justification and veritic luck. It is incompatible with some subspecies of veritic luck, but quite compatible with others. This seems to hold not only for our notion of Type I justification, but to the evaluation of intentional action in general.

In this section we discussed a problem that is particularly salient for Type I justification. In the next section we discuss an objection to our account of Type II justification. The objection concerns the extent to which our account is subject to certain sceptical worries.

7.4 SCEPTICISM

In this section we discuss the relation between Type II justification and scepticism. Before we start, it should be noted that our dual account of justification is compatible with various accounts of the relation between justification and knowledge. The problem that we will discuss here targets accounts of knowledge that require our conditions for Type II justification to be satisfied. By now, we have seen that Sosa's accounts of reflective knowledge and knowl-

edge full well are examples of such accounts. The problem is that on our account of Type II justification, such knowledge is impossible, because the conditions for Type II justification can never be satisfied.

Why think that Type II justification is impossible? Here I focus on one particular argument to this effect—an argument that can be put directly in terms of luck. The argument has the following simple structure:

- P₁ Type II justification requires the elimination of reflective luck.
- P₂ The elimination of reflective luck is impossible.
- C Therefore, Type II justification is impossible.

The argument is valid. P₁ is true by definition of Type II justification. The crucial premise here is thus P₂. P₂ has been defended by Duncan Pritchard (2005a). The argument Pritchard provides for P₂ has the following structure:

- P₃ The elimination of reflective luck requires us to have reflectively accessible grounds for our beliefs that favour them over sceptical scenarios.
- P₄ We cannot have reflectively accessible grounds for our beliefs that favour them over sceptical scenarios.
- C The elimination of reflective luck is impossible.

Why does Pritchard think that the elimination of reflective luck would require us to have reflectively accessible grounds for our beliefs that favour them over sceptical scenarios? The idea here is that if a world w is to be counted as far-off, on the ordering relevant for reflective luck—that is, the ordering from the perspective of the subject herself—the subject needs to have access to reasons that support the fact that w is not the world, or a world very much alike, the

world she is actually living in. If she does not have access to such reasons, then from her own perspective, the subject might as well be living in w . So w will be easily possible from her perspective. This holds true for all worlds, and thus also for worlds that contain sceptical scenarios.

By stipulation, in sceptical scenarios, most of our beliefs are produced in a way that could easily produce false belief. It follows that if the grounds for our beliefs do not favour them over sceptical scenarios, our belief-forming methods could have easily produced false belief, from our perspective. Thus, if the reasons a subject has for her belief do not favour them over sceptical scenarios, then given these reasons only, her belief will be subject to reflective luck.

Why think P_4 is true? That is, why think that we cannot have reflectively accessible grounds for our beliefs that favour them over sceptical scenarios? Here the reason is simply that sceptical scenarios are *stipulated* to be reflectively identical to 'normal' situations, situations in which everything is as it appears to us. So, in the brain in a vat case, it is stipulated that the brain in a vat has access to exactly the same reflectively accessible evidence as we have—assuming that we are not ourselves brains in vats. Similarly, when the New Evil Demon is introduced, it is stipulated that this demon is able to manipulate the world in such a way that there is no reflectively accessible difference between us and the demon's victims.

If our reflectively accessible evidence is the same between the normal and the sceptical case, then that evidence will not favour the normal case over the sceptical case. You do not have evidence for being in the normal case when you would have exactly the same evidence in the sceptical scenario. The evidence you have is equally compatible with both situations and thus does not provide you reason to favour one over the other.

It thus seems that our reflectively accessible evidence does not favour our present beliefs over sceptical scenarios, and that our beliefs will be ineliminably subject to reflective luck.

If Pritchard is right and the elimination of reflective luck is impossible, then this will have consequences for the present theory of justification. If Type II justification requires the elimination of reflective luck, then we have to conclude that Type II justification is impossible. If there are kinds of knowledge that require this kind of justification, such as Sosa's higher levels of knowledge, this will in turn have sceptical consequences for these kinds of knowledge. Something like this seems to be Pritchard's view. In his words, the ineliminability of reflective luck means that we will be condemned to a state of perpetual *Epistemic Angst*, an inability to attain a kind of knowledge for which we can take full responsibility as epistemic agents (Pritchard, 2005a, p. 247). But it would certainly be a bitter pill to swallow. As we have argued, reflective knowledge is epistemically more valuable than animal knowledge.¹⁰ It would be a pity if a kind of knowledge with such high epistemic value would be out of reach in principle.

Do we need to go along with this pessimistic conclusion? There are at least two ways to resist it.

First, we can attack premise P₃. As we saw, the plausibility of this premise depends on the easy possibility of sceptical scenarios, given one's reflectively accessible evidence only. From this it is concluded that one's beliefs could have easily been false from one's perspective, and are thus reflectively lucky. But this inference is invalid because not all worlds that are easily possible from one's perspective are relevant for reflective luck. Only those worlds where the

¹⁰ This does not imply anything on how to two kinds of knowledge compare in terms of other kinds of value. A given item of animal knowledge may have much more *pragmatic* value than an item of reflective knowledge.

subject uses the same method as that she *believes* she uses in the actual world are. This means that sceptical worlds where I form my belief on the basis of a method different from the one I believe I used in the actual world will not be relevant. So if I believe that the way I formed my belief involves faculties that take in light reflected by external objects to form my belief that there is a tree there, then my belief may well fail to be reflectively lucky, since from my perspective *this* method could not have easily produced false belief. Nearby sceptical scenarios are irrelevant, since in these easily possible worlds I use a method different from the one I believe I am using.

On this concept of reflective luck, the inference from the easy possibility of sceptical scenarios to the presence of reflective luck is invalid, and P₃ is thus false. This would allow us to evade the conclusion that Type II knowledge is impossible.¹¹

A second strategy for resisting this conclusion attacks premise P₄ of Pritchard's argument. As we saw, the plausibility of this premise depends on the fact that by definition, we are in possession of the same reflectively accessible evidence in sceptical scenarios as that we are in more 'normal' situations. The cogency of this possibility is far from clear, however, on some recent views regarding the nature of perceptual evidence. In particular, there seems to be some traction for disjunctivist views on perceptual evidence that imply one's evidence is *not* the same between the normal and the sceptical case (Byrne & Logue, 2009).

According to a common disjunctivist view, the reason that one has for believing there is a tree yonder is *that one*

¹¹ This requires that we individuate our belief-forming methods on the basis of factors that are external to the subject. I take it that many of our beliefs about our own belief-forming methods involve such external factors. This does not yet mean that reflective luck itself depends on the external environment. Beliefs *about* external factors are themselves still appropriately within the subject's perspective to count as internal.

sees there is a tree yonder. Disjunctivists argue that this fact is one's evidence for belief and appropriately (that is, reflectively accessible) in one's possession in the normal case. In the sceptical case, however, one does not see the tree but merely *seems* to see the tree. Thus, there is no fact of seeing the tree that can be one's evidence or appropriately in one's possession. What this means is that on this conception of the reasons in one's possession, one *can* have reasons that favour one's beliefs over sceptical scenarios. If there is no tree, I cannot see that there is a tree over yonder, so the fact that I see a tree certainly favours the normal scenario, in which there is a tree, over the sceptical scenario in which there is no such tree.

Drawing on a disjunctivist conception of evidence we can thus resist P₄, and thereby resist the conclusion that our reflectively accessible evidence does not favour normal scenarios over sceptical ones. The disjunctivist conception of evidence is controversial, however. But even on a less controversial conception of one's reflectively accessible evidence, one could still use our first strategy for denying Pritchard's conclusion. What this shows is that there are still multiple ways to account for the possibility of Type II justification in the face of Pritchard's argument to the contrary. We do not need to withdraw into a state of perpetual epistemic *angst*.

7.5 CONCLUDING REMARKS

In this chapter I presented my own views on the relation between justification and luck. I distinguished two kinds of justification: Type I and Type II justification. Type I justification requires the elimination of veritic luck. Type II justification requires the elimination of both reflective and veritic luck.

But this is not all we can say about the relation between luck in justification. We saw that we can provide a more general condition on Type I and Type II justification in terms of the notion of veritic *risk*. Both forms of justification are incompatible with veritic risk, which explains why both forms of justification are incompatible with veritic luck, but it is also able to explain the difference between false justified and false unjustified beliefs.

Furthermore, I argued that we can distinguish between good and bad veritic luck. Type I and Type II justification are incompatible with good veritic luck, but quite compatible with bad veritic luck. Indeed, when the belief is false, it seems it must be subject to bad veritic luck if it is to be Type I or Type II justified at all.

Finally, we considered the threat of scepticism that faces accounts of justification that require the absence of reflective luck, such as our account of Type II justification, as well as some ways to respond to these objections. What we see is that at least in the face of the present objections, the distinction between Type I and Type II justification turns out to be quite defensible.

CONCLUSIONS

Luck is everywhere. It hides in the small crevices between the actual world and worlds much like it. Although we rarely realize it, it shows itself when we cross the street unharmed, or open the fridge after a long day's work to find some food still left in there. It shows itself more clearly when we meet someone special, or when we witness something extraordinary. Much of our personal lives are mired with luck, and so are our professional ones. We are lucky when our home price offer is accepted, or when we are offered a job. But luck also comes in negative varieties. When our papers are rejected, we lament our bad luck. Similarly when we lose our loved ones in a tragic accident, or when we lose our fortune at the gambling tables.

Our lives are thus riddled with luck. But besides this practical role in our daily lives, luck has a theoretical role to play as well. For as recent work has shown, luck is intimately related to some of the most central concepts in philosophy. In ethics, for example, it is argued that luck is incompatible with free will, and undermines moral responsibility (Levy, 2011b). In epistemology, it is argued that luck undermines knowledge. In order to substantiate these claims, it is vitally important that more work is done on the nature of luck and the relation between luck and these various other concepts so important for philosophy.

In this study, my aim was to contribute to this project insofar as it concerns the role of luck in epistemology. Some important work has already been done. In particular, Duncan Pritchard has been influential in distinguishing various forms of luck, and arguing that knowledge is compatible with some of them but not with others. What is central to

Pritchard's account of luck is that lucky events have a certain modal profile. That is, a lucky event is an event that occurs in the actual world, but fails to occur in a relevant set of nearby possible worlds. In this sense, all lucky events could have *easily* failed to occur. Such a modal account of luck, I argued, is superior to rival alternative accounts, formulated in terms of probability or lack of control over the event on the part of the agent.

But Pritchard's account of luck is not quite right in how it handles the various ways in which luck comes in *degrees*. For Pritchard, degree of luck is determined by the modal distance between the nearest possible world where the relevant initial conditions for an event are the same as in the actual world, but in which the event fails to obtain, and the actual world. I argued that this view has some counter-intuitive consequences. To evade these, we should acknowledge that degree of luck also depends on both the *significance* of the event, and on the *proportion* of nearby possible worlds where the relevant initial conditions for the event are the same as in the actual world, but in which the event fails to obtain. Once we do so, it becomes clear that the degree of luck to which an event is subject will depend on subjective factors in more than one way. In particular, the degree of luck to which an event is subject will depend not just on how we classify the event in question, or how we choose to partition the relevant modal space, but also on the event's significance to us.

If we consider the relation between knowledge and luck, it quickly becomes apparent that not all cases of luck are incompatible with the possession of knowledge. For example, the fact that it is a matter of luck that the proposition believed by the agent is true need not necessarily preclude the agent from knowing that proposition. Similarly, when she is lucky to acquire some evidence, she may still come to know on its basis. But there are forms of luck that do preclude knowledge. The most well-known of these is veritic

luck, the kind of luck that is in play when the method that the agent used to form her true belief could have easily produced a false belief instead. More controversial is the compatibility of knowledge with reflective luck, the kind of luck that is in play when *from the agent's perspective*, her belief-forming method could have easily produced a false belief instead of a true one. I argued for a nuanced picture, drawing on Sosa's distinction between animal knowledge and reflective knowledge. Whereas animal knowledge is compatible with reflective luck, reflective knowledge is not.

Traditionally, knowledge has been regarded to amount to justified, true belief. While many take Gettier to have shown that this simple definition fails in its generality, the picture remains influential. In particular, it is widely held that a belief needs to be justified if it is to count as knowledge. But what does this epistemic justification amount to? This question gives rise to one of the most important and at the same time intractable debates in contemporary epistemology: the internalism/externalism-debate. According to internalists, epistemic justification supervenes on facts that are 'internal', in some way or other, to the believing subject. Externalists deny this. This, of course, is not a very informative characterization. I argued that we can provide a better one by focussing on the relation between various accounts of justification and luck. In particular, I argued that whereas internalist concepts of justification tend to be compatible with veritic luck, but incompatible with reflective luck, the converse is true for externalist concepts. What this means is that we can characterize the internalism/externalism-debate in terms of luck: internalism is the view that justification is incompatible with reflective luck (but compatible with veritic luck), externalism the view that justification is incompatible with veritic luck (but compatible with reflective luck). This account identifies core commitments for both internalism and externalism and should be

preferred over the standard way of characterizing the internalism/externalism-debate.

Our findings regarding the relation between internalism, externalism and epistemic luck also support a nuanced story about the import of Gettier cases. For it is commonly accepted that what prohibits knowledge in Gettier cases is the presence of veritic luck. On the assumption that externalist concepts are incompatible with such luck, it follows that Gettier cases provide a counterexample to the tripartite analysis of knowledge only if we interpret the justification condition in an *internalist* way. Externalists should simply deny that Gettier-victims are justified in believing as they do.

Finally, I investigated the prospects of a dual account of justification. Type I justification, the kind of justification required for our most basic kind of knowledge, requires the elimination of veritic luck. Type II justification, the kind of justification required for the higher grades of knowledge, requires the elimination of both veritic *and* reflective luck. I argued that an account of Type I justification in terms of the absence of luck cannot explain the difference between justified and unjustified false beliefs, since all false beliefs are non-lucky by definition. But I argued that we can provide a more informative condition on the requirements of Type I justification in terms of the notion of epistemic *risk*. This notion allows us to provide a uniform explanation of the difference between Type I justified and unjustified belief, while at the same time explaining why Type I justification is incompatible with veritic luck.

In this study my aim was to provide a contribution to anti-luck epistemology by investigating the relation between luck and epistemic justification. Our findings indicate that there are important relations between justification (of either the internalist or externalist variety) and reflective and veritic luck. This is significant, because the received opinion in the literature seems to be that external-

ist concepts of justification fail to rule out veritic luck (e.g. Harper, 1996), whereas internalist concepts of justification fail to rule out reflective luck (e.g. Pritchard, 2005a). I have argued the opposite.

Once we see clearly the relation between internalism, externalism and the various forms of malignant epistemic luck, it becomes possible to see the internalism/externalism-debate in a new light. Knowledge is incompatible with luck, and we want our beliefs to be justified because this will guarantee that they satisfy this criterion. So far, this is a standard interpretation of the function of justification (e.g. Booth, 2011). Our findings make this story more nuanced. We can distinguish different forms of knowledge, and these different forms bear different relations to various forms of luck. In order to exclude these different forms of luck, we need different forms of justification. In particular, the kind of knowledge that we aspire to in science, the kind of knowledge that can be explicitly endorsed and defended—reflective knowledge—requires the absence of both veritic and reflective luck, and therefore Type II justification. Animal knowledge, which constitutes the vast bulk of our knowledge, and is the kind of knowledge that helps us get around in the world with near miraculous efficiency, requires just the elimination of veritic luck, and therefore Type I justification. Both forms of justification are legitimately called forms of *justification*, in that they present necessary epistemic conditions for various forms of knowledge. Internalists and externalists have, however, been partly talking past each other because they fail to make the distinction between different grades of justification and knowledge, with different anti-luck requirements for each grade.

This story is compatible with the insistence that sometimes justification requires the *presence* of some form of veritic luck. For we can distinguish between good and bad varieties of veritic luck. All forms of justification are incom-

patible with good veritic luck, but compatible with bad veritic luck. Indeed, if they are to be justified at all, false beliefs *must* be subject to bad veritic luck. An agent cannot be justified if her belief is formed in a way for which the production of true belief is only a remote possibility. Our methods need to be more reliable than that if they are to provide justification. So veritic luck is relevant for justification in more than one way.

In this study I aimed to provide a foundation for the investigation of the relation between justification and luck. I have only considered what I take to be the most prominent concepts of justification, and the defence of my own view regarding the relation between justification and luck has been somewhat tentative. So much more work can still be done. Nevertheless, the present study illustrates both the importance and fruitfulness of the study of luck in epistemology.

BIBLIOGRAPHY

- Alston, W. P. (1985). Concepts Of Epistemic Justification. *The Monist*, 68(1), 57–89.
- Alston, W. P. (1988a). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2, 257–299.
- Alston, W. P. (1988b). An Internalist Externalism. *Synthese*, 74(3), 265–283.
- Alston, W. P. (1989). *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.
- Alston, W. P. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Cambridge: Cambridge University Press.
- Armstrong, D. M. (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Audi, R. (1994). Dispositional Beliefs and Dispositions to Believe. *Noûs*, 28(4), 419–434.
- Ayer, A. J. (1940). *The Foundations Of Empirical Knowledge*. London: Macmillan.
- Bach, K. (1985). A Rationale for Reliabilism. *The Monist*, 68(2), 246–263.
- Baehr, J. S. (2011). *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford: Oxford University Press.
- Ballantyne, N. (2014). Does Luck Have a Place in Epistemology? *Synthese*, 191(7), 1391–1407.
- Baumann, P. (2014). No Luck With Knowledge? On a Dogma of Epistemology. *Philosophy and Phenomenological Research*, 89(3), 523–551.
- Beaney, M. (2016). Analysis. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Sum. 2016

- ed.). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/sum2016/entries/analysis/>
- Bergmann, M. (2006). *Justification without Awareness: A Defense of Epistemic Externalism*. Oxford: Oxford University Press.
- BonJour, L. (1980). Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy*, 5(1), 53–73.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- BonJour, L. (2002). Internalism and Externalism. In *The oxford handbook of epistemology*. New York, NY: Oxford University Press.
- BonJour, L., & Sosa, E. (2003). *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Hoboken, NJ: Wiley-Blackwell.
- Booth, A. R. (2011). The Theory of Epistemic Justification and the Theory of Knowledge: A Divorce. *Erkenntnis*, 75(1), 37–43.
- Brandom, R. (1998). Insights and Blindspots of Reliabilism. *The Monist*, 81(3), 371–392.
- Brown, J. (2007). Externalism in Mind and Epistemology. In S. Goldberg (Ed.), *Internalism and externalism in semantics and epistemology*. Oxford: Oxford University Press.
- Byrne, A., & Logue, H. (2009). *Disjunctivism: Contemporary Readings*. Cambridge, MA: MIT Press.
- Cappelen, H. (2012). *Philosophy Without Intuitions*. Oxford: Oxford University Press.
- Carter, J. A., & Pritchard, D. (2015). Knowledge How and Epistemic Luck. *Noûs*, 49(3), 440–453.
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. New York, NY: Oxford University Press.
- Chalmers, D. J. (2002). Does Conceivability Entail Possi-

- bility? In T. S. Gendler & J. Hawthorne (Eds.), *Conceivability and possibility*. Oxford: Oxford University Press.
- Chalmers, D. J. (2010). The Two-Dimensional Argument Against Materialism. In *The character of consciousness*. Oxford: Oxford University Press.
- Chalmers, D. J. (2014). Intuitions in Philosophy: A Minimal Defense. *Philosophical Studies*, 171(3), 535–544.
- Chappell, S. G. (2013). Plato on Knowledge in the Theaetetus. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Win. 2013 ed.). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/win2013/entries/plato-theaetetus/>
- Chisholm, R. M. (1977). *Theory of Knowledge* (2nd ed.). Englewood Cliffs, NJ: Prentice-Hall.
- Coffman, E. J. (2010). Misleading Dispositions and the Value of Knowledge. *Journal of Philosophical Research*, 35, 241–258.
- Conee, E. (2004). Externalism, Internalism, and Skepticism. *Philosophical issues*, 14(1), 78–90.
- Conee, E., & Feldman, R. (1998). The Generality Problem for Reliabilism. *Philosophical Studies*, 89(1), 1–29.
- Conee, E., & Feldman, R. (2001). Internalism Defended. *American Philosophical Quarterly*, 38(1), 1–18.
- Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in Epistemology*. Oxford, NY: Oxford University Press.
- Dancy, J. (1985). *An Introduction to Contemporary Epistemology*. Oxford, NY: Blackwell Publishing.
- David, M. (2001). Truth as the Epistemic Goal. In M. Steup (Ed.), *Knowledge, truth, and duty*. New York, NY: Oxford University Press.
- David, M. (2013). Truth as the Primary Epistemic Goal: A Working Hypothesis. In E. Steup, Matthias, Turri, John, Sosa (Ed.), *Contemporary debates in epistemology* (2nd ed.). John Wiley & Sons.

- Davies, M. (2005). *A Frequency Dictionary of Spanish: Core Vocabulary for Learners*. New York, NY: Routledge.
- de Grefte, J. (2017). Epistemic Justification and Epistemic Luck. *Synthese*, 1–16. Retrieved from <http://dx.doi.org/10.1007/s11229-016-1306-7>
- de Moivre, A. (1738). *The Doctrine of Chances: Or, A Method of Calculating the Probabilities of Events in Play* (second ed.). New York, NY: Chelsea Publishing Company.
- Dennett, D. C. (1980). The Milk of Human Intentionality. *Behavioral and Brain Sciences*, 3(3), 428–430.
- Dennett, D. C. (2013). *Intuition Pumps and Other Tools for Thinking*. New York, NY: WW Norton & Company.
- Descartes, R. (1931). *The Philosophical Works of Descartes Rendered Into English* (Vol. 77; E. S. Haldane & G. R. T. Ross, Eds.) (No. 2). Cambridge: Cambridge University Press.
- Dutant, J., & Dorsch, F. (Eds.). (n.d.). *The New Evil Demon Problem*. Oxford University Press.
- Dutilh Novaes, C., & Reck, E. (2017). Carnapian Explanation, Formalisms as Cognitive Tools, and the Paradox of Adequate Formalization. *Synthese*, 194(1), 195–215.
- Edgington, D. (1985). The Paradox of Knowability. *Mind*, 94(376), 557–568.
- Engel, M. (1992a). Is Epistemic Luck Compatible with Knowledge? *The Southern Journal of Philosophy*, 30(2), 59–75.
- Engel, M. (1992b). Personal and Doxastic Justification in Epistemology. *Philosophical Studies*, 67(2), 133–150.
- Engel, M. (2015). *Epistemic Luck*. Retrieved 2015-12-15, from www.iep.utm.edu/epi-luck/
- Feldman, R. (2004). Having Evidence. In *Evidentialism: Essays in epistemology*. Oxford, NY: Oxford University Press.
- Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48(1), 15–34.

- Fumerton, R. (1988). The Internalism/Externalism Controversy. In K. B. Wray (Ed.), *Knowledge and inquiry: Readings in epistemology*. Peterborough: Broadview Press.
- Gettier, E. L. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121–123.
- Ginet, C. (1975). *Knowledge, Perception and Memory*. Dordrecht: D. Reidel Publishing Company.
- Goldberg, S. C. (2015). Epistemic Entitlement and Luck. *Philosophy and Phenomenological Research*, 91(2), 273–302.
- Goldman, A. I. (1976). Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, 73(20), 771–791.
- Goldman, A. I. (1979). What is Justified Belief? In G. S. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: D. Reidel Publishing Company.
- Goldman, A. I. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. I. (1988). Strong and Weak Justification. *Philosophical Perspectives*, 2, 51–69.
- Goldman, A. I. (1994). Naturalistic Epistemology and Reliabilism. *Midwest Studies in Philosophy*(19), 301–320.
- Goldman, A. I. (1999). Internalism Exposed. *Journal of Philosophy*, 96(6), 271–293.
- Goldman, A. I. (2015). Reliabilism, Veritism, and Epistemic Consequentialism. *Episteme*, 12(02), 131–143.
- Goldman, A. I., & Olsson, E. J. (2009). Reliabilism and the Value of Knowledge. In *Epistemic value*. Oxford: Oxford University Press.
- Greco, J. (1993). Virtues and Vices of Virtue Epistemology. *Canadian Journal of Philosophy*, 23(3), 413–432.
- Greco, J. (1995). A Second Paradox Concerning Responsibility and Luck. *Metaphilosophy*, 26(1-2), 81–96.
- Greco, J. (1999). Agent Reliabilism. *Nous*, 33, 273–296.
- Greco, J. (2000). *Putting Sceptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge: Cambridge University Press.

- Haddock, A. (2010). Knowledge and Action. In D. Pritchard, A. Millar, & A. Haddock (Eds.), *The nature and value of knowledge: Three investigations*. Oxford, NY: Oxford University Press.
- Hájek, A. (2007). The Reference Class Problem Is Your Problem Too. *Synthese*, 156(3), 563–585.
- Hales, S. D. (2016). Why Every Theory of Luck is Wrong. *Noûs*, 50(3), 490–508.
- Hall, B. J. (1993). On Epistemic Luck. *Southern Journal of Philosophy*, 32(1), 79–84.
- Harper, W. (1996). Knowledge and Luck. *Southern Journal of Philosophy*, 34(3), 273–283.
- Hazlett, A. (2010). The Myth of Factive Verbs. *Philosophy and Phenomenological Research*, 80(3), 497–522.
- Heller, M. (1999). The Proper Role for Contextualism in an Anti-Luck Epistemology. *Nous*, 33(s13), 115–129.
- Hiller, A., & Neta, R. (2006). Safety and Epistemic Luck. *Synthese*, 158(3), 303–313.
- Hindriks, F. (2011). Control, Intentional Action, and Moral Responsibility. *Philosophical Psychology*, 24(6), 787–801.
- Holliday, W. H. (2014). Partiality and Adjointness in Modal Logic. *Advances in Modal Logic*, 10, 313–332.
- Howard-snyder, D., Howard-snyder, F., & Feit, N. (2003). Infallibilism and Gettier's Legacy. *Philosophy and Phenomenological Research*, 66(2), 304–327.
- Humberstone, I. L. (1981). From Worlds to Possibilities. *Journal of Philosophical Logic*, 10(3), 313–339.
- Ichikawa, J. J., & Steup, M. (2017). The Analysis of Knowledge. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Spr. 2017 ed.). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/spr2017/entries/knowledge-analysis/>
- Jackson, F. (2000). *The Role of Conceptual Analysis*. In *From metaphysics to ethics*. Oxford: Oxford University

- Press.
- James, W. (1897). *The Will to Believe*. In *The will to believe and other essays in popular philosophy*. New York, NY: Longmans, Green, and Co.
- Knobe, J. (2003). Intentional Action in Folk Psychology: An Experimental Investigation. *Philosophical Psychology*, 16(2), 309–324.
- Korcz, K. A. (2000). The Causal-Doxastic Theory of the Basing Relation. *Canadian Journal of Philosophy*, 30(4), 525–550.
- Kripke, S. A. (1980). *Naming and Necessity*. Oxford: Blackwell.
- Kvanvig, J. L. (2003). Propositionalism and the Perspectival Character of Justification. *American Philosophical Quarterly*, 40(1), 3–17.
- Kvanvig, J. L. (2013). Truth Is Not the Primary Epistemic Goal. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed.). Chichester: Wiley-Blackwell.
- Kvanvig, J. L., & Menzel, C. (1990). The Basic Notion of Justification. *Philosophical Studies*, 59(3), 235–261.
- Lackey, J. (2006). Pritchard's Epistemic Luck. *The Philosophical Quarterly*, 56(223), 284–289.
- Lackey, J. (2008). What Luck is Not. *Australasian Journal of Philosophy*, 86(2), 255–267.
- Lakatos, I. (1976). *Proofs and Refutations: The Logic of Mathematical Discovery* (J. Worrall & E. Zahar, Eds.). Cambridge: Cambridge University Press.
- Latus, A. (2003). Constitutive Luck. *Metaphilosophy*, 34(4), 460–475.
- Lehrer, K., & Cohen, S. (1983). Justification, Truth, and Coherence. *Synthese*, 55(2), 191–207.
- Leite, A. (2008). Believing One's Reasons are Good. *Synthese*, 161(3), 419–441.
- Levy, N. (2011a). An Account of Luck. In *Hard luck*. Oxford: Oxford University Press.

- Levy, N. (2011b). *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. Oxford: Oxford University Press.
- Lewis, D. K. (1973). *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Littlejohn, C. (n.d.). A Plea for Epistemic Excuses. In F. D. Julien Dutant (Ed.), *The new evil demon problem*. Oxford University Press.
- Littlejohn, C. (2009). The Externalist's Demon. *Canadian Journal of Philosophy*, 39(3), 399–434.
- Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge University Press.
- MasterRussian.com. (2016). *Top 500 Russian Verbs*. Retrieved 2016-12-12, from http://masterrussian.com/vocabulary/common_{_}verbs.htm
- Matilal, B. K. (1986). *Perception: An Essay on Classical Indian Theories of Knowledge*. Oxford: Oxford University Press.
- McCain, K. (2014). *Evidentialism and Epistemic Justification*. Hoboken, NJ: Routledge.
- McDowell, J. (1994). *Mind and World*. Cambridge: Harvard University Press.
- McHugh, C. (2012). The Truth Norm of Belief. *Pacific Philosophical Quarterly*, 93(1), 8–30.
- Menzies, P. (1998). Possibility and Conceivability: A Response-Dependent Account of their Connections. In R. Casati (Ed.), *European review of philosophy, volume 3: Response-dependence*. Stanford: Csl Publications.
- Nadelhoffer, T. (2004). The Butler Problem Revisited. *Analysis*, 64(283), 277–284.
- Neta, R. (2007). Propositional Justification, Evidence, and the Cost of Error. *Philosophical Issues*, 17, 197–216.
- Pappas, G. (2014). Internalist vs. Externalist Conceptions of Epistemic Justification. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Fall 2014 ed.). Stanford, CA: Metaphysics Re-

- search Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/fall2014/entries/justep-intext/>
- Peacocke, C. (1999). *Being Known*. Oxford: Oxford University Press.
- Plantinga, A. (1993a). *Warrant and Proper Function*. Oxford: Oxford University Press.
- Plantinga, A. (1993b). *Warrant: The Current Debate*. Oxford: Oxford University Press.
- Plato. (1973). *Theaetetus* (J. McDowell, Ed.). Oxford: Clarendon Press.
- Pollock, J. L., & Cruz, J. (1999). *Contemporary Theories of Knowledge*. Lanham, MD: Rowman & Littlefield Publishers.
- Poston, T. (2016). *Internalism and Externalism in Epistemology*. Retrieved 2017-04-06, from <http://www.iep.utm.edu/int-ext/>
- Price, H. H. (1969). *Belief*. London: Allen and Unwin.
- Pritchard, D. (2004). Epistemic luck. *Journal of Philosophical Research*, 29, 193–222.
- Pritchard, D. (2005a). *Epistemic luck*. New York, NY: Oxford University Press.
- Pritchard, D. (2005b). Scepticism, Epistemic Luck, and Epistemic Angst. *Australasian Journal of Philosophy*, 83(2), 185–205.
- Pritchard, D. (2007). Anti-Luck Epistemology. *Synthese*, 158(3), 277–297.
- Pritchard, D. (2008a). I - Radical Scepticism, Epistemic Luck, and Epistemic Value. *Aristotelian Society Supplementary Volume*, 82(1), 19–41.
- Pritchard, D. (2008b). Knowledge, Luck and Lotteries. In V. Hendricks (Ed.), *New waves in epistemology* (pp. 1–22). Basingstoke: Palgrave Macmillan.
- Pritchard, D. (2010). Knowledge and Understanding. In D. Pritchard, A. Haddock, & A. Millar (Eds.), *The nature and value of knowledge: Three investigations*. Oxford

- University Press.
- Pritchard, D. (2012a). Anti-Luck Virtue Epistemology. *Journal of Philosophy*, 109(3), 247–279.
- Pritchard, D. (2012b). *Epistemological Disjunctivism*. Oxford: Oxford University Press.
- Pritchard, D. (2014). The Modal Account of Luck. *Metaphilosophy*, 45(4-5), 594–619.
- Pritchard, D. (2015). The Modal Account of Luck. In D. Pritchard & L. J. Whittington (Eds.), *The philosophy of luck*. Malden, MA: Wiley-Blackwell.
- Pryor, J. (2001). Highlights of Recent Epistemology. *British Journal for the Philosophy of Science*, 52(1), 95–124.
- Putnam, H. (1975). *Mind, Language and Reality*. Cambridge: Cambridge University Press.
- Radford, C. (1966). Knowledge: By Examples. *Analysis*, 27(1), 1–11.
- Reichenbach, H. (1949). *The theory of probability, an inquiry into the logical and mathematical foundations of the calculus of probability*. Berkeley: University of California Press.
- Rescher, N. (2001). *Luck: The Brilliant Randomness of Everyday Life*. Pittsburgh, PA: University of Pittsburgh Press.
- Rescher, N. (2014). The Machinations of Luck. *Metaphilosophy*, 45(4-5), 620–626.
- Riggs, W. D. (2007). Why Epistemologists are so Down on Their Luck. *Synthese*, 158(3), 329–344.
- Riggs, W. D. (2009). Luck, Knowledge, and Control. In *Epistemic value*. Oxford: Oxford University Press.
- Rumfitt, I. (2015). *The Boundary Stones of Thought: An Essay in the Philosophy of Logic*. Oxford: Oxford University Press.
- Russell, B. (1948). *Human Knowledge: Its Scope and Limits*. New York, NY: Simon and Schuster.
- Sainsbury, M. R. (1997). Easy Possibilities. *Philosophy and Phenomenological Research*, 57(4), 907–919.

- Schwitzgebel, E. (2015). Belief. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Summer2015 ed.). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/sum2015/entries/belief/>
- Searle, J. R. (1992). *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.
- Sennett, J. (1992). Toward a Compatibility Theory for Internalist and Externalist Epistemologies. *Philosophy and Phenomenological Research*, 52(3), 641–655.
- Simion, M., Kelp, C., & Ghijsen, H. (2016). Norms of Belief. *Philosophical issues*, 26(1), 375–392.
- Smith, M. (2010). What Else Justification Could Be. *Noûs*, 1(1), 10–31.
- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.
- Sosa, E. (1985). The Coherence of Virtue and the Virtue of Coherence. *Synthese*, 64(1), 3–28.
- Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.
- Sosa, E. (1993). Proper Functionalism and Virtue Epistemology. *Noûs*, 27, 51–65.
- Sosa, E. (1999). How to Defeat Opposition to Moore. *Nous*, 33(s13), 141–153.
- Sosa, E. (2007). A Virtue Epistemology. In *A virtue epistemology*. Oxford: Oxford University Press.
- Sosa, E. (2009a). Knowing Full Well: The Normativity of Beliefs as Performances. *Philosophical Studies*, 142(1), 5–15.
- Sosa, E. (2009b). Philosophical Skepticism and Externalist Epistemology. In *Reflective knowledge: Apt belief and reflective knowledge, volume ii*. Oxford: Oxford University Press.
- Sosa, E. (2009c). *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. Oxford: Oxford University

- Press.
- Sosa, E. (2009d). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Oxford University Press.
- Sosa, E. (2010). *Knowing Full Well*. Princeton, NJ: Princeton University Press.
- Sosa, E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Stalnaker, R. C. (1968). A Theory of Conditionals. In N. Rescher (Ed.), *Studies in logical theory* (Vol. 2). Oxford: Blackwell.
- Stanley, J., & Williamson, T. (2013). Knowing How. *Journal of Philosophy*, 98(8), 411–444.
- Statman, D. (1991). Moral and Epistemic Luck. *Ratio*, 4(2), 146–156.
- Sturgeon, S. (1993). The Gettier Problem. *Analysis*, 53(3), 156–164.
- Sutton, J. (2005). Stick to what You Know. *Nous*, 39(3), 359–396.
- Sutton, J. (2007). *Without Justification*. Cambridge, MA: MIT Press.
- Swain, M. (1979). Justification and the Basis of Belief. In G. Pappas (Ed.), *Justification and knowledge*. Boston: D. Reidel.
- Swain, M. (1981). *Reasons and Knowledge*. Ithaca, NY: Cornell University Press.
- Swinburne, R. (2001). Introduction. In *Epistemic justification* (pp. 1–19). Oxford: Oxford University Press.
- Teigen, K. H. (2005). *When a Small Difference Makes a Big Difference: Counterfactual Thinking and Luck*. New York, NY: Routledge.
- Turri, J. (2010). On the Relationship between Propositional and Doxastic Justification. *Philosophy and Phenomenological Research*, 80(2), 312–326.
- Unger, P. (1968). An analysis of Factual Knowledge. *The Journal of Philosophy*, 65(6), 157–170.

- Weatherson, B. (2014). Centrality and Marginalisation. *Philosophical Studies*, 171(3), 517–533.
- Wedgwood, R. (2002). Internalism Explained. *Philosophy and Phenomenological Research*, 65(2), 349–369.
- Weinberg, J. M., Nichols, S., & Stich, S. (2001). Normativity and Epistemic Intuitions. *Philosophical Topics*, 29(1-2).
What can the Oxford English Corpus tell us about the English language? (n.d.). Retrieved 2016-11-30, from <https://en.oxforddictionaries.com/explore/what-can-corpus-tell-us-about-language>
- Whittington, L. J. (2016). Luck, Knowledge and Value. *Synthese*, 193(6), 1615–1633.
- Williams, B. A. O. (1992). Introduction. In B. A. O. Williams (Ed.), *Theaetetus*. Indianapolis, IN: Hackett Publishing Company.
- Williams, B. A. O., & Nagel, T. (1976). Moral Luck. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 50, 115–151.
- Williamson, T. (n.d.). Justifications, Excuses, and Sceptical Scenarios. In J. Dutant & F. Dorsch (Eds.), *The new evil demon problem*. Oxford University Press.
- Williamson, T. (1994). *Vagueness*. New York, NY: Routledge.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford, NY: Oxford University Press.
- Williamson, T. (2009). Probability and Danger. *The Amherst Lecture in Philosophy*, 4, 1–35. Retrieved from <http://www.amherstlecture.org/williamson2009/>
- Williamson, T. (2016). Knowledge First Epistemology. In *The routledge companion to epistemology*. Abingdon: Routledge.
- Yablo, S. (1993). Is Conceivability a Guide to Possibility? *Philosophy and Phenomenological Research*, 53(1), 1–42.
- Zagzebski, L. (1994). The inescapability of Gettier problems. *Philosophical Quarterly*, 44(174), 65–73.

- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. New York, NY: Cambridge University Press.
- Zagzebski, L. (1999). What Is Knowledge? In J. Greco & E. Sosa (Eds.), *The blackwell guide to epistemology*. Oxford: Blackwell Publishing.
- Zimmerman, M. J. (1987). Luck and Moral Responsibility. *Ethics*, 97(2), 374–386.

SUMMARY

One of the oldest questions in philosophy, and the question that is central to this thesis, is the question: what is knowledge? While the question is short and simple, the answers that philosophers have formulated since Socrates put forward the question in the *Theaetetus* have tended to be long, complex and inconclusive.

That does not mean that no progress has been made. Over the course of time, something like a consensus has been reached among epistemologists that knowledge requires at least three criteria to be met: if a subject is to have knowledge that p , it must be the case that i) she believes p , ii) her believe that p is true, and iii) she must be *justified* in believing p .

It is the last of these three criteria that is the focal point of this dissertation. Traditionally, the justification condition has been interpreted as a requirement to have evidence for one's belief. For it seems that no true belief can amount to knowledge without the subject having some evidence in favour of it. If I make a wild guess, without any evidence in favour of this belief, that Justin Bieber will be forgotten twenty years from now, this belief will not amount to knowledge even if I turn out to be right.

The interpretation of epistemic justification (as opposed to *moral* justification, for example) in terms of the evidence one has for one's belief is an interpretation that is central to so-called *internalist* theories of justification. These are theories that posit that epistemic justification depends solely on facts that are in some way internal to the believing subject. After an introduction in Chapter 1, I discuss the two most prominent internalist concepts of epistemic justification in Chapter 2: accessibilism and mentalism. I

argue that these two different forms of internalism are motivated by two different conceptions of epistemic justification. For the accessibilist, epistemic justification is a matter of fulfilling your epistemic *duties*, while the mentalist considers justification directly in terms of a certain conception of the evidence one possesses. In Chapter 2, I also discuss the positions of opponents of internalism, the *externalists*. Although externalism is often simply construed as the denial of internalism, I show that different versions of externalism are motivated by different positive views on the nature of justification. Thus, reliabilism is motivated by a truth-conducive view of justification, while virtue-epistemological approaches, as the name indicates, emphasise the importance of epistemic virtues.

With this analysis of the debate between internalists and externalists in hand, we take a step back from epistemology in Chapter 3. Already in the *Theaetetus* we find the suggestion that one of the functions of epistemic justification is the exclusion of luck. Evidence for your belief, after all, makes it less likely that it was only a matter of luck that you formed a true belief. This relation between luck and justification is investigated in the following chapters. In Chapter 3, I start with providing an analysis of the notion of luck itself. I argue that we are best off with a modal notion of luck—an analysis in terms of *possibilities*—which essentially states that an event is lucky only if it is easily possible that the event could not have occurred. I draw on the analysis of Duncan Pritchard, but make some crucial modifications. Among other things, I argue that it is necessary not only to take into account the minimal difference between the current world and the possible worlds where the event does not occur, but that the *proportion* of possible worlds where the event does not take place is important as well.

In Chapter 4, I bring the discussion of luck back to epistemology, and distinguish different forms of luck that may

be relevant for the possession of knowledge. More concretely, I distinguish two forms of luck, *veritic* luck and *reflective* luck, which undermine knowledge. A belief falls a prey to veritic luck in case it is only a matter of luck that the way the belief was formed produced a true belief. Reflective luck can then be construed as the subjective variant of veritic luck. A belief is reflectively lucky if, from the perspective of the subject itself, it is a matter of luck that her belief-forming method produced a true belief. Where most epistemologists agree that veritic luck is incompatible with knowledge, it is still an open question whether reflective luck is. I argue on the basis of Ernest Sosa's recent work that it is useful to distinguish different forms of knowledge, where the weaker forms of knowledge, but not the stronger forms, are compatible with reflective luck.

In Chapters 5 and 6, I return to the debate between internalists and externalists, and connect the findings of the previous chapters to the various positions within this debate. Specifically, in Chapter 5, I argue that internalist notions of justification exclude reflective luck, but do not eliminate veritic luck. In Chapter 6, I argue that the opposite is true for externalist views of justification. They usually eliminate veritic luck but not reflective luck. This allows us to give a new characterization of the debate between internalists and externalists: internalists and externalists disagree on the question *what kind of luck* justification should eliminate.

In Chapter 7, I review the main implications of my findings. First, I argue that we can do justice to the main intuitions of both internalists and externalists by distinguishing two types of justification: a weak type (Type I justification) and a strong type (Type II justification). Where externalist criteria for justification are sufficient for Type I justification, Type II justification will also require the fulfilment of internalist criteria. Second, I argue that the possibility of justified false beliefs gives us a reason to investigate the relationship between epistemic luck and epistemic *risk*. I

propose a possible relationship between these two notions. Third, I deal with the relationship between epistemic justification and the difference between good and bad luck. I argue that true beliefs are justified only in the absence of good luck, while false beliefs are justified only in the *presence* of bad luck. Finally, I consider Duncan Pritchard's claim that it is impossible to exclude reflective luck, and the skepticism that seems to follow from it. I argue that there are several ways to resist Pritchard's skeptical conclusions.

In Chapter 8, I draw some conclusions.

SAMENVATTING

Een van de oudste vragen van de filosofie, en de vraag die in dit proefschrift centraal staat, is de vraag: wat is kennis? Zo simpel en kort als de vraag zelf is, zo ingewikkeld en uitgebreid zijn de filosofische antwoorden hierop gebleken sinds Socrates hem voor het eerst stelde in de *Theaetetus*.

In de loop der tijd is er onder filosofen consensus ontstaan over het feit dat, voor we van kennis kunnen spreken, er voldaan moet worden aan ten minste drie voorwaarden: i) kennis is een overtuiging van een subject, zonder dergelijke overtuiging geen kennis, ii) alleen overtuigingen die waar zijn kunnen met recht kennis genoemd worden, en iii) alleen ware overtuigingen die *gerechtvaardigd* zijn kunnen kennis vormen.

In dit proefschrift is vooral het laatste van deze criteria belangrijk. Traditioneel werd dit rechtvaardigingscriterium vooral gezien als een criterium van goed bewijs. Immers, het lijkt problematisch om te spreken van kennis, zelfs al heb je een ware overtuiging, als je geen enkel *bewijs* voor je overtuiging hebt. Je kan niet weten dat de aarde rond is als dit slechts een wilde gok van je is, zonder dat je enig bewijs voor deze (ware) overtuiging hebt.

De interpretatie van epistemische rechtvaardiging ('epistemisch' om het onderscheid met andere vormen van rechtvaardiging, zoals *morele* rechtvaardiging duidelijk te maken) in termen van het bewijs dat je voor je overtuiging hebt is een interpretatie die vooral tot uitdrukking komt in zogenaamde *internalistische* theorieën van rechtvaardiging. Dit zijn theorieën die stellen dat epistemische rechtvaardiging slechts afhangt van zaken die op enigerlei wijze intern zijn aan het kennend subject. Nadat ik in Hoofdstuk 1 een inleiding gegeven heb, bespreek ik in Hoofdstuk 2

de twee meest voorkomende vormen van internalisme: het accessibilisme en het mentalisme. Ik betoog dat deze twee verschillende vormen van internalisme gemotiveerd worden door twee verschillende concepten van epistemische rechtvaardiging. Voor de accessibilist is epistemische rechtvaardiging een kwestie van het nakomen van je epistemische *plichten*, terwijl de mentalist rechtvaardiging direct opvat in termen van een bepaalde opvatting van bewijs. Ook bespreek ik in Hoofdstuk 2 de posities van de tegenstanders van het internalisme, de *externalisten*. Alhoewel externalisme vaak simpelweg wordt opgevat als de ontkenning van de internalistische these dat rechtvaardiging alleen afhangt van wat intern is aan het subject, laat ik zien dat verschillende versies van het externalisme wel degelijk een positieve opvatting hebben over wat rechtvaardiging is, en bovendien van elkaar onderscheiden kunnen worden aan de hand van deze opvattingen. Zo wordt het reliabilisme gemotiveerd door een waarheidsbevorderende opvatting van rechtvaardiging, terwijl deugdepistemologische benaderingen, zoals de naam al aangeeft, het belang van epistemische deugden benadrukken.

Met deze analyse van het internalisme-externalisme debat in de hand nemen we in Hoofdstuk 3 even afstand van de epistemologie. Al in de *Theaetetus* vinden we namelijk de suggestie dat een van de functies van epistemische rechtvaardiging het uitsluiten van *toeval* is. Bewijs voor je overtuiging, immers, maakt het minder waarschijnlijk dat je slechts toevallig een ware overtuiging gevormd hebt, zoals bij een wilde gok wel het geval is. In Hoofdstuk 3 analyseer ik dan ook de notie van toeval. Ik beargumenteer dat we het beste af zijn met een modale notie van toeval—een analyse in termen van alternatieve mogelijkheden—die in essentie stelt dat een gebeurtenis slechts een kwestie van toeval is als de gebeurtenis gemakkelijk niet had kunnen plaatsvinden. Ik baseer me hier op de analyse van Duncan Pritchard, maar pas deze op een aantal cruciale punten aan.

Zo betoog ik onder andere dat het noodzakelijk is niet alleen rekening te houden met het minimale verschil tussen de huidige wereld en de mogelijke werelden waar de gebeurtenis niet voorkomt, maar dat ook de proportie van mogelijke werelden waar de gebeurtenis niet plaatsvindt belangrijk is.

In Hoofdstuk 4 pas ik mijn analyse van toeval toe op de epistemologie, en onderscheid ik verschillende vormen van toeval die relevant zijn in het debat over de vraag wat kennis is. Meer precies onderscheid ik twee vormen van toeval, *veristisch* toeval en *reflectief* toeval, die kennis ondermijnen. Een overtuiging valt ten prooi aan veristisch toeval in het geval het slechts toevallig is dat de methode waarmee deze opvatting tot stand gebracht is een ware overtuiging opleverde. Reflectief toeval kan worden opgevat als de subjectieve variant van veristisch toeval. Een overtuiging is reflectief toevallig als het *vanuit het perspectief van het subject zelf* slecht toevallig is dat de methode waarmee zij haar overtuiging vormde een ware overtuiging voortbracht. Daar waar de meeste epistemologen het erover eens zijn dat veristisch toeval kennis ondermijnt, is het vooralsnog een open vraag of reflectief toeval compatibel is met het hebben van kennis. Ik beargumenteer aan de hand van recent werk van Ernest Sosa dat het nuttig is in dit verband verschillende vormen van kennis te onderscheiden, waarbij de zwakkere vormen van kennis wél, maar de sterkere vormen van kennis níet compatibel zijn met reflectief toeval.

In Hoofdstuk 5 en 6 kom ik terug op het onderscheid tussen internalisme en externalisme, en verbind de bevindingen uit de vorige hoofdstukken aan de verschillende posities binnen dit debat. Specifiek betoog ik in Hoofdstuk 5 dat internalistische opvattingen van rechtvaardiging reflectief toeval wel uitsluiten, maar veristisch toeval niet. In hoofdstuk 6 betoog is het omgekeerde voor externalistische opvattingen van rechtvaardiging. Zij sluiten in de re-

gel wel veristisch toeval uit, maar niet reflectief toeval. Dit stelt ons in staat een nieuwe karakterisering van het debat tussen internalisten en externalisten te geven: internalisten en externalisten verschillen met elkaar van mening over de vraag *welke vorm* van toeval de rechtvaardiging van onze overtuigingen uit dient te sluiten.

In hoofdstuk 7 ga ik na wat de belangrijkste implicaties van mijn bevindingen zijn. Ten eerste betoog ik dat we recht kunnen doen aan de belangrijkste intuïties van zowel internalisten als externalisten door twee typen rechtvaardiging te onderscheiden: een zwakke vorm (Type I rechtvaardiging) en een sterke vorm (Type II rechtvaardiging). Daar waar externalistische criteria voor rechtvaardiging voldoende zijn voor Type I rechtvaardiging, zal voor Type II rechtvaardiging ook aan de internalistische criteria voldaan moeten worden. Ten tweede betoog ik dat de mogelijkheid van onware gerechtvaardigde overtuigingen ons een reden verschaft om te onderzoeken wat de relatie is tussen epistemisch toeval en epistemisch *risico*. Ik doe een voorstel voor een mogelijke relatie. Ten derde behandel ik de relatie tussen epistemische rechtvaardiging en het verschil tussen positief en negatief toeval ('good luck' en 'bad luck', respectievelijk). Ik beargumenteer dat ware overtuigingen slechts gerechtvaardigd zijn als er geen sprake is van positief toeval (good luck), terwijl onware overtuigingen slechts gerechtvaardigd zijn als er *wel* sprake is van negatief toeval (bad luck). Ten slotte behandel ik de claim van Duncan Pritchard dat het onmogelijk is reflectief toeval uit te sluiten, en het scepticisme dat daaruit voortvloeit. Ik betoog dat er meerdere manieren zijn om aan Pritchards sceptische conclusies te ontkomen.

In hoofdstuk 8 trek ik mijn conclusies.