

EPISTEMIC LUCK, EPISTEMIC AGENCY, AND SKEPTICISM

by

Michaela Mueller

Copyright © Michaela Mueller 2007

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2007

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Michaela Mueller entitled Epistemic Luck, Epistemic Agency, and Skepticism and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

Terry Horgan Date: 4/30/07

Keith Lehrer Date: 4/30/07

John Pollock Date: 4/30/07

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Dissertation Director: Terry Horgan Date: 4/30/07

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at the University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the copyright holder.

SIGNED: Michaela Mueller

ACKNOWLEDGMENTS

Every text that is written reflects something about the person who wrote it and about the place in which this person has lived. For more than six years Tucson, Arizona has been my home, both philosophically and personally. In these years I have been able to grow into a more mature and better philosopher and person, and I am very grateful for the various opportunities that allowed this. It seems to me that this dissertation reflects what is good in the philosophical and personal environment that I have found here.

There are more specific philosophical debts that need to be acknowledged. The dissertation was directed by Terry Horgan, who has influenced it through suggestions at every stage and without whom the dissertation would not have taken this form. The other members of my dissertation committee, Keith Lehrer and John Pollock have also shaped the dissertation through both encouragement and criticism. The members of the dissertation reading group have provided feedback on each of the main chapters of the dissertation.

I wish to thank very much all those who through their support enabled me to write this dissertation.

TABLE OF CONTENTS

LIST OF SUMMARY BOXES.....	8
ABSTRACT.....	9
 1. INTRODUCTION.....	 10
1.1 The gap.....	10
1.2 The skeptic.....	12
1.3 Epistemic luck.....	13
1.4 Epistemic agency.....	16
1.5 An account of perceptual knowledge.....	17
 2. EPISTEMIC LUCK AND SKEPTICISM.....	 22
2.1 The problem and the proposal: Why luck-precluding conditions are necessary conditions for knowledge.....	22
2.2 The connection between skepticism and epistemic luck.....	25
2.3 Accepting the gap.....	33
2.4 Different notions of epistemic luck as a strategy for identifying a flaw in the skeptical argument.....	37
2.5 Non-skeptical epistemic luck.....	43
2.5.1 Searching for the conditions given which it is not lucky that the belief is true.....	43
2.5.2 Luck-precluding conditions.....	47
2.5.3 Comparison with other views.....	53
2.6 Revisiting the skeptical argument.....	62
2.7 The problem and the proposal revisited.....	69
 3. EPISTEMIC LUCK AND GETTIER CASES.....	 72
3.1 The problem and the proposal: How luck-precluding conditions allow us to avoid Gettier cases	72
3.2 The presence of non-skeptical epistemic luck.....	75
3.2.1 Types and tokens of evidence.....	76
3.2.2 Types and tokens of luck-precluding conditions.....	81
3.2.3 Luck-precluding conditions as explanations.....	87
3.3 The type of luck present in Gettier cases is non-skeptical epistemic luck...	89
3.3.1 The proposal.....	89
3.3.2 Several Gettier cases.....	95
3.3.3 An argument.....	101
3.4 Given the absence of non-skeptical epistemic luck no Gettier case can be constructed.....	106
3.4.1 Expected luck-precluding conditions.....	107
3.4.2 Gettier cases revisited.....	114
3.4.3 Can there be a Gettier counterexample to the proposed account of perceptual knowledge?.....	119

TABEL OF CONTENTS - *Continued*

3.5 Comparison with other views.....	132
3.6 The problem and the proposal revisited.....	143
4. EPISTEMIC AGENCY AND JUSTIFICATION.....	145
4.1 The problem and the proposal: What kind of justification is necessary for perceptual knowledge?.....	145
4.2 Epistemic agency.....	147
4.2.1 The capacity to receive sensory experiences.....	148
4.2.2 The capacity to read sensory experiences.....	154
4.2.3 The capacity to evaluate readings of sensory experiences.....	156
4.2.4 Comparison with other views.....	166
4.3 Evidence-based justification.....	167
4.3.1 The proposal.....	167
4.3.2 Evidence.....	171
4.3.3 Comparison with other views.....	179
4.3.4 Distinguishing evidence-based justification from evidence.....	183
4.3.5 Different brain in vat situations.....	190
4.3.6 Evidence-based justification.....	200
4.4 How evidence-based justification is ideally tailored to the luck-precluding condition.....	213
4.5 The problem and the proposal revisited.....	223
5. EPISTEMIC AGENCY AND SKEPTICISM.....	226
5.1 The problem and the proposal: A set of necessary and sufficient conditions for perceptual knowledge and a reply to some skeptical worries.....	226
5.2 Necessary conditions for perceptual knowledge.....	228
5.2.1 True belief.....	229
5.2.2 Evidence.....	230
5.2.3 Luck-precluding conditions.....	244
5.2.4 Evidence-based justification.....	252
5.2.5 True belief revisited.....	260
5.3 Sufficient conditions for perceptual knowledge.....	261
5.3.1 An argument.....	261
5.3.2 Problematic forms of epistemic luck.....	264
5.3.3 Counterexamples based on epistemic luck.....	266
5.3.4 Epistemic luck and the necessary conditions.....	268
5.3.5 Comparison with other views.....	273
5.4 A reply to some skeptical worries.....	281
5.4.1 Two senses of possibility of perceptual knowledge.....	281
5.4.2 Knowledge of the nature of our epistemic agency.....	284
5.4.3 We know that we are capable of perceptual knowledge.....	287
5.4.4 Our perceptual knowledge.....	289

TABLE OF CONTENTS - *Continued*

5.5 The problem and the proposal revisited.....	291
6. CONCLUSION.....	294
6.1 Other kinds of knowledge, both of the physical world and non-physical things.....	294
6.2 Other unfinished work.....	305
REFERENCES.....	311

LIST OF SUMMARY BOXES

2. EPISTEMIC LUCK AND SKEPTICISM

SUMMARY BOX 1,	A skeptical argument based on epistemic luck.....	33
SUMMARY BOX 2,	The distinction between skeptical and non-skeptical epistemic luck.....	43
SUMMARY BOX 3,	Luck-precluding conditions.....	53
SUMMARY BOX 4,	Criticism of the skeptical argument.....	68

3. EPISTEMIC LUCK AND GETTIER CASES

SUMMARY BOX 5,	Type-token transparent evidence and type-token opaque luck-precluding conditions.....	87
SUMMARY BOX 6,	Luck-precluding conditions are necessary conditions for perceptual knowledge.....	94
SUMMARY BOX 7,	Expected luck-precluding conditions.....	113
SUMMARY BOX 8,	A Gettier counterexample?.....	131

4. EPISTEMIC AGENCY AND JUSTIFICATION

SUMMARY BOX 9,	The nature of our epistemic agency.....	165
SUMMARY BOX 10,	Distinguishing evidence from evidence-based justification.....	189
SUMMARY BOX 11,	The definition of evidence-based justification.....	212
SUMMARY BOX 12,	Evidence-based justification and luck-precluding conditions.....	222

5. EPISTEMIC AGENCY AND SKEPTICISM

SUMMARY BOX 13,	Evidence is necessary for perceptual knowledge.....	244
SUMMARY BOX 14,	The luck-precluding condition is necessary for perceptual knowledge.....	251
SUMMARY BOX 15,	Evidence-based justification is necessary for perceptual knowledge.....	259
SUMMARY BOX 16,	An argument that the necessary conditions are jointly sufficient for perceptual knowledge.....	263
SUMMARY BOX 17,	An argument that we are capable of perceptual knowledge.....	290

ABSTRACT

In my dissertation I develop an account of perceptual knowledge through thinking about epistemic luck, epistemic agency, and skepticism. Two conditions are, as I claim, necessary and sufficient to render a true belief an instance of perceptual knowledge. These conditions are the luck-precluding condition and evidence-based justification. The luck-precluding conditions are external conditions, consisting in physical regularities in the world that allow us to arrive through our evidence at the truth in a systematic way. I claim that these luck-precluding conditions also allow us to avoid Gettier cases. Evidence-based justification requires evidence as an internal condition. However, there are some external constraints as to when our evidence can reasonably be expected to lead to the kind of justification required for perceptual knowledge. This account of perceptual knowledge blends together internal and external requirements. The claim is that only given this match of internal and external requirements can we argue that the resulting set of necessary conditions is also sufficient for perceptual knowledge in that it allows us to avoid all problematic forms of epistemic luck.

1. INTRODUCTION

The project of my dissertation is to develop an account of perceptual knowledge. This account is developed in four main chapters, each of which responds to a problem that any such account faces. I claim that a solution to these problems can be found by thinking about epistemic luck and epistemic agency. In this introduction I will sketch the main ideas of the dissertation.

1.1 The gap

The basic assumption and starting point for the theory in this dissertation is that our evidence is inevitably limited in that it cannot by itself guarantee that the belief for which we have that evidence is true. This means that it is always possible that we have the evidence and yet the belief for which we have the evidence is false. Consider for example that you are driving by what you take to be a barn. Your belief is that there is a barn in the field. Your evidence for this belief does not guarantee that the belief is true. This is because your evidence consists in a sensory experience as of a part of a barn. It might be that what you see is not a barn but a mere barn façade.

We are in the following situation: on the basis of mere glimpses, of little bits of evidence we draw conclusions that go far beyond that evidence. We have a sensory experience as of a part of a barn and conclude that there is a barn in the field. However, this evidence does not guarantee that there indeed is a barn in the field. It might be a mere façade. I

take this situation to be inevitable. We always have evidence that we read as indicating a truth when this evidence is compatible with there being no such truth.

One might ask: Why not get out of the car and walk around the barn, open its door, knock on its walls, and so on? We would then have sensory evidence as of a structure as opposed to a façade. However, we would still have merely more evidence, which is itself also compatible with there being no barn. Maybe someone created a strange barn illusion that looks from all sides exactly like a barn and that even responds to touch, but that is no more than a mirror image that comes with unusual properties including responsiveness to touch.

So it seems that even though one can pile up more and more evidence for the belief that there is a barn in the field, any such set of evidence is compatible with there being no barn in the field. Furthermore, this possibility, that what the evidence indicates is false, does not depend on a problematic distinction between mind and world. There are familiar illusions, which, if they occur, would mean that our evidence is present but the belief for which we have the evidence is false. There is nothing mysterious or problematic about the idea that we might be looking at a mere barn façade, or that we might be looking at an illusion of some kind. For example, if we were to read a sign saying that in this region mere barn facades have been put up, we will no longer conclude on the basis of a casual glimpse that there is a barn in the field. If we learn of an illusion that is three-dimensional and responsive to touch, some kind of a scientific experiment conducted in this area, we

will hesitate to say that something that looks from all sides and feels like a barn is a barn. So it is not a mere philosophical distinction between mind and world that creates this problem. The problem is at the very heart of all of our epistemic endeavors. It is not created by philosophy; it exists because of our actual epistemic situation, that our evidence is inevitably limited in that it cannot guarantee that the belief for which we have the evidence is true.

The basic assumption made is thus that there this gap between our evidence and the truth of the belief for which we have the evidence. The question that now arises is whether this gap is a problem. As we will see in a minute, the skeptic claims it is a problem. The dissertation tries to show that this gap does not represent a problem, that this gap is compatible with our having perceptual knowledge.

1.2 The skeptic

We assume that there is such a gap. Somehow the skeptic wants to say that if there is such a gap then we do not have any knowledge. The skeptic wants to say that if our evidence for a belief is compatible with that belief being false, then we do not know that it is true. The question we need to ask is why that should be the case. Why would the gap render it impossible that we have any knowledge? As we have seen, the existence of this gap is something intuitive and present in our everyday life. However, we do not normally intuitively and in everyday life conclude that we do not have any knowledge because our

evidence is inevitably limited. So it seems that the skeptic needs to tell us why this gap is incompatible with knowledge.

As it seems to me, what the skeptic should say is that if given a set of evidence our belief might be true but might also be false, then it is just lucky if it is indeed true, and this luck is incompatible with knowledge. So the idea is that if given our evidence it is possible for our belief to be false, then if it is true, one is just lucky and does not really have knowledge. Epistemic luck might thus provide the bridge between the existence of this gap and the skeptic's claim that there is no knowledge.

We can thus construct a skeptical argument in the following way. We begin with a statement of the gap: Our evidence for our beliefs is compatible with these beliefs being false. Next we have the idea that if there is such a gap then it is a matter of epistemic luck if the belief is true: If given a set of evidence a belief might be true and might be false, then it is a matter of epistemic luck if it is true. The skeptic now just needs to say that this epistemic luck is incompatible with knowledge, and we get the conclusion that we do not have any knowledge. The second chapter of my dissertation examines this argument and develops a non-skeptical response to it.

1.3 Epistemic luck

The main question that the assumption of the gap and the skeptical argument based on epistemic luck raise is, naturally, what exactly epistemic luck is and whether all of its

varieties are incompatible with knowledge. I call those forms of epistemic luck that are incompatible with knowledge “problematic forms of epistemic luck,” as opposed to the unproblematic ones, those that are compatible with knowledge. The main question we need to ask is what characterizes problematic forms of epistemic luck.

This question is not somehow a side issue in epistemology. Instead the distinction between problematic and unproblematic forms of epistemic luck is crucial for any attempt to develop a set of necessary and sufficient conditions for knowledge. This is particularly obvious when we consider Gettier cases. In a Gettier case we have a situation in which someone has a justified true belief, but no knowledge because it is somehow a matter of luck that the justified belief is also true.

Gettier cases are just one type of counterexample for a proposed set of necessary and sufficient conditions for knowledge. In the third chapter of my dissertation I develop an understanding of the kind of epistemic luck involved in a Gettier case and a proposal of how to avoid this kind of problematic epistemic luck. However, Gettier cases share with other possible counterexamples to a set of necessary and sufficient conditions for knowledge that they rely on epistemic luck. In the fifth chapter of my dissertation I give arguments that each of the proposed conditions for knowledge is necessary. These arguments all claim that if this necessary condition is absent then a counterexample to the account of knowledge can be given that is based on a problematic form of epistemic luck.

Thus epistemic luck is crucial for arguing that proposed conditions for knowledge are necessary.

Also in the fifth chapter I develop an argument that the necessary conditions are together sufficient for knowledge. Again, this argument draws on epistemic luck. It claims that to the proposed set of conditions for knowledge there can be no counterexample that is based on a problematic form of epistemic luck. This argument is possible only given an understanding of what renders a form of epistemic luck problematic. As I claim, problematic forms of epistemic luck arise because the internal and the external requirements for knowledge are not sufficiently tailored to each other, they do not fit together as seamlessly as possible. In the fifth chapter of my dissertation I claim that the account developed here fits together the internal and external requirements for knowledge as seamlessly as at all possible.

Thus we can see that epistemic luck is crucial for developing an account of perceptual knowledge. However, the dissertation does not rely on mere intuitions of when something is too lucky to be knowledge. Instead these intuitions are complemented by a theoretical insight as to what renders a form of epistemic luck problematic. The intuitions merely point us in a certain direction, it is a theory and whatever its merits are that allow me to claim that the proposed set of necessary and sufficient conditions is not subject to a problematic form of epistemic luck.

1.4 Epistemic agency

In order to develop a set of necessary and sufficient conditions for perceptual knowledge, one needs to ask what kind of justification is required for such knowledge. Thinking about epistemic luck leads to the proposal of an external requirement for knowledge, what I call “luck-precluding conditions.” Now, if it is correct that problematic forms of epistemic luck arise due to a mismatch of the internal and external requirements for knowledge, then we would expect that this external requirement, the luck-precluding conditions must somehow be matched with an internal requirement. This internal requirement is provided by what I call “evidence-based justification,” the kind of justification required for perceptual knowledge.

Evidence is understood as an internal condition, consisting in sensory experiences and beliefs, where these are somehow things that go on in the mind. For perceptual knowledge we need such evidence, in particular sensory experiences. To have perceptual knowledge that there is a barn in the field we need a sensory experience as of a barn. I develop an account of this evidence-based justification on the basis of our epistemic agency.

By “epistemic agency” I mean those capacities that we have that are relevant to whether or not we are able to gain a justified belief and knowledge. In the case of perceptual knowledge there are three such capacities: the capacity to receive sensory experiences as input, the capacity to read this input, and the capacity to evaluate readings of this input.

So we might have a sensory experience as of a barn, we read it as indicating that there is a barn in the field, and then ask whether we should believe what we have read off of the sensory experience. Sometimes we will not believe what we have read off of a sensory experience, as when we see a stick partly immersed in water as bent, but, remembering the familiar illusion, do not believe that the stick is bent.

Due to the nature of our epistemic agency we are able to receive sensory experiences as evidence, to read them and to evaluate our readings of them. This is what renders us capable of acquiring a justified belief and knowledge. It is evidence that confers evidence-based justification on us, provided that certain external constraints are met. These external constraints cannot give us justification; they merely limit the cases in which evidence can give us justification. Basically what these external constraints demand is that we are at all capable of gaining perceptual knowledge.

The name “epistemic agency” for these capacities is appropriate because there is a certain decision required on our part, namely whether to accept the evidence presented to us or not. This is a genuine decision in that it is always possible that further evidence would defeat our current evidence. At some point we have to decide as epistemic agents whether or not to accept what evidence we have or whether to persist in searching for further confirmation. This decision will often not be conscious.

1.5 An account of perceptual knowledge

Suppose that, driving around, you have a glimpse of what you take to be a barn and arrive at the belief that there is a barn in the field. Suppose that the belief is true. There are two further conditions that are required for perceptual knowledge that there is a barn in the field. These are the luck-precluding conditions and evidence-based justification.

You truly believe that there is a barn in the field on the evidence of a sensory experience of a part of a barn, a barn façade. The luck-precluding conditions form a tight connection between this evidence and the truth that it indicates. The idea is that the world has to be a certain way so that on the basis of this evidence we will arrive at a true belief. This tight connection between evidence and truth is supposed to be developed through a physical regularity in the world. In the example, the barn façade must be a part of a barn rather than a mere façade. In addition, the barn façade must somehow cause in you the sensory experience as of a barn. So there is a causal regularity and a part-whole regularity that each must be present, both in general and in this instance. If the world in which we live is indeed characterized by these regularities, then the world is such that on the basis of a sensory experience as of a barn we will arrive at the true belief that there is a barn in the field.

The luck-precluding conditions are strictly external conditions. However, the mere obtaining of these external conditions is not sufficient for knowledge. For there might be very strange external regularities which are somehow beyond our cognitive horizon. Consider for example a somewhat weird Gettier case, where you are looking out on a

field and truly believe that there is a sheep in the field. However, your sensory experience as of a sheep is caused by a sheepdog. Furthermore, there is a regularity where sheepdogs are followed at a safe distance by sheep, too far away to be recognizable as such. Now the situation is that you have a sensory experience as of a sheep and that there indeed is a sheep in the field. Furthermore, there is a regularity linking what causes your sensory experience, the sheep dog, to what renders the belief true, the unrecognizable sheep in the distance. However, I take it that we would not want to say that we have knowledge in this situation.

As we can see, all the external conditions for knowledge are met in this case. There is a sheep, it is linked by a regularity to our evidence. Weirdly enough, all the conditions for knowledge that the world needs to provide are provided. However, there is still no knowledge because the regularity is beyond our cognitive horizon. If we share this intuition about this case, that there is no knowledge even though all the external conditions for knowledge are present, we have further support for the theoretical insight that internal and external conditions need to be matched, that they need to fit together as seamlessly as possible for us to have knowledge.

Now what is missing in the Gettier case just described is that somehow there needs to be an expectation or assumption that the luck-precluding condition is present. It is somewhat difficult to say exactly what kind of a thing this expectation or assumption is. Surely, in ordinary perception we do not have any beliefs about how our cognitive apparatus works.

In ordinary perception we are not even aware of the distinction between sensory experience and physical reality. This is because we see the world through the sensory experience, and this is not somehow like looking at a picture where we can also look at the original.

As the Gettier case just described shows, we need somehow an expectation or assumption of this regularity in the physical world that allows us to arrive through our evidence at a true belief. However, this assumption or expectation is not an occurring belief, it is not somehow a premise through which we reason to our belief about the physical world. It is present only in so far as we would respond to certain defeaters. In the example, if the epistemic agent learns that sheepdogs are not followed by unrecognizable sheep in the distance, he would not withdraw the belief that there is a sheep in the field. So this regularity which is required for the evidence to lead to a true belief is not within the cognitive horizon of the epistemic agent in the sense that in acquiring a defeater that this regularity does not obtain he would not withdraw the belief.

In addition to the external requirements, in particular of the luck-precluding conditions, we need to require internally that these luck-precluding conditions are expected or assumed to obtain. Suppose that someone has evidence for a true belief and the luck-precluding condition obtains. The epistemic agent expects or assumes that the luck-precluding condition obtains if evidence that it does not obtain would be regarded as a defeater for the belief on the basis of the evidence. If learning that the regularity on which

the luck-precluding condition is based does not obtain is not regarded as a defeater then this luck-precluding condition is beyond the cognitive horizon of the epistemic agent.

The account of perceptual knowledge combines internal requirements, belief and evidence, with external requirements, truth and luck-precluding conditions. The luck-precluding conditions allow us on the basis of the evidence to arrive at a true belief. Evidence-based justification involves an expectation or assumption that these luck-precluding conditions obtain. The internal conditions thus expect the external conditions, and the external conditions underwrite and support the internal conditions. In this the internal and external conditions for knowledge are as well matched and fitted together as at all possible.

Let me briefly say what can be found in the following five chapters of my dissertation. In the second chapter I discuss the gap and the skeptical argument based on the gap. A response to this skeptical argument, based on a distinction between two kinds of epistemic luck, is developed. In the third chapter we see that this distinction between two kinds of epistemic luck also allows us to diagnose why in a Gettier case we have no knowledge and to propose a response as to how to avoid Gettier cases. In the fourth chapter the account of evidence-based justification is developed. In the fifth chapter, arguments are given that each of the conditions for perceptual knowledge is individually necessary and that they are jointly sufficient. In the concluding sixth chapter other kinds of knowledge are compared with perceptual knowledge.

2. EPISTEMIC LUCK AND SKEPTICISM

2.1 The problem and the proposal: Why luck-precluding conditions are necessary conditions for knowledge

Suppose that in 1985 a novel is published that is set in California of 2005. Many interesting things happen in this novel: people love in vain, they are involved in elaborate misunderstandings, their kids are rebellious, and – Arnold Schwarzenegger is elected Governor of California. Now suppose that on the other end of the globe, there is a secluded spot where international news never finds its way and where novels published 20 years ago are considered “new arrivals.” In this place there lives a person, named “Lucky,” who reads this novel in 2005, but for some reason mistakes it for an account of contemporary life in California, not a work of fiction but a report of what is going on in California right now. As a result, Lucky comes to believe that Arnold Schwarzenegger is Governor of California. The belief is true. But obviously Lucky is just lucky; that her belief is true is a matter of epistemic luck.

As this example shows, some forms of epistemic luck are obviously incompatible with knowledge. Lucky does not know that Arnold Schwarzenegger is Governor of California, because it is merely lucky that her belief is also true. In some cases we have this intuition that if it is merely lucky that a belief is true, this belief is not knowledge. At a minimum we want to say that knowledge requires a belief that is true but not as a matter of a form

of epistemic luck that is incompatible with knowledge. This chapter tries to understand better the ties between epistemic luck and knowledge.

As should be clear, this requires us to think about the varieties of epistemic luck, only some of which are clearly incompatible with knowledge. It might be a matter of luck that one witnesses a car accident in that one might easily not have been in that place at that time. However, this luck does not seem to prevent us from knowing a lot of things about the car accident. I will later, in the fifth chapter of my dissertation, try to say more precisely what kinds of epistemic luck are compatible and what kinds are incompatible with knowledge. In this chapter we will merely look at two varieties of epistemic luck, namely “skeptical epistemic luck” and “non-skeptical epistemic luck.”

This distinction is made in response to a skeptical argument. The skeptic might say that there is always a certain element of epistemic luck involved when we arrive at a true belief. According to the skeptic, epistemic luck is unavoidable. However, an example like the Lucky case given above might be taken to show that epistemic luck is incompatible with knowledge. These two premises would lead to the skeptical conclusion that we do not have any knowledge.

As already suggested, this argument is flawed in that it does not recognize the varieties of epistemic luck, not all of which are incompatible with knowledge. The proposal in response to this skeptical argument is, thus, that we need to distinguish between skeptical

and non-skeptical epistemic luck. Skeptical epistemic luck is unavoidable, but compatible with knowledge. Non-skeptical epistemic luck is incompatible with knowledge, but not unavoidable. This distinction allows us to criticize the skeptical argument.

The main result of this discussion is the proposal of a necessary condition for perceptual knowledge, called “luck-precluding condition.” These luck-precluding conditions allow us to avoid non-skeptical epistemic luck. In later chapters we will have to think more about other forms of epistemic luck that might also be incompatible with knowledge. If this were not too cumbersome, one would have to call this necessary condition more appropriately “non-skeptical-luck-precluding condition.”

Let me sketch what I want to do in this chapter. In the next section I will explain how skepticism and epistemic luck are related to one another. We will see that the skeptic can give an argument for his skepticism that is based on epistemic luck. This argument is also based on what I call the “gap.” In section 2.3 I will explain that I regard this gap as a fact about our actual epistemic situation concerning the physical world. In the remaining sections I will develop the distinction between skeptical and non-skeptical epistemic luck. This distinction allows us to criticize the skeptical argument based on the gap and epistemic luck.

2.2 The connection between skepticism and epistemic luck

In this section I will explain a skeptical argument for the conclusion that we do not have any perceptual knowledge. This skeptical argument begins with the statement that there is a gap between our evidence and the truth of the beliefs for which we have this evidence. This gap consists in the fact that it is possible to have this evidence when the belief for which we have the evidence is false.

To explain this gap, let us consider an ordinary case of perceptual knowledge, my knowledge that I am sitting at a desk. My evidence for this true belief consists in a variety of sensory experiences as of a desk. It is possible for me to have this evidence when the belief that this evidence indicates is false. Descartes has sketched some such possibilities. I might be dreaming instead of sitting at my desk. I might be deceived in some way by an evil demon so as to confuse me into thinking that I am now sitting at my desk. I might be a brain in a vat that is deceived into thinking that there is a desk in front of it when there is no such thing. In all of these cases I would have the evidence, the sensory experiences as of a desk, when in fact the belief that this evidence indicates, that there is a desk, is false.

These scenarios just sketched are somewhat fanciful, and they might suggest that this gap is somehow raised only by philosophical speculation and something removed from our everyday situation. However, there is nothing fanciful about the possibility of a variety of illusions, including mere barn facades, mirror images, tricks of the light, hallucinations

induced by alcohol, projected images, and so on and so forth. All of these illusions are everyday examples, and it does not seem that they are the mere products of philosophical speculation.

However, the gap does not just arise because of the existence of such illusions. After all, these illusions are not phenomena that we encounter every day. We do not run very often into mere barn facades, deceptive mirror images, or alcohol induced hallucinations.

However, this is not what the gap is all about. Rather it is all about the inevitable limitations of our evidence. We always only have glimpses, little bits of evidence, on the basis of which we draw conclusions that go beyond these glimpses and little bits of evidence.

Suppose that we go on a drive and, looking out of the window, we arrive at a variety of perceptual beliefs: There is a house, a barn, a tree, a village, a gas station, a deer, and so on and so forth. We arrive at these beliefs on the basis of casual glimpses, of sensory experiences as of a house, or a barn, or a tree, and so on. The possibilities given which these glimpses would be misleading are somewhat strange. One might be driving through a movie set rather than a village, the tree might not be a tree but a sculpture of a tree, the barn might consist in a mere façade. We normally reject such strange possibilities, but the gap is merely made explicit through such strange possibilities. It is not that the gap exists because of these strange possibilities, but that these strange possibilities are possibilities because of the gap. It is because one in fact sees only a part of a barn that it is a

possibility that one is looking at a mere façade. Similarly, it is because one recognizes the shape and size of the tree that one understands the possibility that it might be a sculpture rather than a tree.

It is a perfectly ordinary situation in which one concludes on the basis of a sensory experience as of a part of a barn that there is a barn in the field. One's evidence for the belief that there is a barn in the field is compatible with there being no such thing, and instead a mere façade or some kind of an illusion. One might seek further confirmation for the belief that there is a barn. One might walk around it, open its door, knock on its walls. But all of this would only increase the evidence, but not close the gap. The evidence would still be compatible with the falsity of the belief for which we have the evidence. This is because our evidence for perceptual beliefs ultimately consists in sensory experiences, where these are probably states of the mind. However, these states of the mind are not identical to the truth about the physical world they indicate. We might have sensory experiences as of a barn even if there is no barn in the field. Evidence is not identical to the truth that this evidence indicates. Because of this it is possible that the evidence exists while what it indicates is false.

This gap is not somehow created by philosophy; it is not something that we through philosophizing bring into existence. Rather this gap characterizes our inevitable epistemic situation. If the skeptic begins his argument with a statement of the gap, it seems to me that we cannot deny that there indeed is this gap.

However, the skeptic wants to say that we do not have any knowledge. Even if we acknowledge that there is this gap, that our evidence is always and inevitably compatible with the falsity of what this evidence indicates, the skeptic still needs to get from this claim to the conclusion that we do not have any knowledge. It is not self-evident that we do not have any knowledge just because there is this gap. The existence of the gap is supported by our everyday intuitions, by our ordinary beliefs. But we do not in everyday life assume that because our evidence is limited we have no knowledge. Instead we confidently assert our knowledge of a wide variety of things, knowing that our evidence for these beliefs is inevitably limited. So the skeptic needs to say why it is that the gap precludes us from having any knowledge.

This is where epistemic luck comes in. The skeptic could say that if given our evidence it is possible that the belief for which we have the evidence is false, then it is merely a matter of epistemic luck if the belief is true. So the idea is that the skeptic might say that the gap means that given our evidence the belief for which we have the evidence might be true but might also be false. Furthermore, if a belief might be true and might be false given our evidence, then if it is true this is merely a matter of epistemic luck. Now, epistemic luck seems to be incompatible with knowledge. If it is merely a matter of epistemic luck that a belief is true, then, it seems, we do not know that it is true. Thus there is an argument for skepticism about knowledge that draws on the gap and epistemic luck.

Let us state this argument more precisely.

- (1) My evidence for a belief about the physical world is inevitably compatible with this belief being false.
- (2) If my evidence for a belief about the physical world is inevitably compatible with this belief being false, then, if the belief is true, this is a matter of epistemic luck.
- (3) If it is a matter of epistemic luck that a belief is true, then I do not know that this belief is true.
- (4) Therefore, even if the belief about the physical world for which I have the evidence is true, I do not know that it is true.

The skeptical argument thus says that even if I have good evidence for a true belief, this true belief is not knowledge because of the gap and epistemic luck. In the remainder of this chapter I will identify a way of showing this argument to be flawed. Let me conclude this section by emphasizing why this skeptical argument and attempts to identify a flaw in it are important.

As should be clear from the preliminary explanation above, I take the gap to be indeed a characteristic of our epistemic situation. It seems to me that it is impossible to ever close this gap, that this gap is an unavoidable element of our actual epistemic situation. Any attempt to develop a theory of perceptual knowledge must somehow deal with this gap,

must somehow find a response to this gap. This is why I begin this dissertation with a discussion of this skeptical argument.

Furthermore, I take it that there is a lot of intuitive evidence for the claim that at least certain varieties of epistemic luck are indeed incompatible with knowledge. Lucky in the example given at the beginning of the chapter does not know what she truly believes. In addition to a true belief it can not, or so it seems, also be the case that this belief is true as a matter of a problematic form of epistemic luck. In this chapter I will identify one such form of problematic epistemic luck, and suggest a necessary condition for knowledge that would eliminate this problematic form of epistemic luck.

Before beginning the discussion of this skeptical argument, let us see that this skeptical argument is a version of the Cartesian-type argument for skepticism as it is discussed in contemporary epistemology. Peter Unger (1974) discusses the following argument for skepticism:

- (1) If someone knows something to be so, then it is all right for the person to be absolutely certain that it is so.
- (2) It is never all right for anyone to be absolutely certain that anything is so.
- (3) Therefore, nobody ever knows that anything is so.

One way in which one might support the second premise in Unger's argument is by drawing on the first premise in the skeptical argument given above. According to the first premise in the skeptical argument there is this gap, which means that it is compatible with

our evidence that the belief for which we have this evidence is false. One might now say that it is because of this gap that one is never entitled to absolute certainty in matters of the physical world. Our epistemic situation is characterized by this gap, and so it is, as one might say, not such so as to make it rational for us to be absolutely certain about anything. One can also tie Unger's second premise to the notion of epistemic luck. One might say that it is because epistemic luck is an unavoidable feature of our epistemic situation that we are never entitled to be absolutely certain. There seems to be a connection between entitlement to certainty and epistemic luck in that if there is luck involved, then we are somehow not entitled to certainty.

The first premise in Unger's argument, then, is a version of the second and third premise in the skeptical argument just presented. The second and the third premise in the skeptical argument basically say that we need to close the gap, for if it is not closed then it is just lucky that our beliefs are true. Unger's first premise similarly says that it is unacceptable to not have certainty; for knowledge we need to be certain, we need to close the gap. If we tie certainty to epistemic luck, we can see that because epistemic luck and knowledge are incompatible, we need certainty for knowledge. Certainty is a way of eliminating epistemic luck, and if such epistemic luck is incompatible with knowledge then it might be that certainty is required for knowledge. – I do not want to claim that my skeptical argument is identical to Unger's, but it seems clear that they are based on similar ideas and that they draw their plausibility from a similar problem that seems to affect our epistemic situation.

Keith DeRose (1995) discusses the following skeptical argument:

- (1) I do not know that I am not a brain in a vat.
- (2) If I do not know that I am not a brain in a vat, then I do not know that I have hands.
- (3) Therefore, I do not know that I have hands.

The first premise in DeRose's argument is based on the problem of the gap. It is plausible to say that I do not know that I am not a brain in a vat because there is this gap between my evidence to the effect that I have hands and its being true that I have hands. I might have all the sensory experience as of having hands when I do not in fact have hands. This is a statement of the problem of the gap, namely that my evidence is compatible with the falsity of the belief for which I have the evidence. In this case, my evidence that I have hands is compatible with my not having hands. For example, if I were a brain in a vat, then I would have all the same sensory evidence to the effect that I have hands, but it would be false that I have hands. So the plausibility of the first premise in DeRose's argument depends on the problem of the gap stated in the first premise of the skeptical argument given above. The second premise in DeRose's argument then claims that this problem of the gap is incompatible with our knowledge of perfectly ordinary things, as for example that we have hands. The line of reasoning on which this second premise depends differs from the epistemic luck approach of the skeptical argument above, but they are similar in that they each try to spell out how the gap is a problem and leads to skepticism. Again, I do not claim that the arguments are identical, just that they share

some ideas and that thus the skeptical argument given above is not completely removed from the contemporary discussion of skepticism.

Summary Box 1: A skeptical argument based on epistemic luck

The skeptical argument discussed in this chapter goes as follows:

- (1) My evidence for a belief about the physical world is inevitably compatible with this belief being false.
- (2) If my evidence for a belief about the physical world is inevitably compatible with this belief being false, then, if the belief is true, this is a matter of epistemic luck.
- (3) If it is a matter of epistemic luck that a belief is true, then I do not know that this belief is true.
- (4) Therefore, even if the belief about the physical world for which I have the evidence is true, I do not know that it is true.

2.3 Accepting the gap

I assume and accept that the gap characterizes our actual epistemic situation and that this is inevitable. The skeptic wants to say that the gap represents a problem. In later sections I will try to show that accepting the gap does not lead to skepticism, that one can assume that there is such a gap without thereby having to endorse the claim that there is no knowledge of the physical world. In this section I will try to explain more fully what this gap consists in and what underlies my assumption that the skeptic is correct in his claim

that there is such a gap, although incorrect in his claim that this gap leads inevitably to skepticism.

We have to distinguish between the question of whether a gap characterizes our epistemic situation concerning the physical world, and the question of whether this gap represents a problem in that it leads inevitably into skepticism. I assume in response to the first question that there is such a gap, that the gap furthermore is inevitable and can never be closed. However, in response to the second question I will attempt to show that one can agree that there is such a gap and still avoid skepticism.

Let me again say what is meant by the gap. I assume that we have evidence for our beliefs about the physical world. The gap means that it is logically possible to have this evidence when the belief about the physical world that is supported by the evidence is false. Whatever evidence we have for our beliefs about the physical world is compatible with these beliefs about the physical world being false.

The simplest way of supporting this assumption of the gap is by saying that the evidence is not identical to the truth that we learn on the basis of it. So we have some evidence, glimpses and other sensory experiences. On their basis we learn about truths in the physical world, that there are barns and houses, trees and so on and so forth. The evidence that we have is not identical with the truth that we learn on the basis of it. Our sensory experiences are not identical to the trees and houses, barns and so on. It seems at

least that they are distinct in kind: barns and houses and trees exist in both space and time. Sensory experiences exist in the mind, they exist in time, but probably not in space. Even if they have an existence in space, for example if they were identical to some brain activity pattern, they still would not be identical with the trees and houses and barns. Surely, whatever goes on in the mind or the brain is distinct from trees, houses, barns and so on. One way of supporting the assumption of the gap is thus by saying that since the evidence and the truth supported by the evidence are not identical it is logically possible for the evidence to obtain when the truth does not.

Another way of supporting the assumption is by saying that the beliefs that our evidence supports inevitably involve commitments that go beyond what is contained in the evidence. We always see only a part of a barn or a house, but we believe on the basis of such evidence that there is a barn and a house, not merely those parts that we at this moment see. I do not usually believe that there is a barn façade rather than a full structure, although the façade is all I see. I recognize an object as shaped like a tree and being of a tree-like size, but I will usually believe that there is tree, not merely something shaped like a tree and of the size of a tree.

When we at all adopt beliefs about the physical world rather than about our mental states, we will do so on the basis of evidence that might obtain when nothing that it indicates is true. This is the gap. This gap is inevitable; it cannot be closed. For in an attempt to close

it we could only add further evidence, but the new set of evidence is again compatible with the falsity of what it indicates as true.

One might question the assumption that our evidence for our beliefs about the physical world consists in sensory experiences, where these are somehow mental states. What exactly distinguishes mind from body, how they are related, whether this distinction at all makes sense is not something that I can investigate. However, the gap does not depend on a problematic distinction between physical and mental things. It can be stated in everyday terms. There is always the possibility, even though we usually discount it and it seems in ordinary circumstances very slight, that one's evidence, however carefully assembled, is misleading in that what indicates is false. However careful we are, we can think up a scenario in which we would have this evidence and yet what it indicates is false.

Again, this gap should not be exaggerated, for we do not usually consider it to be incompatible with knowledge and I will try to show that this is right, that it does not lead to skepticism. However, the existence of the gap has to be acknowledged, as it seems to me. Even though when we are very careful the possibility that our evidence is misleading will be slight and in all respects negligible, it is still a possibility. This is what the gap is all about. Our evidence does not guarantee that the belief that it indicates as true is indeed true. It is possible to think of a scenario in which we would have the evidence we have and yet what it indicates is false.

In the following I will assume that there indeed is such a gap, as described just now. I take it to characterize our actual epistemic situation, and this seems to me to be one of the most basic facts about our epistemic situation concerning the physical world. The account of perceptual knowledge developed in my dissertation accepts this fact and attempts to show how knowledge is possible even though there is this gap.

2.4 Different notions of epistemic luck as a strategy for identifying a flaw in the skeptical argument

The skeptical argument stated above draws on the gap and on epistemic luck. I accept the gap. If given this we still want to identify a flaw in the skeptical argument, we need to think about epistemic luck. There are two claims made about epistemic luck in the skeptical argument. First, the claim that if there is a gap between our evidence and a true belief, then it is a matter of epistemic luck if the belief is true. Second, the claim that epistemic luck is incompatible with knowledge. The first claim, given that I accept the gap as unavoidable, would mean that epistemic luck is also an unavoidable feature of our actual epistemic situation. If this first and the second claim are both true, then, it seems, we have to accept skepticism.

The evidence for the first claim would seem to be the following. If given my evidence it is possible that the belief for which I have the evidence is true, but also possible that it is false, then if the belief is in fact true this is a matter of epistemic luck. The idea is thus that we offer to the skeptic our evidence for the true belief. The skeptic responds by

saying that the belief even given this evidence might be false. We have to agree because of the gap. Now the skeptic can plausibly say that if the belief might be false then if it is true there is an element of epistemic luck involved.

The evidence for the second claim would seem to be the following. One can construct a variety of examples in which we would say that it is a matter of epistemic luck that the belief in question is true. The Lucky case at the beginning of the chapter, a lucky guess, a Gettier case are all cases where it seems that it is a matter of epistemic luck that the belief in question is true. Furthermore, these are cases in which we would ordinarily say that we lack knowledge precisely because of this element of epistemic luck. The Lucky case, a lucky guess, a Gettier cases are all cases where the epistemic agent in question lacks knowledge precisely because it seems to be a matter of epistemic luck that the belief is true. Now the skeptic can say that these cases show that epistemic luck and knowledge are incompatible.

As I have already suggested, the key for identifying a flaw in this skeptical argument is distinguishing between different forms of epistemic luck. It is correct that in the Lucky case, in a lucky guess, in a Gettier case there is no knowledge because the form of epistemic luck present in those cases is indeed incompatible with knowledge. However, this does not mean that the form of epistemic luck that is inevitably present because of the gap is also incompatible with knowledge. The skeptic would have to show that the

mere presence of the gap means that a form of epistemic luck is present that is incompatible with knowledge.

We might want to ask at this point what in general epistemic luck is and what some of its varieties are. Probably the most general thing one can say about epistemic luck is that it concerns luck in matters of true beliefs. Epistemic luck is present in all of the following cases. Someone makes a guess for which they have no evidence, and the guess turns out to be true. Someone acquires a true belief as a matter of some lucky circumstances, being in the right place at the right time. Someone acquires some evidence for a true belief as a matter of some lucky circumstance, being in the right place at the right time. Someone has insufficient evidence or ignores an obvious defeater for a true belief. Someone arrives through confused reasoning at a true belief. Someone is in a Gettier situation. There are probably many more, but these seem to be the most obvious cases. In the fifth chapter I will sketch what renders some of these cases of epistemic luck incompatible with knowledge, while others are compatible with knowledge.

The thing to realize at this point is that there is not just one form of epistemic luck, but that it comes in many different shapes and varieties. Furthermore, we do not ordinarily think that someone who acquired some evidence by being due to a lucky circumstance in the right place at the right time lacks knowledge because of this element of epistemic luck. While the skeptic is thus right that some forms of epistemic luck are incompatible with knowledge, there seem to be a variety of forms of epistemic luck that are perfectly

compatible with knowledge. The question that we now need to raise is whether the form of epistemic luck present due to the gap is incompatible with knowledge.

Of course I believe that the gap is compatible with knowledge. So I believe that the element of epistemic luck present because of the gap is in itself compatible with knowledge. In order to argue for this claim, let me draw a distinction between two varieties of epistemic luck, namely skeptical and non-skeptical epistemic luck. Skeptical epistemic luck is the kind of epistemic luck that the skeptic needs to show to be incompatible with knowledge for his argument to work out.

Skeptical epistemic luck is the form of epistemic luck that is present whenever our evidence is compatible with the falsity of what it indicates. Whenever the gap is present, we will have an element of skeptical epistemic luck if we end up with a true belief. So when we have evidence for a true belief, but the belief might be false given the evidence, then its being true is a matter of skeptical epistemic luck. We can see that this is the form of epistemic luck that is due merely to the gap. This form of epistemic luck is present whenever there is a gap between the evidence and the truth that the evidence indicates. Thus, this is the form of epistemic luck that the skeptic would need to show to be incompatible with knowledge.

Non-skeptical epistemic luck is supposed to spell out those conditions given which a case that is also a case of skeptical epistemic luck is incompatible with knowledge. So the idea

is that some instances of skeptical epistemic luck are instances of knowledge, while others are instances of a lack of knowledge. This is due to the presence or absence of another form of epistemic luck, namely non-skeptical epistemic luck. This would mean that it is not the gap as such or skeptical epistemic luck as such that is incompatible with knowledge. Whether or not we have knowledge in a situation characterized by the gap and skeptical epistemic luck depends on whether the situation is also one characterized by non-skeptical epistemic luck.

This will have to be developed and argued for at some length. Quick intuitive support can be provided by thinking again about the example at the very beginning of this chapter. Lucky's belief that Arnold Schwarzenegger is Governor of California is not just lucky in the skeptical sense. My belief that Arnold Schwarzenegger is Governor of California is also lucky in the skeptical sense, but it is importantly different from Lucky's belief. The main question for the remainder of this chapter is to spell out this difference. The key idea is to say that Lucky's true belief is both an instance of skeptical and non-skeptical epistemic luck, and therefore not knowledge. My true belief is an instance of skeptical, but not non-skeptical epistemic luck, and therefore knowledge.

The main thing that needs to be done in the remainder of this chapter is to develop this notion of non-skeptical epistemic luck. This is a task that will be continued in the third chapter.

We can also state the strategy for responding to the skeptical argument as involving splitting the skeptical argument into two arguments. The first argument goes as follows:

- (1) My evidence for a belief about the physical world is inevitably compatible with this belief being false.
- (2) If my evidence for a belief about the physical world is inevitably compatible with this belief being false, then, if the belief is true, this is a matter of skeptical epistemic luck.
- (3) If it is a matter of skeptical epistemic luck that a belief is true, then I do not know that this belief is true.
- (4) Therefore, even if the belief about the physical world for which I have the evidence is true, I do not know that it is true.

In this argument, the third premise is not true.

The second argument goes as follows:

- (1) My evidence for a belief about the physical world is inevitably compatible with this belief being false.
- (2) If my evidence for a belief about the physical world is inevitably compatible with this belief being false, then, if the belief is true, this is a matter of non-skeptical epistemic luck.
- (3) If it is a matter of non-skeptical epistemic luck that a belief is true, then I do not know that this belief is true.
- (4) Therefore, even if the belief about the physical world for which I have the evidence is true, I do not know that it is true.

In this argument the second premise is not true.

Summary Box 2: The distinction between skeptical and non-skeptical epistemic luck

In order to criticize the skeptical argument I distinguish between skeptical and non-skeptical epistemic luck.

Skeptical epistemic luck is present whenever our evidence is compatible with the falsity of the belief for which we have the evidence. Skeptical epistemic luck is thus present whenever the gap is present. My claim is that this form of epistemic luck is unavoidable, but compatible with knowledge.

Non-skeptical epistemic luck is not present whenever the gap is present; instead it depends on further conditions. My claim is that this form of epistemic luck is incompatible with knowledge, but not unavoidable.

2.5 Non-skeptical epistemic luck

2.5.1 Searching for the conditions given which it is not lucky that the belief is true

The skeptic says that if given our evidence a belief might be true and might be false, then it is a matter of skeptical epistemic luck if the belief is true. This type of luck arises merely because of the gap, and it is an unavoidable feature of our actual epistemic situation. In order to show that skeptical epistemic luck is compatible with knowledge we need to think about another form of epistemic luck. This will allow us to arrive at the condition given which despite skeptical epistemic luck we do have knowledge.

I agree with the skeptic that given our evidence alone the belief for which we have the evidence might be true and might also be false. However, might it not be possible to add something to this evidence so that given this evidence plus the addition the belief must be true? If there is such an addition we can imagine ourselves in the following dialogue with the skeptic: We offer our evidence for a particular belief. The skeptic says that this evidence is still compatible with the belief being false. We have to agree. However, suppose that we specify further conditions given which together with the evidence the belief might not be false. In that case the skeptic cannot respond by saying that given the evidence plus this further condition the belief might still be false. So we would have conditions given which the belief for which we have the evidence has to be true. Now it seems that for the truth of the belief to be a matter of epistemic luck, we would have to be able to say that it might be false and that thus it is lucky if it is true. Given the evidence and the further conditions, this would not be possible. So there would be a sense in which it would not be a matter of epistemic luck if the belief in question is true. The belief would, rather, have to be true given the evidence and the further condition.

What drives the skeptical intuition about epistemic luck is the following. If we are able to say that a belief might be true but might also be false, then its being true is lucky. So if we can specify conditions given which the belief might not be false, then its being true is not lucky. This is correct for at least one notion of epistemic luck. We now need to ask what this additional condition could possibly be.

In accepting the gap we have accepted that whatever goes on in our mind might be in place even if the physical world in which we live is radically different from what we believe it to be. By evidence I mean internal states. In thus accepting the gap as inevitable I also accept that no additional internal states, no additional evidence could ever be this further condition. No internal condition can be added so that a belief about the physical world must be true. Because of the gap there are no such internal conditions that can by themselves guarantee that the belief about the physical world is true.

If no internal condition can do the trick, it seems clear that we are left only with an external condition. This external condition is the way the world needs to be so that given our evidence we will arrive at a true belief about the physical world. The idea is that we have certain evidence, and if the world is a certain way, then this evidence will lead us to the truth. We cannot somehow merely in virtue of our mental states, what goes on in our minds gain knowledge of the physical world. Instead, certain conditions need to be in place, the world has to be a certain way so that we will arrive at the truth by means of our evidence. I call those conditions “luck-precluding conditions.”

Before exploring what these luck-precluding conditions are, let us consider whether some other external condition could do the trick. Suppose that we say that some form of reliability is what is in addition to evidence for a true belief required for knowledge. Now the skeptic will point out that the reliability in question will not be perfect reliability, but that there is a sizable chance that given the evidence for and reliability of the belief, the

belief is false. Again we can imagine ourselves in a dialogue with skeptic. We offer both our evidence and the reliability. The skeptic responds that even given the evidence and the reliability the belief might still be false, and that thus it is lucky if the belief is true. Thus reliability would not give us the absence of the kind of epistemic luck where the skeptic can say that the belief might be false and so that it is lucky if it is true.

One might ask whether we could not argue that this kind of epistemic luck is compatible with knowledge. The strategy proposed by means of these luck-precluding conditions is to arrive at conditions given which the belief in question must be true. Instead one might propose that one could show that even if the belief in question might be false, this is compatible with knowledge. To me it seems preferable to try to eliminate this important kind of epistemic luck that allows the skeptic to say that if the belief might be false it is lucky that it is true. There is intuitive plausibility to the claim that such luck is incompatible with knowledge, in particular since it allows Gettier cases to be constructed. Showing that the kind of luck still present given evidence and reliability is compatible with knowledge would mean to show that Gettier cases are compatible with knowledge. This will be discussed in more detail in the third chapter.

It seems to me the tighter we can render the connection between evidence and true belief, the closer we come to the intuitive idea that the belief must be true, the easier will it be to argue that the remaining epistemic luck is compatible with knowledge. However, if we render the connection between evidence and true belief one that is itself merely statistical

or probable, it will be very difficult to argue that the remaining epistemic luck is compatible with knowledge. This then is what leads to the attempt to find the luck-precluding conditions given which and the evidence the belief in question has to be true.

2.5.2 Luck-precluding conditions

In this section I will explain what luck-precluding conditions are. Further details will follow in the third chapter. As I have already said, luck-precluding conditions are external conditions; they are in particular not premises that the epistemic agent needs to believe in order to arrive at perceptual knowledge. We will see in the third chapter that these external conditions have to be expected or assumed to obtain by the epistemic agent for perceptual knowledge, but the luck-precluding condition itself is a strictly external condition.

The most important thing that can be said about luck-precluding conditions is that they are regularities that, if they obtain, allow us to arrive through our evidence at the truth that this evidence indicates. We imagine ourselves having evidence for a true belief. Maybe we are driving around and have a casual glimpse of a barn. We believe truly that there is a barn in the field on the basis of a sensory experience as of a barn. The luck-precluding conditions are those physical regularities that, if they obtain, will allow us to arrive by means of this evidence at the true belief. In the example we would need a causal regularity, where parts of the barn cause in us a sensory experience as of a barn, and a

part-whole regularity, where the barn façade is attached to a barn rather than a mere façade.

These regularities have to be actual features of our world. Furthermore, they are contingent features of our world in that they might not have obtained. These regularities include causal regularities, part-whole regularities, regularities of temporal sequence, and possibly other contingent regularities that are based on physical features of the world. These regularities are not somehow abstract laws, or beliefs or the product of theories of an epistemic agent. Instead, they are features of the physical world. In the example, the part-whole regularity between barn facades and barn structures is a regularity that is present in virtue of physical features of the world, there being a barn rather than a mere façade as a matter of a regularity.

The claim is now that our evidence plus these luck-precluding conditions are not compatible with the falsity of the belief for which we have the evidence. The luck-precluding conditions are those features of the world, which, if they obtain, allow us to arrive through our evidence at the truth. Suppose thus that we have sensory experience as of a part of a barn. Suppose furthermore that there exists a regularity in the physical world according to which sensory experiences as of barns are caused by a part of a barn, and this part of a barn is attached to a barn rather than a mere façade. This regularity has to be present both in general and in this particular instance. For the regularity to be

present in a particular instance involves (among other things) that both the elements linked in the regularity are present in this instance.

So the idea is that we have evidence and this evidence is part of a regularity or set of regularities that link it, the evidence, to the truth that it, the evidence, indicates. We have the evidence, the sensory experience as of a barn. We have the truth, there is a barn in the field. The luck-precluding conditions are physical regularities that link this evidence to this truth. In the example this would be a causal and a part-whole regularity. The idea is that the luck-precluding conditions are those features of the world that allow us to arrive at the truth on the basis of the evidence. These features are contingent in that we can imagine possible worlds in which they do not obtain.

The idea is furthermore that given the evidence and these luck-precluding conditions the belief for which we have the evidence has to be true. The luck-precluding conditions guarantee that through this evidence we will arrive at a true belief. Now one might object that even if the world is generally such that sensory experiences as of barn facades lead to the true belief that there is a barn out there, sometimes we are tricked by an illusion. The luck-precluding condition does not hold without exception. But then how can a statistical regularity guarantee rather than render likely the truth of the belief in question?

In response, the regularity on the level of types of evidence and true beliefs is indeed only statistical. If we consider a lot of different cases in which someone has a sensory

experience as of a certain kind, there will be some such cases in which the belief for which we have the evidence is not true. A luck-precluding condition is statistical on the level of types of evidence and beliefs. Not all facades are parts of structures, not all sensory experiences are caused by what seems to be their cause. However, the luck-precluding condition needs to obtain both statistically on the type level and actually on the token level. Again, the luck-precluding condition is not an abstract entity, or belief or theory of an epistemic agent. It is instead a physical regularity that is based on whether facades are parts of structures, what causes some events, which events follow other events, and so on and so forth. If we consider many different cases, we will have to say that these regularities are only statistical. However, the luck-precluding condition has to obtain also actually in the case at hand. This means that in the specific case there indeed needs to be a part of a barn that causes in us the sensory experience as of a barn, and that is part of a real structure rather than a mere façade.

There are thus two requirements for a luck-precluding condition to obtain. First, it must obtain on the type level. This means that as a statistical regularity, usually and mostly the physical regularity is present. Second, it must obtain on the token level. This means (among other things) that as an actual fact of the physical world what is linked by the regularity has to be present.

Let us see how this idea of luck-precluding conditions can give us a response to the skeptic. The skeptic might say the following. It is logically possible for us to have the

evidence and end up with a false belief. So the belief might be true but might also be false given the evidence, and thus it is lucky if it is true. But then it seems that we have no knowledge that the belief is true. We can now answer the skeptic: Even though it is possible for us to have the evidence and end up with a false belief, if the luck-precluding condition obtains then the belief for which we have the evidence has to be true. If the actual world is characterized by the luck-precluding condition then the belief for which we have the evidence will be true in the actual world. Why should the mere logical possibility that the evidence leads to a false belief endanger our knowledge? The skeptic now cannot respond in the following way: Given the evidence plus the luck-precluding condition the belief might be true but might also be false. Thus there is a lot of epistemic luck involved if it turns out to be true. This is not a possible move for the skeptic because the belief for which we have the evidence must be true if the luck-precluding condition is present. It is not possible for us to have this evidence and for the luck-precluding conditions to obtain and for the belief to be false.

Consider again the distinction between skeptical and non-skeptical epistemic luck. Skeptical epistemic luck is present whenever the gap is present. This means that skeptical epistemic luck is present whenever it is possible that we have this evidence while our belief is false. Non-skeptical epistemic luck is present when the luck-precluding condition is absent. So non-skeptical epistemic luck is present when the world is not such that the evidence will lead us to a true belief. Now it seems clear that non-skeptical epistemic luck is incompatible with knowledge. If the world is such that our evidence is

systematically misleading, then we cannot gain knowledge on the basis of this evidence. So non-skeptical epistemic luck is incompatible with knowledge because the luck-precluding condition is necessary for knowledge. This means that it is necessary for us to gain knowledge on the basis of some evidence that the world is such that this evidence is neither in general nor in this instance misleading. The claim is furthermore that if non-skeptical epistemic luck is absent then skeptical epistemic luck by itself, although still present, is compatible with knowledge. So the mere presence of the gap, the mere presence of skeptical epistemic luck is compatible with knowledge.

There is an objection to which this account is subject. This objection targets the idea of a regularity. What exactly is required for there to be a regularity? One might imagine that one lives in a little barn-façade county that is part of a district in which barn facades are parts of barns, which is part of a state in which there are a lot of mere facades, and so and so forth. This is the question of what exactly is required for there to be this part-whole regularity that facades are parts of barns. I do not think that a precise quantitative answer can be given to this objection.

However, in the third chapter of the dissertation, I explain the sense in which the luck-precluding condition must be expected or assumed to obtain by the epistemic agent in question. There we will see that evidence that a luck-precluding condition is absent will be considered a defeater for the belief about the barn by those epistemic agents who expect the luck-precluding condition to obtain. The rough and ready answer to this

objection is thus that the regularity has to obtain as expected. This allows us to say that for example two mere barn façade on the other end of the planet would not count as a defeater to whether there is a barn in front of us here. However, a lot of mere barn facades right in this neighborhood would count as a defeater. This is still very vague, but I am not sure how to make it more precise. There probably is no quantitative answer to this question, and a more precise answer would depend on explaining in detail what exactly would be a defeater in an objective sense.

Summary Box 3: Luck-precluding conditions

Luck-precluding conditions are actual contingent regularities in the physical world. They include causal regularities, part-whole regularities, and regularities of temporal sequence. These luck-precluding conditions are those features of the world that allow us through our evidence to arrive at a true belief. These luck-precluding conditions ensure that our evidence is neither in general nor in this instance misleading. Non-skeptical epistemic luck is present when the luck-precluding conditions are absent.

2.5.3 Comparison with other views

In this section I will distinguish the proposal just explained from other theories in epistemology. One might ask what distinguishes these luck-precluding conditions from other proposed externalist conditions. The idea is that the luck-precluding conditions make sure that the world is actually such that through the evidence I will arrive at a true belief. The luck-precluding conditions are those features of the world that make it so that

de facto my evidence leads to a true belief in a systematic way. Reliabilism proposes a similar connection of our beliefs to the truth. What then is the main difference between reliability and these luck-precluding conditions?

There are mainly two differences between luck-precluding conditions and a reliability requirement. First, a luck-precluding condition has to obtain also in the particular case at hand. It provides more than just a statistical connection to truth. A reliably produced belief is likely, but not guaranteed to be true. It also does not make much sense to say that a particular belief is somehow reliable. Reliability is a statistical matter; it concerns the likelihood that the beliefs produced by a specific cognitive process are true. A luck-precluding condition on the other hand consists in a regularity in the physical world that has to obtain both in general and in the particular case at hand.

This is familiar from all kinds of regularities. When we come to know that there is a regularity between smoking and lung cancer, first we need to observe that there are many instances in which people who smoke also developed lung cancer. When we conclude that there is such a regularity, we do so by observing many instances of it. The luck-precluding condition is based on a regularity that has to obtain both in general and in the particular instance at hand, which involves also that the two elements connected by the regularity are both present in the case at hand. There is no such thing as reliability in one instance. All we can say about one instance is that a belief turns out to be true. Reliability is the statistical requirement that many of the beliefs produced by a particular cognitive

process are true. The luck-precluding condition thus provides a stronger connection to truth than reliability.

Second, a luck-precluding condition is not some kind of a justification. The proposal made here is not that we somehow need both an internal and an external justification for our beliefs about the physical world. It is, however, true that both internal and external requirements need to obtain for perceptual knowledge. Evidence as an internal requirement is combined with the luck-precluding condition as an external requirement. This external requirement is not a justification, but makes sure that the evidence will lead us to a true belief. In the following chapter we will see more clearly why luck-precluding conditions are not some kind of a justification.

The picture that emerges is thus importantly different from the one painted by reliabilists. It can be summarized as follows: as epistemic agents we do our part, we arrive at evidence for our beliefs. However, this evidence will lead us to a true belief only if the luck-precluding conditions also obtain. Any evidence that they do not obtain is a defeater for our belief. If we assume that the evidence is not misleading then we assume that the luck-precluding conditions obtain.¹ This will have to be developed further in the fourth chapter of my dissertation.

¹ This general idea also underlies Keith Lehrer's account of knowledge. Compare Keith Lehrer (2000).

One of the main objectives of this chapter is to introduce the distinction between skeptical and non-skeptical epistemic luck. The claim is that skeptical epistemic luck is compatible with knowledge, while non-skeptical epistemic luck is incompatible with knowledge. This strategy of defining different senses of epistemic luck, only some of which are incompatible with knowledge, has also been pursued by Duncan Pritchard and Mylan Engel. Both distinguish between two kinds of epistemic luck, one of which they claim to be incompatible with knowledge, while the other is compatible with knowledge. We should now see what the similarities and differences are between these three sets of distinctions among different kinds of epistemic luck.

In his paper “Is Epistemic Luck Compatible with Knowledge?” Mylan Engel proposes the distinction between veritic epistemic luck and evidential epistemic luck. Veritic epistemic luck arises when it is lucky given the justification that the belief for which we have the justification is true. Furthermore, Engel thinks that externalism, for example reliabilism avoids this kind of epistemic luck. The idea is that reliabilism allows us to say that the belief in question is likely to be true and thus that it is not just lucky that the belief is true. So the idea is that if it is likely that the belief is true given the justification then it is not lucky that the belief is true.

In a sense both skeptical and non-skeptical epistemic luck are a kind of veritic luck. Skeptical luck arises when the evidence is compatible with the falsity of the belief for which we have the evidence. If given the evidence it is possible that the belief for which

we have the evidence is false, then it is a matter of skeptical epistemic luck if the belief in question is true. According to Engel, veritic epistemic luck arises when there is no connection between the justification and the truth. The skeptic interprets this, according to my distinction, as demanding that the evidence should be incompatible with the falsity of the belief for which we have the evidence.

Now all instances of non-skeptical epistemic luck are also instances of skeptical epistemic luck, although it is not true that all instances of skeptical epistemic luck are also instances of non-skeptical epistemic luck. So non-skeptical epistemic luck arises when there is no luck-precluding connection that makes sure that the evidence leads us to a true belief. In a sense, skeptical epistemic luck challenges Engel's claim that we need merely a statistical connection to truth. Skeptical epistemic luck can only be avoided if given the evidence the belief for which we have the evidence must be true. Non-skeptical epistemic luck similarly challenges that it is a statistical connection to truth that we want, in that given the luck-precluding condition and the evidence, the belief in question must be true. As I have presented the skeptic, his claim is that if it is possible for a belief to be true but also false, then if it is true this is a matter of epistemic luck. Engel does not seem to address this claim. For, surely, even if there is a statistical connection between justification and true belief, if this statistical connection allows that the belief might be false, then the skeptic can still claim that it is just lucky if the belief is true.

Engel also introduces the notion of evidential epistemic luck, where it is lucky that someone is in the epistemic position he is in. This means that if it is a matter of luck that someone has a certain body of evidence or reliably working cognitive processes, this is an instance of evidential epistemic luck. I agree with his general assessment that such forms of epistemic luck are compatible with knowledge. In the fifth chapter I will provide a more systematic account of what renders a form of epistemic luck compatible or incompatible with knowledge, and this type of evidential luck will be seen to be compatible with knowledge.

Engel's general conclusion is that only externalism can be expected to allow us to avoid those forms of epistemic luck that are incompatible with knowledge. As we will see, Pritchard suggests a similar strategy, where somehow external conditions do all the heavy lifting to give us knowledge. There is something right about this in that we need certain external conditions to be in place for us to avoid the problematic forms of epistemic luck. As I suggest in this and the following chapters, luck-precluding conditions need to obtain for us to gain knowledge, where they are understood to be external conditions. However, as we will see more clearly in the fifth chapter, external conditions are only one element of the solution that would allow us to avoid problematic forms of epistemic luck. In the fifth chapter I will suggest that we need to tailor internal and external requirements for knowledge to each other, so that they fit together as seamlessly as possible. Only that can then allow us to avoid problematic forms of epistemic luck. So there is a clear requirement of an external condition for knowledge. However, without the internal

requirements also in place, we will not be able to avoid problematic forms of epistemic luck.

Let us now turn to Duncan Pritchard's (2005) distinction between veritic and reflective epistemic luck. Let us contrast the distinction between skeptical and non-skeptical epistemic luck with the distinction between veritic and reflective epistemic luck as introduced by Duncan Pritchard. Veritic epistemic luck is defined as: "It is a matter of luck that the agent's belief is true." (Pritchard 2005, p. 146) Reflective epistemic luck is defined as: "Given only what the agent is able to know by reflection alone, it is a matter of luck that her belief is true." (Pritchard 2005, p. 175)

Pritchard distinguishes both of these kinds of epistemic luck from other forms of unproblematic epistemic luck. The unproblematic senses of epistemic luck include the following. There is a sense in which it is lucky that an epistemic agent is alive, that he lives in this world at this time, that his immediate surroundings are of a specific kind, that he gains the evidence that he gains, and so on. But, as Pritchard and I agree, none of this is really problematic for knowledge. That a person is lucky to be alive is not relevant to whether or not he can come to know things. Similarly, that we are born into a world and in a specific time are all matters of luck, but this is perfectly compatible with our gaining knowledge. Similarly that some beliefs turn out to be true while others turn out to be false is by itself not a threat to knowledge. There is a sense in which somehow features of the world make our beliefs true, and it is lucky that these features of the world obtain.

Similarly it might be lucky that I remember something at the right time, or that I am in the right spot to witness a car accident, but none of this is incompatible with my gaining knowledge on the basis of such a lucky memory or such a sensory experience.

From these unproblematic senses of luck Pritchard distinguishes veritic and reflective epistemic luck as possibly incompatible with knowledge. Even though this might not be obvious from the definitions stated above we can restate the distinction in the following way: According to Pritchard, veritic luck can only and fully be avoided through an external justification, for example some form of reliability of our cognitive processes. Pritchard specifies the type of external justification that is present when veritic epistemic luck is absent as sensitivity. So the idea is that veritic luck is present if and only if an external justification is absent. Reflective luck on the other hand could, according to Pritchard, only be avoided through an internal justification that is incompatible with the falsity of the belief for which we have the justification. The basic idea that renders this distinction possible in this way is the distinction between an internal and an external justification. Furthermore, in making this distinction the assumption is that the problematic form of epistemic luck, the veritic epistemic luck, can only be avoided through an external justification.

There is a certain similarity between reflective epistemic luck and skeptical epistemic luck in that both arise because of the gap between what is internally available and the truth of the belief for which we have the evidence. However, Pritchard depends on an

external justification that is supposed to take care of epistemic luck even in the absence of any internal requirements. In contrast, I introduce an external condition that just complements the internal requirements and that is not a form of external justification. We will in the fifth chapter try to understand more fully what renders a form of epistemic luck problematic, and I will then explain why it seems to me that there really is no problematic form of epistemic luck that can be avoided through an external justification alone.

Pritchard seems to think that an external justification is what is doing the heavy lifting to avoid epistemic luck and that evidence as an internal requirement is not necessary for knowledge. I think that evidence is necessary, but that there is a gap between it and the true belief it indicates. This gap is rendered compatible with knowledge, not through an external justification but through an external condition of a different kind, the luck-precluding condition.

It is interesting to see that both Pritchard and I take skepticism seriously. However, Pritchard finds that internal evidence cannot speak against skeptical scenarios, and so he concludes that there is no internal requirement that would avoid reflective or veritic epistemic luck. This is very similar to my position in that I claim that there is this gap which renders our evidence compatible with the falsity of the belief for which we have the evidence. However, we each draw opposite conclusions from this. Pritchard develops an account of externalist justification, the sensitivity account, and claims that it eliminates

veritic epistemic luck. Furthermore, he says that it is impossible to avoid reflective epistemic luck, and that thus evidence as an internal requirement is always deficient and, as it seems to him, therefore not a necessary condition for knowledge. Basically Pritchard sees that given our internal evidence there is still the gap, and then concludes that it is not such evidence but external justification that gives us knowledge.

In contrast, I regard evidence as a necessary condition for knowledge, and it seems to me to play an important role in avoiding non-skeptical epistemic luck. However, something must be added to the evidence so that this gap is compatible with knowledge. This is what leads me to introduce an external condition, the luck-precluding condition as also required for knowledge. If we want to put it briefly, Pritchard sees what I call the problem of the gap and thinks that we need no internal but just external justification for knowledge. I also see the problem of the gap but claim that it shows that internal requirements, the evidence must be supplemented by an external condition that allows us to arrive at a true belief on the basis of our evidence.

2.6 Revisiting the skeptical argument

We have now developed the distinction between skeptical and non-skeptical epistemic luck. Let us revisit the skeptical argument and see how this distinction allows us to criticize it. Again, we have to consider the skeptical argument in its two versions. Here is the version of the skeptical argument that employs non-skeptical epistemic luck.

- (1) My evidence for a belief about the physical world is inevitably compatible with this belief being false.
- (2) If my evidence for a belief about the physical world is inevitably compatible with this belief being false, then, if the belief is true, this is a matter of non-skeptical epistemic luck.
- (3) If it is a matter of non-skeptical epistemic luck that a belief is true, then I do not know that this belief is true.
- (1) Therefore, even if the belief about the physical world for which I have the evidence is true, I do not know that it is true.

Let us examine the second premise. It says that if there is a gap between the evidence and the true belief it indicates, then it is a matter of non-skeptical epistemic luck if the belief is indeed true. What we have seen is that even though there is a gap between the evidence and true belief, if luck-precluding conditions obtain, then it is not a matter of non-skeptical epistemic luck if the belief is indeed true. Given the evidence alone there is a gap between it and the truth of the belief for which we have the evidence. However, the gap by itself does not lead to non-skeptical epistemic luck. For non-skeptical epistemic luck to be present we need the absence of luck-precluding conditions, not merely the gap.

Let us now look at the second version of the skeptical argument, which employs the skeptical notion of epistemic luck.

- (1) My evidence for a belief about the physical world is inevitably compatible with this belief being false.
- (2) If my evidence for a belief about the physical world is inevitably compatible with this belief being false, then, if the belief is true, this is a matter of skeptical epistemic luck.
- (3) If it is a matter of skeptical epistemic luck that a belief is true, then I do not know that this belief is true.
- (4) Therefore, even if the belief about the physical world for which I have the evidence is true, I do not know that it is true.

In this argument the second premise is true. The question is, however, if the third premise is also true. Suppose that it is not a matter of non-skeptical epistemic luck that a belief is true. So we suppose that we have some evidence for a belief and luck-precluding conditions obtain. Even though it is still a matter of skeptical epistemic luck if the belief is true, the situation now does not seem to be incompatible with knowledge. We have evidence for a belief. The world is in fact such so that this evidence will lead us to a true belief. At this point it is not clear whether this is sufficient for knowledge, but the skeptic has to say that this situation is incompatible with knowledge, that it is sufficient for our not having knowledge. And that now seems implausible. So I suggest that this situation, although maybe not containing everything we would need for knowledge, at least seems to be compatible with knowledge. It does not seem as though the situation we have described is one where we need to despair of ever gaining knowledge. Thus the third premise in this argument is false.

Above we mentioned that the skeptical argument criticized in this chapter is similar to skeptical arguments discussed by Unger and DeRose. Let us now see whether what we have said so far would also furnish us with a criticism of these two skeptical arguments. Unger's (1974) argument goes as follows:

- (1) If someone knows something to be so, then it is all right for the person to be absolutely certain that it is so.
- (2) It is never all right for anyone to be absolutely certain that anything is so.
- (3) Therefore, nobody ever knows that anything is so.

The question to think about is whether certainty has to reflect the absence of skeptical or merely non-skeptical epistemic luck. If we are certain only if it is not a matter of skeptical epistemic luck that a belief is true, then we would have to accept Unger's second premise. Skeptical epistemic luck is unavoidable given the gap. However, as I have suggested and will try to show more clearly in the third chapter, skeptical epistemic luck is compatible with knowledge. So if certainty depends on the absence of skeptical epistemic luck, but skeptical epistemic luck is compatible with knowledge, this would give us a way of criticizing the first premise. We could say that certainty based on the absence of skeptical epistemic luck is not required for knowledge and that thus the first premise is not true.

Now, what about saying that certainty demands the absence of non-skeptical epistemic luck? As we have seen, non-skeptical epistemic luck is incompatible with knowledge and so certainty that depends on the absence of non-skeptical epistemic luck would be

required for knowledge. However, this non-skeptical epistemic luck is not unavoidable; it might be absent, and so if this is what is required for certainty then Unger's second premise seems problematic. For we might have a belief that is true without any non-skeptical epistemic luck. In that case we would be entitled to certainty in that sense of certainty. And so the second premise in Unger's argument would not be true. – I do not claim that this is all one would need to say to show Unger's argument to be flawed. However, it might at least seem that the ideas and in particular the distinction between skeptical and non-skeptical epistemic luck might be helpful for responding to Unger's skeptical arguments.

Let us now look again at DeRose's (1995) skeptical argument:

- (1) I do not know that I am not a brain in a vat.
- (2) If I do not know that I am not a brain in a vat, then I do not know that I have hands.
- (3) Therefore, I do not know that I have hands.

Now, if I am indeed a brain in a vat, then I do not know that I am not a brain in a vat. But suppose that it is true that I have hands and that I have this belief and all the normal evidence for the belief. This evidence is also evidence for the belief that I am not a brain in a vat. I take it that the plausibility of DeRose's skeptical argument depends on the idea that this evidence is not good enough to give us knowledge because it is compatible with my indeed being a brain in a vat. This is the problem of the gap all over again. What I try to argue is that if our evidence is supported by an appropriate external condition, the luck-precluding condition, then we can know that we have hands and that we are not

brains in vats. This knowledge however is not a matter merely of what is internally available to us. An external condition has to be added to this internal condition, and it is possible that we live in a world in which this external condition is absent.

One of the things that I want to argue for in this dissertation is that we can know that we have hands and that we are not brains in vats on the basis of strong evidence and favorable external circumstances. Part of this argument depends on the distinction between skeptical and non-skeptical epistemic luck, and the idea that skeptical epistemic luck by itself is compatible with knowledge. The plausibility of the claim that we do not know that we are not brains in vats depends on the limits of our internal evidence, the gap, and the intuition that something more is required for knowledge, that the gap must be closed by what is internally available to us. This however seems impossible; internal conditions by themselves cannot close the gap. And so we need to consider what external conditions might instead create the conditions in which what is internally available will lead us to a true belief. – Again I do not claim that this is everything one would need to say to criticize this skeptical argument, but it might seem that the ideas presented in this chapter might be helpful for responding also to this other skeptical argument.

Let me emphasize that there is no charge of equivocation here, in the sense that there would be a willful glossing over or ignoring of the distinction between skeptical and non-skeptical epistemic luck. All I am trying to do is suggest that there is a plausible distinction between these two defined senses of epistemic luck, and that it might be

interesting to consider whether other skeptical arguments might also be criticized on the basis of this distinction.

Summary Box 4: Criticism of the skeptical argument

Consider the skeptical argument where “epistemic luck” always means “non-skeptical epistemic luck.” Here is the criticism of this argument: Given the evidence alone there is a gap between it and the truth of the belief for which we have the evidence. However, according to the non-skeptical account of epistemic luck the evidence can be supported by luck-precluding conditions. If these luck-precluding conditions obtain, then non-skeptical epistemic luck is absent. So the second premise in this argument is not true.

Consider the skeptical argument where “epistemic luck” always means “skeptical epistemic luck.” Here is the criticism of this argument: Suppose that it is not a matter of non-skeptical epistemic luck that a belief is true. So we suppose that we have some evidence for a belief and luck-precluding conditions obtain. Even though it is still a matter of skeptical epistemic luck if the belief is true, the situation now does not seem to be incompatible with knowledge. We have evidence for a belief, and the world is in fact such so that this evidence will lead us to the truth. It is not clear whether this would be sufficient for knowledge, but the skeptic has to say that this situation is incompatible with knowledge, that it is sufficient for our not having knowledge. And that now seems implausible. Thus the third premise in this argument is not true.

2.7 The problem and the proposal revisited

In this chapter we have looked at two varieties of epistemic luck, skeptical and non-skeptical epistemic luck. Skeptical epistemic luck is present merely in virtue of the gap. Since I regard the gap as an unavoidable aspect of our actual epistemic situation, the skeptic is right in his claim that a certain amount of epistemic luck is always present when we arrive at a true belief. Lucky's belief in the example at the beginning of the chapter is true as a matter of skeptical epistemic luck, but so is our belief that Arnold Schwarzenegger is Governor of California. However, I take it that we would intuitively want to say that Lucky's belief is not knowledge because of epistemic luck, but that our belief is knowledge despite the element of epistemic luck. This leads us to think about how to distinguish between those instances of skeptical epistemic luck in which there is no knowledge and those in which there is knowledge.

The idea of non-skeptical epistemic luck is supposed to develop the conditions given which despite skeptical epistemic luck we can have knowledge. Non-skeptical epistemic luck is present whenever luck-precluding conditions are absent. These luck-precluding conditions are physical regularities in the world, including causal or part-whole regularities, or regularities of temporal sequence. These regularities connect our evidence, the sensory experiences, to the truth that they indicate. If these regularities obtain in the physical world, then our evidence will lead us to a true belief. The intuitive idea of the luck-precluding conditions is that they are those features of the world that allow us to arrive at a true belief on the basis of our evidence.

In the Lucky example there is no regularity that connects the events described in a novel written 20 years ago to what is going on in California right now. The world is not characterized by a regularity that would render reading this novel written many years ago good evidence for truths about what is happening in California. This means that it is also a matter of non-skeptical epistemic luck that Lucky's belief is true. If we have the intuition that Lucky does not have knowledge, then the absence of non-skeptical epistemic luck is a necessary condition for knowledge. Since non-skeptical epistemic luck is present whenever a luck-precluding condition is absent, we get the result that the presence of luck-precluding conditions is necessary for knowledge. What this demands is that the world has to be characterized by regularities that allow us to arrive on the basis of our evidence at the truth. These regularities are necessary for us to gain perceptual knowledge.

What we have seen now is at least that examples like the Lucky case can be understood without assuming that all kinds of epistemic luck, or skeptical epistemic luck as such, are incompatible with knowledge. We again get the result that while some varieties of epistemic luck are clearly incompatible with knowledge, and this includes non-skeptical epistemic luck, others are compatible with knowledge, including skeptical epistemic luck.

The skeptical argument discussed in this chapter can then plausibly be criticized on the basis of this distinction. The skeptic would have to show that all instances of skeptical

epistemic luck are incompatible with knowledge, and this cannot be done by pointing to the Lucky case or some other similar case. The intuition that Lucky lacks knowledge because of epistemic luck is strong and plausible, but the example does not support the more radical claim that all forms of epistemic luck or all instances of skeptical epistemic luck are incompatible with knowledge. However, the skeptic would need to make this stronger claim beyond the obvious truth that some forms of epistemic luck, including those present in the Lucky case, are incompatible with knowledge.

Let me conclude this chapter by sketching what I want to do in the third chapter. The third chapter studies the kind of epistemic luck present in Gettier cases. Furthermore, I claim that this is non-skeptical epistemic luck, as defined in this chapter. Gettier cases are a challenge to whether we can ever come up with sufficient conditions for knowledge. For it seems that there is a Gettier counterexample to almost every proposed set of sufficient conditions for knowledge. By examining how Gettier cases arise, I hope to be able to come up with those conditions for perceptual knowledge that would not be subject to a Gettier-style counterexample.

3. EPISTEMIC LUCK AND GETTIER CASES

3.1 The problem and the proposal: How luck-precluding conditions allow us to avoid Gettier cases

Suppose that a journalist, named “Sokal,” decides one day to test the journalistic integrity and intellectual standards of the very reputable newspaper that employs him. In order to do so this journalist calls up his newspaper only a few minutes before it goes to print and gives the following excited and completely fictitious report: “Today an escaped circus elephant crossed I-10 between Tucson and Phoenix and caused a 32 car pileup. Luckily nobody, including the elephant, was hurt.” To his superiors at the newspaper this journalist is known as reliable and trustworthy, and so it is decided at the last moment to put his report on the front page of the newspaper.

Now suppose that a reader of this newspaper, named “Lucky,” has ever since she first got the newspaper delivered 32 years ago found it to be thoroughly reliable and trustworthy. Today Lucky is running a little late and has barely time to glance at the headlines. All she reads is: “Escaped Circus Elephant Causes 32 Car Pileup on I-10.” Thus Lucky acquires the justified belief that an escaped circus elephant caused a 32 car pileup on I-10. Now suppose that, although not on I-10 between Tucson and Phoenix, but somewhere close to the state border to California, an escaped circus elephant caused a 32-car pileup on I-10. So Lucky’s belief is justified and true, but it seems that it is too lucky to be knowledge.

This is an example of a Gettier case.² Lucky has a justified true belief but no knowledge of the elephant caused accident. Moreover, it seems that it is because of epistemic luck that Lucky has no knowledge. We want to say that it is a matter of mere epistemic luck that her justified belief is true, and that this epistemic luck is incompatible with knowledge.

The question arises what characterizes the type of luck that allows us to construct a Gettier case. We should want to explain precisely this type of luck and the conditions given which it would be absent. The proposal in this chapter is that Gettier cases arise because of non-skeptical epistemic luck. It is this type of luck that renders it possible to construct a Gettier case.

Now, Gettier cases are usually taken to be counterexamples to a proposal of sufficient conditions for knowledge. Often it is said in response to a Gettier case that at best the proposed conditions are necessary, but not jointly sufficient. The project of the dissertation is to develop a set of necessary and sufficient conditions for perceptual knowledge. This means, among other things, that I need an argument to the effect that for the proposed set of conditions for perceptual knowledge there is no Gettier case that would show them to be not sufficient. In order to do so, I propose to examine more closely the type of luck involved in Gettier cases.

² The Gettier cases were first presented in Edmund Gettier (1963).

It does not seem to me that merely by showing that there can be no Gettier case to a proposed set of conditions for perceptual knowledge one would already have shown that these conditions are individually necessary and jointly sufficient. This is a task I will leave to the fifth chapter of my dissertation. The project of this chapter is more limited. I will characterize more precisely non-skeptical epistemic luck. Then I will explain that if non-skeptical epistemic luck is present, we can construct a Gettier case. Furthermore, I will present an argument that if non-skeptical epistemic luck is absent and certain further conditions obtain, then no Gettier case can be constructed. Again, this is not already an argument that any proposed set of conditions for perceptual knowledge would be jointly sufficient. It is merely an analysis of Gettier cases and the proposal that if non-skeptical epistemic luck is absent and certain further conditions obtain, then there is no room for Gettier cases.

The project of this chapter can also be briefly stated in the following ways: Linda Zagzebski (1994) argues that the Gettier cases show that there is always an element of luck involved in our knowledge. This chapter shows that this luck cannot be non-skeptical epistemic luck, and that the presence of skeptical epistemic luck does not generate a Gettier intuition. Peter Klein (1971) argues that Gettier cases arise because of our intuition that a felicitous coincidence is incompatible with knowledge. This chapter shows how the presence of non-skeptical epistemic luck allows for felicitous coincidences that are incompatible with knowledge. However, if non-skeptical epistemic luck is absent and certain further conditions obtain, then there is no room for such

felicitous coincidences. Keith Lehrer (2006) argues that to avoid Gettier cases the justification must be the explanation for the true belief. This chapter explores in what way the luck-precluding condition would explain how given the evidence we arrive at a true belief, and how this avoids Gettier cases.

3.2 The presence of non-skeptical epistemic luck

In this section I want to understand better what exactly happens when non-skeptical epistemic luck is present in a given case. In later sections I will propose that the type of luck involved in Gettier cases is non-skeptical epistemic luck. However, before looking at Gettier cases, let us try to continue in detail the characterization of non-skeptical epistemic luck begun in the preceding chapter.

To begin let us again consider the distinction between skeptical and non-skeptical epistemic luck. Skeptical epistemic luck is present merely in virtue of the gap. This means that whenever it is compatible with my evidence that the belief for which I have the evidence is false, skeptical epistemic luck is present if I end up with a true belief. Non-skeptical epistemic luck is present when a luck-precluding condition is absent. We now need to consider how the luck-precluding condition might be absent.

Luck-precluding conditions are regularities in the physical world. Examples are, again, causal, part-whole regularities, or regularities of temporal sequence. In the second chapter I said that for a luck-precluding condition to be present it has to be present both on the

type and on the token level. The regularity has to be present both in general and in the particular case at hand. The regularity could thus be absent in two ways. Either it is not present in general and therefore absent. Or it is present in general but not in this instance and therefore absent. Consider again the regularity that smoking causes lung cancer. This regularity could be absent in two ways. We can imagine a world in which smoking does not increase the probability of lung cancer. In that world the regularity would be in general absent. We can also imagine a world in which although in general smoking causes lung cancer, one individual, maybe in virtue of his genetic makeup, is able to smoke without ever ending up with lung cancer. In that case the regularity would be present in general but absent in the particular case.

We thus see that a luck-precluding condition is not present if it is absent on the token level. In the next section I will look at evidence and then contrast it with the luck-precluding condition. We will see that they differ in a certain property concerning types and tokens of evidence and luck-precluding conditions. This will also be seen to support the claim that the luck-precluding condition is not some kind of an external justification.

3.2.1 Types and tokens of evidence

In this section I will discuss that evidence has a certain property that I call “type-token transparency.” To begin, let us consider how we judge whether something is evidence for a belief. It is common when we ask whether we have evidence for a particular belief to think about similar beliefs where we have the same evidence. When we ask whether to

believe what a witness tells us, we ask whether he has told the truth about similar matters in the past. When we ask whether to trust our memory in a particular matter, we ask whether we were able to trust it in the past on similar matters in similar circumstances. When we ask whether we should believe what a newspaper writes or what an expert testifies we again ask about similar cases.

What exactly renders these similar cases similar? Similar cases are instances where we have the same type of evidence for the same type of belief. In the memory case we ask whether the same type of evidence, a memory, for a similar belief, a belief of the same type, has led me to a true belief in the past. In the witness case we ask whether the same type of evidence, testimony by this witness, for a similar belief, a belief of the same type, has led me to a true belief in the past. So when we look into similar cases we look for two things: same type of evidence, for example memory, sensory experience, testimony, and same type of belief, for example about a field of expertise, or about what a witness saw and heard, or something similar. Sometimes we will have to delineate the types in question so that they cover only a few instance, or alternatively that they cover a wide range of similar cases. This talk of types of evidence and types of belief is somewhat vague, and I will later on explain further how we delineate these types.

For now what we need to see is that when judging whether something is evidence for a particular belief we will often proceed by looking at such similar cases, cases where we have the same type of evidence for the same type of belief. This reflects a property of

evidence as such. Evidence is always attached to types of evidence for types of beliefs. When a memory of the time of a lunch appointment is evidence for my belief that the lunch appointment is at such and such a time, this means that a similar type of memory will also be evidence for a similar type of belief in other cases.

Our evidence is, however, defeasible. So if our circumstances change, if we acquire a defeater, we might not believe what our evidence indicates. Nevertheless there is a sense in which this type of evidence defeasibly indicates the truth of a type of belief. Evidence is not a matter of particular beliefs and how they are related to particular sensory experiences and other beliefs. Evidence always comes in types, all tokens of which have the same evidential relationships to the same beliefs. This is supported by our intuition that if I can trust my memory or senses or testimony on something once, I can trust it again, unless of course there is other evidence, for example defeaters, available to me.

We can make this more precise. Given any type of evidence and any type of belief, if this evidence indicates that the belief is true in one token, so does the same type of evidence for the same type of belief indicate the truth of this belief in any other token. Evidence is a matter of types of evidence and types of beliefs. If we can say that a type of evidence indicates the truth of a type of beliefs, then for any token of an evidence-belief pair of this type, the evidence indicates the truth of the belief. Furthermore, if any token of an evidence-belief pair is such that the evidence indicates the truth of the belief, then this

type of evidence indicates the truth of this type of belief. I will call this property of evidence “type-token-transparency.”

Type-token-transparency means that evidence is something general; it does not attach idiosyncratically to particular belief-evidence pairs. Instead it attaches to types. This is something that we all accept when we try to assess whether some evidence indicates that a belief is true. From the type-token transparency it follows that if some evidence indicates the truth of a belief, while some other evidence does not indicate the truth of some other belief, then these two evidence-belief pairs belong to different types. When we assess whether we have good evidence for a belief we often try to delineate such types. For example, I might trust my memory on some matters, but not when I have had a lot of beer, or when I was half asleep, or when it happened a long time ago and I did not attach any importance to it at that time, and so on. This is a very familiar process. We ask whether a given memory belongs to a type of evidence for a type of belief that we can trust. Delineating types of evidence that indicate the truth of types of beliefs is a familiar part of our epistemic life.

Something similar goes on in the case of sensory perception. We trust visual perception in many cases, but not when it is very dark, or there is a blinding light, or when we have had too much beer, or when we do not get a good look, or something similar. Suppose that in a country all the barn facades are attached to barns, but some of the house-facades are not attached to houses. When visiting this country, we would probably begin by

taking all the sensory evidence of facades to indicate that there is a building there. Then we might delineate the type more carefully and might conclude that sensory experience as of a barn-façade is good evidence that there is a barn, but sensory evidence of a house-façade is not good evidence of a house. Again, this involves delineating the type of evidence for a type of belief that it indicates as true.

The belief for which we have such evidence might still turn out to be false. It might be that in general a barn-façade is good evidence for there being a barn, but maybe one epistemology-minded farmer decided to put up a barn-façade that is not attached to a barn, in order to mislead those who are innocent of epistemology. Upon seeing this non-attached barn façade, you would have evidence for believing that there is a barn, but the belief would be false. It seems to me that we would intuitively say when confronted with such a situation that we ended up with a false belief, but not that we did not have evidence for this false belief. If, however, we find out one day that we have for the past 25 years lived in barn façade country, without ever realizing it, we will probably conclude that we never really had evidence for the belief that there is a barn. So if something is generally misleading, we will stop considering it as evidence for the belief in question. However, if something is misleading only in isolated instances, we will say that we have evidence for a false belief. Again, this underlines the idea that evidence is a matter of types.

Type-token transparency is the property of evidence that renders it possible to have good evidence for a false belief. Furthermore, this is, in part, what allows us to design Gettier cases. Because of type-token transparency it is possible to have good evidence for a false belief. In the example given above, Lucky has good evidence for her belief in what the newspaper writes, given that she has always found over a long period of time the newspaper to be reliable and trustworthy. In this instance, however, given full facts about this particular evidence token, we would expect that this belief for which Lucky has good evidence is false. After all, the headline in the newspaper is due to a fictitious report by Sokal. It is a matter of mere epistemic luck that the justified belief then turns out to be true. Again, this epistemic luck seems to be incompatible with knowledge.

3.2.2 Types and tokens of luck-precluding conditions

We have seen that evidence is type-token transparent. Now we need to see that luck-precluding conditions are not type-token transparent, instead they are type-token opaque. To begin, let us think more carefully about these luck-precluding conditions. We might also note at this point that it is this difference between evidence and luck-precluding conditions that makes the decision not to call the luck-precluding condition some kind of an external justification more than a merely terminological decision. Luck-precluding conditions are type-token opaque while even external justification, as it is usually understood, is type-token transparent. For example, if a belief is externally justified just in case it is produced by a reliable cognitive process, then this allows for instances where we have a justified but false belief. A belief might be produced by a generally reliable

cognitive process and still be false in one instance. Even external justification seems to be type-token transparent, but the luck-precluding conditions are, as we will see, type-token opaque. This makes it clear that the luck-precluding condition is not some additional kind of external justification. It is an external condition of a different kind, and not an external justification.

Let us now see how it is that luck-precluding conditions are type-token opaque. Luck-precluding conditions are contingent regularities in the world. There are different kinds of such regularities. We have causal regularities, part-whole regularities, and regularities of temporal sequence, and there are probably other types of contingent regularities. The point is that the physical world in which we live is full of such regularities. We will see clearly in this chapter that without such regularities we could not gain any perceptual knowledge.

Luck-precluding conditions are the physical regularities that link the evidence to the true belief. They are those conditions in the physical world given which our evidence will lead us to a true belief. As already suggested, these regularities might be absent in two ways. They can be generally absent, or they might be generally present but absent in a particular case.

Let us first look at examples where the regularity is in general absent. There is no general regularity of going for a walk in the evening causing lung cancer. A luck-precluding

condition needs to be a regularity that links evidence to the truth it indicates. In such an example, we can imagine a world in which it is never the case that a barn façade is attached to a barn. All barn facades in that world are mere facades, maybe they function as decorative screens of some kind. In that world the luck-precluding condition linking evidence of a barn façade to the truth that there is a barn would be absent because the general regularity is absent.

Let us now consider a case where the regularity is generally present, although it is absent in a particular case. We can again consider the philosopher farmer who delights in presenting various illusions and misleading evidence on his farm to trick unsuspecting tourists. This philosopher farmer might be imagined to put up a single barn façade that is not attached to a barn. We imagine that there are no other philosopher farmers around and that all other barn facades are parts of barns. In this case the regularity between barn facades and barns would be present in general, but not on the farm of the imagined philosopher farmer.

Again let us remember that the luck-precluding conditions make sure that our evidence will lead us to a true belief. If the luck-precluding condition is absent, then our evidence will be misleading. If the luck-precluding condition is absent on the type-level, then in general a type of evidence is misleading. If the luck-precluding condition is absent on the token-level then a particular token of evidence is misleading, although in general this type of evidence is not misleading. This independence of types and tokens of luck-

precluding conditions is what is meant when we say that luck-precluding conditions are type-token opaque.

Let us discuss one example in detail. Suppose that I am going on a cross-country drive and upon seeing a barn-façade I arrive at the belief that there is a barn in the field. My evidence consists in a sensory experience as of a barn façade. My belief for which I have this evidence is that there is a barn in the field. There are two luck-precluding conditions that need to be in place for non-skeptical epistemic luck to be absent in this case. First, it must be that there is a luck-precluding condition that consists in the contingent regularity of barn facades causing in me a sensory experience as of a barn façade. Furthermore, there must be a luck-precluding condition that consists in the contingent regularity of barn facades being parts of barns. So even in this simple example, we have two layers of luck-precluding conditions. In more complicated examples, we will have even more layers of luck-precluding conditions.

Let us pick one of these luck-precluding conditions, and discuss what happens if it is absent either on the type or the token level. For simplicity let us discuss the luck-precluding condition that depends on the part-whole relationships of barns and their facades. It might be that there is some place where this regularity does in general not hold. In that case the luck-precluding condition would be absent on the type level. If my evidence for the belief that there is a barn consists in my seeing a barn façade, this evidence would not in that case be linked by a luck-precluding condition to the true belief

that there is a barn out there. This is the barn-façade county case. In this case the world renders evidence as of a barn façade misleading concerning the belief that there is a barn in the field. The world would not be such so that on the basis of seeing a barn façade we will be led to a true belief that there is a barn in the field.

It might also be that there is a general regularity but in one instance it is not present. This would be the case in a country where barn facades are usually attached to barns, but one farmer has set up a single barn-façade that is not attached to a barn. In that case the luck-precluding condition is present on the type level, but not on the token level. It is in general the case that the evidence of seeing a barn façade indicates the truth that there is a barn, but the luck-precluding condition is absent in this particular case. So in this particular case the evidence is misleading. The world renders in this instance the evidence of the barn façade misleading concerning the belief that there is a barn in the field.

We should note that we usually assume that the luck-precluding conditions are present. The situation of driving around and coming to believe that there is a barn in the field is perfectly ordinary. However, we will usually only see a façade, a part of a barn that we take to be a part of a structure rather than a mere façade. In section 2.4 we will discuss in detail what this assumption or expectation that the luck-precluding condition obtains consists in.

We should say what is the case when a luck-precluding condition is present on the token level. That a luck-precluding condition is present on the token level means more than merely that the two elements connected by the regularity are each present. The luck-precluding condition consists in a regularity of a certain nature, for example a causal, part-whole, or temporal sequence regularity. It is possible that both the elements of the regularity are present, when they are not connected by the regularity. Suppose, for example, that I have a sensory experience as of a sheep that is caused by a sheepdog, and that there is a sheep in the far distance. In that case both elements connected by the regularity are present (the sensory experience as of a sheep and the sheep), but they are not properly connected through a causal regularity (the sheep does not cause the sensory experience as of a sheep). For a regularity to be present on the token level three conditions need to be satisfied. First, both the elements connected by the regularity need to be present. Second, they need to be connected in the way that characterizes the regularity, for example one needs to be the cause of the other, one needs to be a part of the other, or one must follow the other in close temporal sequence. Third, there needs to be a general regularity of the same type. There are Gettier cases in which the luck-precluding condition is absent on the token level because of the second point of the three just mentioned.

Summary Box 5: Type-token transparent evidence and type-token opaque luck-precluding conditions

Evidence is type-token transparent: Given any type of evidence and any type of belief, if the evidence indicates the truth of the belief in one token, so does the same type of evidence for the same type of belief in any other token.

Luck-precluding conditions are type-token opaque: The luck-precluding conditions are contingent regularities that can be absent in two ways. The luck-precluding condition might be absent on the type-level, if there is no general regularity. This means that in general this type of evidence is misleading. The luck-precluding condition might be absent on the token-level if there is a general regularity, which is, however, not present in the particular case at hand. This means that although in general a certain type of evidence leads to a certain type of true belief, in this instance the evidence is misleading.

3.2.3 Luck-precluding conditions as explanations

In this section I will explain in what sense luck-precluding conditions can be understood as explanations. Again, the luck-precluding conditions are not beliefs or premises or theories of an epistemic agent. They are not something that occurs in the mind of an epistemic agent. Instead they are external conditions. They are regularities in the physical world, and they are present in virtue of physical feature of the world. Examples are part-whole or causal regularities, or regularities of temporal sequence. So the luck-precluding conditions are not explanations in the sense that they are somehow a theory or a belief of

an epistemic agent that supports his evidence. They are not explanations that anyone needs to possess. They are not explanations that anyone needs to come up with.

Nevertheless there is a sense in which the luck-precluding conditions explain how on the basis of the evidence an epistemic agent arrives at the truth. The luck-precluding conditions are some of the physical regularities in the world. They, so to speak, set the stage for an epistemic agent to be able to arrive at the truth on the basis of his evidence. They allow us to move from our evidence to the truth. They, so to speak, provide the channels or venues for us to reach by means of our evidence the truth. In that sense do the luck-precluding conditions explain how we ever manage to reach the truth on the basis of our evidence.

The picture that emerges is the following. We only have these little bits of evidence, these glimpses and casual touches, snatches of sound, and so on. On the basis of this evidence we arrive at beliefs about something that goes beyond this evidence. We believe things about houses and barns, trees and flowers, towns and rivers on the basis of evidence that is compatible with the falsity of all of these beliefs. We might indeed be dreaming or be deceived in some way, in some or possibly all such cases. So how do we ever arrive at the truth on the basis of this evidence? The answer is that the world needs to be such that this evidence will lead us to the truth. And this is the work the luck-precluding conditions do.

The physical world is characterized by certain regularities. Some of these regularities allow us to arrive at true beliefs on the basis of a certain type of evidence. These regularities provide the channels or venues for us to arrive at the truth on the basis of certain evidence. In that, these regularities, the luck-precluding conditions, are the explanations of how given our evidence we ever manage to reach the truth.

3.3 The type of luck present in Gettier cases is non-skeptical epistemic luck

3.3.1 The proposal

In this section I will do two things. First, we will ask what type of luck is present in a Gettier case. I will claim that it is non-skeptical epistemic luck that is present in Gettier cases. Second, this will suggest that the luck-precluding condition is a necessary condition for perceptual knowledge. If in a Gettier case non-skeptical epistemic luck is present, and if non-skeptical epistemic luck is present whenever a luck-precluding condition is absent, then, if we lack knowledge in a Gettier case, the luck-precluding condition is necessary for knowledge.

Let me explain further what supports the claim that in a Gettier case non-skeptical epistemic luck is present. Consider again the Gettier case sketched at the beginning of this paper. In this case we assume that Lucky is justified in her belief in the elephant-caused accident. She is justified because it is a reputable newspaper and she has always found it to be reliable and trustworthy. Furthermore, the luck-precluding condition is present as a general regularity. It is in general, as we suppose, the case that reports in this

newspaper are mostly true. We suppose that there is a general regularity between a printed report in this newspaper and the report being true. This luck-precluding condition is present on the type level.

However, in this instance Sokal has called in his fictitious report. So in this instance there is no connection between the printed report and its truth. We suppose that Sokal just makes up his report, he has no reason for calling in this particular report rather than any other. There is no connection between a printed report in the newspaper due to a fictitious report by Sokal and the report being true. There is no regularity that would link evidence due to a fictitious report by Sokal to the truth. We can make this more evident by supposing that Sokal calls in other fictitious reports, and none of them happens to be true. Thus it seems that the luck-precluding condition is absent. This evidence, due to a fictitious report by Sokal, is not linked to the truth. The world is not in general such so that a fictitious report by Sokal will lead us to the truth. This, as a first attempt, is what supports the claim that non-skeptical epistemic luck is the type of epistemic luck that leads to a Gettier case.

We can now also explain the proposal that luck-precluding conditions are necessary conditions for perceptual knowledge. If a luck-precluding condition is absent, then the world renders our evidence misleading. There is no systematic connection between the evidence and the truth in the absence of the luck-precluding conditions. We can thus explain the Gettier intuition. It seems to be a common intuition that in a Gettier case there

is no knowledge. If a Gettier case arises because of the absence of a luck-precluding condition, we can see why this should be so. In the absence of a luck-precluding condition our evidence is misleading. Even though Lucky has good evidence for her belief in the elephant-caused accident, it is just lucky that her belief turns out to be true. This is because the world just does not usually allow us to arrive on the basis of this evidence at the truth.

If it is correct that Gettier cases involve the absence of a luck-precluding condition, we have mutual support for the Gettier intuition and the claim that luck-precluding conditions are necessary conditions for knowledge. The Gettier intuition is the intuition that there is no knowledge in a Gettier case. If Gettier cases involve the absence of luck-precluding conditions, then this Gettier intuition supports the claim that luck-precluding conditions are necessary conditions for knowledge. Furthermore, it seems that we can explain clearly why the absence of luck-precluding conditions means that there is no knowledge. If the world renders our evidence misleading, then surely we cannot gain knowledge on the basis of this evidence. Thus, if Gettier cases involve the absence of luck-precluding conditions, then this analysis would support and explain the Gettier intuition.

Let us think again about what a luck-precluding condition does. Basically it is the way the world needs to be so that the evidence will lead us to a true belief. So we have evidence for a particular belief. In the case of perceptual knowledge the evidence

includes a sensory experience as of, for example, typing on a computer. This is evidence for a particular belief about the physical world, for example that I am typing on a computer. What a luck-precluding condition does is it allows me to arrive through this evidence at a true belief in a systematic way.

As explained, the luck-precluding condition consists in a regularity in the physical world. The idea is that the world needs to be a certain way so as to allow us through the evidence to arrive at a true belief. Again, the luck-precluding conditions are not premises or beliefs or theories of an epistemic agent. They are strictly external conditions, the physical regularities in the world. I do not think that there is an inference involved in our gaining perceptual knowledge. We have sensory experience as evidence for a specific belief about the physical world. For this evidence to lead us to the true belief, the world needs to be a certain way. It is not the case that we somehow need the luck-precluding condition as a premise in an argument for a belief about the physical world. No such argument seems to me to be involved in perceptual knowledge.

What happens when the luck-precluding condition is absent is basically that this evidence is either in general or in this instance not a way of arriving at a true belief. Suppose again that I have evidence for the belief that I am typing at a computer. This evidence consists in the various sensory experiences as of my typing at a computer. But now suppose that the luck-precluding condition is absent. Let us again discuss the two ways in which the luck-precluding condition might be absent.

Suppose that the regularity is in general absent, and that thus the luck-precluding condition is absent on the type level. This is to suppose that, for example, there is never a computer that causes in me sensory experiences as of typing on a computer. We can imagine a person who has a body and lives in a normal world and is normal in all respects except when it comes to interacting with a computer. Suppose that this person leads a normal life; he gets up in the morning, has coffee. However, when the time comes to type on the computer there is a glitch in his brain so that while he is in fact sitting in the kitchen he has all the sensory experiences as of writing on a computer. So this person leads a normal life except that instead of typing on his computer, he just sits at the kitchen table while undergoing all the sensory experiences as of typing on a computer. In this case the luck-precluding condition that would link the evidence to the true belief is absent on the type level. This means that the evidence is not a way of arriving at a true belief. For this person the sensory experiences as of typing on a computer are not evidence that leads him to a true belief that he is typing on a computer.

Suppose now that even though the luck-precluding condition is present on the type level, it is absent in a particular instance. We can imagine a person who just in one instance has had a lot of beer, and, while still under the influence of this beer, he sits down at the kitchen table and undergoes all the sensory experiences as of typing on a computer. This person would in general be able to arrive at a true belief on the basis of the sensory experience as of typing on a computer. However, in this particular instance the luck-

precluding condition is absent. This means that in this instance the evidence is not a way of arriving at a true belief.

A Gettier case, as I propose, is such a case where the luck-precluding condition is absent and yet, due to luck, the belief in question turns out to be true. In a Gettier case we have good evidence for a belief. However, the world is such so as to render this evidence misleading, it is not a way of arriving at the truth. Nevertheless the belief turns out to be true, due to luck. In a Gettier case the evidence is not a way of arriving at the truth, and that the belief turns out to be true is a mere matter of luck. We can say more precisely that it is a matter of non-skeptical epistemic luck.

Let me sketch what I would like to do in the following sections. I will first go through a number of Gettier cases in the literature and try to show that they all involve the absence of a luck-precluding condition. This will lend support to the claim that it is non-skeptical epistemic luck that is present in a Gettier situation. I will then give a more general argument to the effect that the type of luck involved in a Gettier case is non-skeptical epistemic luck.

Summary Box 6: Luck-precluding conditions are necessary conditions for perceptual knowledge

The main argument is the following: A Gettier case arises because of non-skeptical epistemic luck. This means that in a Gettier case a luck-precluding condition is absent. If

that is correct and if there is no knowledge in a Gettier case (I call this claim the “Gettier intuition,”) then luck-precluding conditions are necessary conditions for perceptual knowledge. Furthermore, we can explain why there is no knowledge when a luck-precluding condition is absent, thus supporting the Gettier intuition.

3.3.2 Several Gettier cases

In this section I will discuss several of the Gettier case in the literature and show that in each of them the luck-precluding condition is absent. This would mean that in these Gettier cases non-skeptical epistemic luck is present. This in turn supports the claim that the type of luck present in a Gettier case is non-skeptical epistemic luck. In the next section I will give an argument to that effect.

In Gettier’s first case we are to assume that we have strong evidence for a belief about Jones, namely that he, Jones, will get a certain job and that he, Jones, has ten coins in his pocket. Furthermore, we infer from this belief about Jones another belief, namely that the man who will get the job has ten coins in his pocket. Suppose now that the belief about Jones is false, but the inferred belief is true because Smith, who happens to get the job, also has ten coins in his pocket. We want to say that in such a situation we would have a justified true belief but no knowledge that the person who gets the job has ten coins in his pocket.³

³ This case is presented in Edmund Gettier (1963).

For the evidence to lead us to the truth in this instance, it would need to be the case that checking how many coins Jones has in his pocket is a way of arriving at the truth that the person who gets the job, in this case Smith, has ten coins in his pocket. We can imagine a world in which there is such a regularity. Maybe for some jobs it is required of all job candidates to have a specific number of coins in their pockets. Then in counting the coins in the pocket of one job candidate we would be sure to learn how many coins the successful candidate has in his pocket. This would mean that there is a luck-precluding condition on the type level.

However, as the case is designed it is clear that there is no such regularity that links how many coins Jones has in his pocket to how many coins any of the other job candidates has in his pocket. If someone other than Jones is chosen to receive the job then counting the coins in Jones' pocket is not a way of arriving at a true belief about how many coins the successful job candidate has in his pocket. The luck-precluding condition is absent in this case even though the inferred belief is true and there is good evidence for it. However, if Jones were the successful job candidate this would mean that the evidence for the true inferred claim is in fact a good way of arriving at this true belief. So if Jones were indeed the one chosen for the job, then counting the coins in his pocket would be a good way of arriving at the truth about how many coins the person who gets the job has in his pockets. The luck-precluding condition would be present in that case. However, as things are supposed to turn out in this case, the luck-precluding condition is absent, and there is no link between the evidence and the truth of the inferred belief.

In the second Gettier case we are supposed to imagine that we have strong evidence for the belief that Jones owns a Ford. From this justified belief we infer that either Jones owns a Ford or Brown is in Barcelona. We have no evidence whatsoever for believing that Brown is in Barcelona. Suppose now that it is false that Jones owns a Ford, but it is true that Brown is in Barcelona. This means that the inferred disjunction is a justified true belief, but again it surely is not knowledge.⁴

Again it seems that the luck-precluding condition is absent. What renders the disjunction true is that Brown is in Barcelona. However, the evidence for the disjunction consists in evidence to the effect that Jones owns a Ford. There is no luck-precluding connection between the evidence and the truth of the disjunction. We can imagine a case where there would be such a luck-precluding condition. Maybe Jones always house-sits for Brown when and only when the latter is in Barcelona, and Brown allows Jones to drive his car, a Ford, during that time. In that case evidence that Jones owns a Ford (or at least drives a Ford) would indicate that Brown is in Barcelona. However, it is clear that in the case described there is no such connection. Again we can see that the luck-precluding connection would be present if Jones indeed owns a Ford, for then evidence for the belief that Jones owns a Ford is a way of arriving at the truth of the disjunction. However, this luck-precluding condition is missing in this case.

⁴ This case is presented in Edmund Gettier (1963).

So there is some plausibility to the claim that in these Gettier cases we have a situation where we have evidence for a true belief, but the luck-precluding condition is absent. This means that in those two original Gettier cases non-skeptical epistemic luck is present. This is evidence that the type of luck present in Gettier cases is non-skeptical epistemic luck. Let us look at two more examples of such a Gettier case, as discussed in the literature, to further confirm this conclusion.

Suppose that you are looking out on a field and you have the sensory experience as of seeing a sheep at some distance. Suppose now that what you take to be a sheep is indeed a sheepdog. However, further in the distance there is a sheep even though you do not recognize it as such. So you believe that you see a sheep, and this belief is both justified and true, but again it does not seem to be an instance of knowledge.⁵

This is a case of perceptual knowledge, and since I wish to give an account of perceptual knowledge in my dissertation, I need to provide a response to this Gettier case. You have a sensory experience as of a sheep, and this is generally good evidence for the belief that there is a sheep in front of you. That is, we suppose that there is on the type level a luck-precluding condition that is based on the regularity of sheep causing in us a sensory experience as of sheep. However, in this instance the luck-precluding condition is absent on the token level. It is a sheepdog that causes in you the sensory experience as of a sheep. The luck-precluding condition is absent in this instance because you have a

⁵ This case is presented in Roderick Chisholm (1966).

sensory experience as of a sheep but it is not a sheep that causes in you the sensory experience as of a sheep. Now we suppose furthermore that there is a sheep that does not cause in you a sensory experience as of a sheep, maybe because it is too far away or because you are not looking carefully enough at it, or for some other reason. This is the unrecognized sheep in the distance that makes your belief acquired through mistaking a sheepdog for a sheep come out true.

Again there is no luck-precluding condition here. Mistaking a sheepdog for a sheep is not a way of arriving at a true belief that there is an unrecognized sheep in the distance. We might again imagine a world in which there would be such a luck-precluding condition. Maybe in that world sheepdogs are always followed at a safe distance by sheep so that whenever you mistake a sheepdog for a sheep there is indeed a sheep in the distance. Again it is obvious that such a luck-precluding condition is absent in the Gettier case. We can also see that if the sheepdog were not a sheepdog but a sheep, the luck-precluding condition would be present, and the evidence would be linked to the true belief that there is a sheep in front of you.

Let us consider one further case. Suppose that you go to the library and you see Tom Grabit, whom you know well from a class even though you know nothing about his personal life, take a book and conceal it underneath his jacket. When it turns out that the book has been stolen, you speak up and say that you saw Tom Grabit take it. Your belief is justified, and it is also true. However, suppose that Tom also has an identical twin who

was in the library on that day and who is in the habit of taking books from libraries. If this information were made available to you, it would seem that you do not know that Tom rather than his twin took the book.⁶

It seems that in this case the luck-precluding condition is present that links the sensory experience as of Tom Grabit taking the book to the truth, that Tom Grabit took the book. However, there is some defeating evidence available to you. Here we need to remember that the luck-precluding conditions are the way the world needs to be so that given some evidence you will arrive at a true belief. In this case you have evidence that is in fact linked to the truth of the belief for which you have the evidence. So you have good evidence and the world is in fact such that it will lead you to the truth. However, the defeating information makes you doubt that you have such good evidence. The defeating information makes you believe that your evidence is in fact that you have a sensory experience as of either Tom or his twin, where you do not know which, taking the book. And there is no luck-precluding condition that links this evidence to the true belief that Tom took the book. If your evidence consists in fact in a sensory experience as of either Tom or his twin, where you do not know which, taking the book, then this evidence is not linked by a luck-precluding condition to the truth that Tom took the book. We can again imagine a world in which there is such a luck-precluding condition. Maybe whenever the twin takes a book he disguises himself so as to be as dissimilar to Tom as he can be. This would mean that if you have a sensory experience as of Tom or his twin (where you

⁶ This case is presented in Keith Lehrer and Thomas Paxson, Jr. (1969).

cannot say which it is) taking a book, you will arrive at the true belief that it is Tom who takes the book. For if his twin had taken the book, you would have a sensory experience as of someone taking the book who is looking very dissimilar from Tom. Again it is clear that there is no such luck-precluding condition present in the case at hand. So the evidence that given the defeating information you think you have is not linked by a luck-precluding condition to the true belief.

I take it that in these two further Gettier cases discussed in the literature, there is plausibility to the claim that the luck-precluding condition is absent in each of them. Since non-skeptical epistemic luck is present whenever the luck-precluding condition is absent, we get the result that in these Gettier cases non-skeptical epistemic luck is present. It is, of course, not possible to go through all discussed or imaginable Gettier cases here. What we need now is an argument that the type of luck that allows Gettier cases to arise is non-skeptical epistemic luck.

3.3.3 An argument

In this section I will present an argument that in a Gettier case non-skeptical epistemic luck is present. In this argument I will think more about what exactly is the case when non-skeptical epistemic luck is present and whether all such cases lead to a Gettier intuition. Again, by a Gettier intuition I mean the intuition that there is no knowledge in a Gettier case.

Let us suppose that someone has good evidence for a true belief but the luck-precluding condition is absent. Again, the luck-precluding condition can be absent in two ways, and so we need to discuss two cases. Suppose first that the luck-precluding condition is absent on the type level. This means that there is no general regularity that would link the type of evidence in question to the type of true belief in question. This means that the evidence is misleading. Suppose for example that your sensory experiences as of a computer are systematically misleading in that there is no luck-precluding condition that links your sensory experiences to computers. Instead you sit at your kitchen table and just undergo a hallucination as of typing on your computer. We suppose that this is what generally happens, that there is no general regularity that would link your evidence (experiences as of typing on a computer) to a true belief (you are typing on a computer).

We have thus described a situation in which the luck-precluding condition is absent and in which therefore non-skeptical epistemic luck is present. In order to arrive at a Gettier case we merely need to render the belief for which we have the evidence true in one instance. So suppose that in one instance you, as always in the morning, are sitting at your kitchen table undergoing all the sensory experiences as of typing on a computer. Generally your belief that you are typing on a computer is false. However, today your neighbor has left his computer on the tabletop and your fingers, usually just “typing” on the bare table, are really pressing keys on a computer keyboard. You would have good evidence for the true belief that on this day you are typing on a computer.

We thus have now taken a case in which non-skeptical epistemic luck is present and transformed it into a Gettier case just by rendering the belief in question true in one instance. As it seems to me, we have a Gettier intuition about this case. We have the intuition that in this case you do not know that you are typing on your computer. If we have to say why we think that there is no knowledge in this case, it seems to me that the only thing to point to would be the absence of the luck-precluding conditions. We want to say that this evidence is because of the absence of the luck-precluding condition not a way of arriving at the truth. We want to say that in this instance it is merely a matter of luck that the belief is true. It is not that the evidence can lead us to the truth. In fact, the evidence is not a way of arriving at the truth in the circumstances imagined. If the luck-precluding condition is absent on the type level then the evidence and the truth that it indicates are mismatched. This is, as it seems, what explains the Gettier intuition about this case.

I take it thus that if the luck-precluding condition is absent we can transform this case into a Gettier case, and we would have a Gettier intuition about this case. Furthermore, this Gettier intuition can plausibly be explained by the absence of the luck-precluding condition. Because the luck-precluding condition is absent, there is no knowledge in a Gettier case. This again suggests that it is non-skeptical epistemic luck that leads to Gettier cases.

We have now considered the case where the luck-precluding condition is absent on the type level. Let us now consider the case where the luck-precluding condition is present on the type-level but absent on the token level. We suppose that in general the evidence is linked to the truth, but the regularity in the world on which this link depends is missing in this one instance. Again suppose that generally when you have the sensory experiences as of typing on a computer, you are indeed typing on a computer. However, in this one instance you have a vivid hallucination in which you, while in fact sitting at your kitchen table, undergo all the sensory experiences as of typing on a computer. The regularity that connects the evidence (sensory experiences as of typing on a computer) to the truth (you are typing on a computer) is present in general, but absent in this instance. We have thus described a case in which the luck-precluding condition is absent.

Let us now transform this case into a Gettier case. To do so we just need to render this belief for which we have good evidence also true. Suppose that, unbeknownst to you, a very advanced scientist has embedded in your kitchen table a keyboard that is linked to a computer. So suppose that while you are “typing” on the kitchen table, touching what you would assume to be the bare wood of the kitchen table, some device inside of your kitchen table picks up the pattern of typing and sends it to a computer. In that case it would be true that you are typing on a computer even though you are hallucinating that you are typing on a computer. In this strange case again we would have good evidence for a true belief. Thus we have transformed the case in which the luck-precluding condition is absent into a Gettier case.

Furthermore, I take it that we have a Gettier intuition about this case. If we have no idea that there might be this embedded keyboard and are just touching what we would, if we were not hallucinating, take to be the bare wood of the kitchen table, then we do not know that we are typing on a computer. Hallucinating and tapping out a pattern on a kitchen table is usually not linked to the truth that one is typing on a computer. The world is not such so as to render this kind of hallucination a good way of arriving at the truth that one is typing on a computer. Again, it seems that the absence of the luck-precluding condition would be the natural explanation for the Gettier intuition. It is because the world is not such so as to render this kind of hallucination a way of arriving at the truth, that there is no knowledge in this case even if the belief happens to be true.

The luck-precluding condition links evidence and truth. However, in this case what renders it true that you are typing on a keyboard is so remote from your evidence, which is just a sensory experience as of typing on your normal computer, that this evidence cannot give us knowledge. Again this suggests that in a Gettier case non-skeptical epistemic luck is present and that it is this non-skeptical epistemic luck that creates the Gettier intuition.

The argument that I have presented in this section can be briefly stated as follows: If we describe a case in which non-skeptical epistemic luck is present, then we can transform it into a Gettier case merely by rendering the belief in question in one instance true. The

Gettier intuition about the resulting case can plausibly be attributed to the absence of a luck-precluding condition. This suggests that the type of luck that leads to a Gettier case is non-skeptical epistemic luck.

3.4. Given the absence of non-skeptical epistemic luck no Gettier case can be constructed

Weirdly enough, in the preceding section I have used Gettier cases to argue that luck-precluding conditions are necessary conditions for perceptual knowledge. However, Gettier cases are usually perceived as the challenge to find sufficient conditions for knowledge. In this section I will discuss the conditions given which a Gettier counterexample to a proposed account of perceptual knowledge is not possible.

The main idea is that if non-skeptical epistemic luck is absent and certain further conditions are present, then we cannot create a Gettier intuition about the resulting case. So the idea is that luck-precluding conditions plus certain further conditions are not subject to a Gettier counterexample that would show them to be not sufficient for perceptual knowledge. In the fifth chapter I will try to give an argument that the necessary conditions developed in this dissertation for perceptual knowledge are jointly sufficient. The project of this section is more limited. All I want to do is to think about whether we can construct a Gettier case if non-skeptical epistemic luck is absent and certain further conditions obtain. My claim is that this is not possible. This strengthens also the claim that Gettier cases depend on non-skeptical epistemic luck.

3.4.1 Expected luck-precluding conditions

In this section I will think about which conditions might be such that when they are present no Gettier case can be constructed. This is a step in search of sufficient conditions for perceptual knowledge. I have tried to show that the luck-precluding conditions are necessary conditions for perceptual knowledge. I have also assumed without argument (which will, however, be supplied in the fifth chapter) that a true belief and evidence for it are also necessary for perceptual knowledge. Let us now think about Gettier cases and see which conditions would, if present, allow us to avoid Gettier cases.

In one of the preceding sections several Gettier cases from the literature were discussed. In this preceding section I was merely interested in arguing that in each of them non-skeptical epistemic luck is present. Let us now revisit two of these Gettier cases and discuss what would have to be added to these cases so that we would no longer have a Gettier intuition about it. So the idea is to think about these Gettier cases and to try to see which conditions have to be added to these Gettier cases to transform them into instances of knowledge, to de-Gettierize them, so to speak.

In the sheepdog case we said that there is no luck-precluding condition that links evidence derived from mistaking a sheepdog for a sheep to a true belief about an unrecognized sheep in the distance. However, now suppose that such a luck-precluding condition is present. So we suppose that the world is in fact such that sheepdogs are

followed at a safe distance by sheep, so that when you mistake a sheepdog for a sheep you indeed see an unrecognized sheep in the distance. We would now have evidence for a true belief and the luck-precluding condition obtains. Would someone in such a situation now know that there is a sheep in the field?

I take it that we would still have a Gettier intuition about this case. That is, we would still think that there is no knowledge in this case even though there is this weird regularity of sheepdogs and sheep that follow them at a safe distance. The mere obtaining of this weird regularity is not sufficient in this case to give us knowledge.

Let us confirm this intuition by thinking about one more of the Gettier cases. In the second of the original Gettier cases the luck-precluding condition that links evidence that Jones drives a Ford to the truth that Brown is in Barcelona is absent. However, suppose now that the luck-precluding condition is present. Suppose that the world is in fact such that Brown lets Jones house-sit and drive his Ford whenever and only when he is in Barcelona. In that case we would have evidence for a true belief and there would be a luck-precluding condition. But again we surely have the intuition that this case is not an instance of knowledge.

Let us reflect a little bit on what the picture of knowledge is that we are painting here. So we have evidence for a true belief. The evidence is, however, compatible with the falsity of the belief for which we have the evidence. So we need the luck-precluding condition to

connect the evidence to the true belief. This is done by physical regularities in the world that render it such that de facto the evidence will lead to a true belief. What we have thought about so far is what this luck-precluding condition exactly is and what happens when it is absent. Now we need to think about an internal condition that needs to complement the external requirement of a luck-precluding condition.

The main idea is that sufficient conditions for perceptual knowledge require both internal and external conditions. Furthermore, these internal and external conditions need to fit together, they need to be matched with one another. The mere obtaining of a weird luck-precluding condition is not given evidence for a true belief sufficient for perceptual knowledge. This is because this luck-precluding condition might be somehow beyond the cognitive horizon of the epistemic agent. So the main idea is that the epistemic agent needs to expect or assume that the luck-precluding condition obtains. This is what we have to explain in some detail now.

The epistemic agent needs to expect or assume that a specific luck-precluding condition obtains. This expectation has to be reasonable in a sense to be explained in the fourth chapter. To see how this might lead to sufficient conditions for perceptual knowledge let us add this further condition to the two Gettier cases just discussed. As the case is described, the evidence for the belief that there is a sheep in front of you consists in a sensory experience as of a sheep that is in fact caused by a sheepdog. Given this evidence we have no expectation whatsoever that there might be a luck-precluding condition that

links sheepdogs to unrecognized sheep in the distance. The epistemic agent does not expect or assume that this luck-precluding condition obtains. In the second Gettier case just described, a variety of evidence is available for the belief that Jones owns a Ford. You have seen him drive a Ford on various occasions; you have heard him refer to “his Ford” on various occasions; maybe you have seen some paperwork that seems to confirm that he owns a Ford, and so on. Again it is clear that this evidence cannot lead to an expectation that the luck-precluding condition linking evidence for Jones’ driving a Ford to Brown’s being in Barcelona is present. Even if this regularity in fact characterizes our world, the evidence would not lead to the expectation that it is present.

The idea is thus that the mere obtaining of the external luck-precluding condition if it is not also internally expected or assumed to obtain is not sufficient for knowledge. There needs to be an internal element that involves somehow an expectation or assumption that the luck-precluding condition obtains. Lots of weird luck-precluding conditions probably characterize our world. If we have no idea that they obtain, then it seems that they by themselves cannot give us knowledge.

We need to explain now in some detail what is meant by an expectation or assumption that the luck-precluding condition obtains. Let us consider again the barn façade case. So you are driving around and you arrive at the belief that there is a barn in the field on the basis of a sensory experience as of a part of a barn, a façade. When you adopt on the basis of such evidence the belief that there is a barn in the field, you assume that what

you see is part of a structure rather than a mere façade. You assume that there is a real building rather than a movie set or a fragment of a building that has fallen into disrepair so that only the façade is left. Similarly, when you arrive at the belief that there is a tree in the field, you assume that it is a real plant rather than a sculpture that has the shape and size of a tree.

I take it that this is a perfectly ordinary phenomenon. We always see only parts of buildings and assume that all the other unseen parts that make it a building are also present. We do not feel the need to test against the hypothesis that the tree might be a sculpture before we adopt the belief that there is a tree in the field. Thus there is a familiar and ordinary sense in which we assume that the luck-precluding conditions obtain.

Now, we do not ordinarily have occurring beliefs about these luck-precluding conditions. It is not that when driving around we reason that this part of the barn visible from the road is a part of a real building rather than a movie set. The expectation or assumption of the luck-precluding condition does not take the form of an occurring belief or a premise that we need to keep in mind to arrive at beliefs about the physical world. We do not usually have occurring beliefs of that sort at all. But then what is the nature of this assumption or expectation that the luck-precluding condition obtains?

It seems to me that we expect a specific luck-precluding condition to obtain if we would regard any evidence that it does not obtain as a defeater for the belief in question.

Consider again the situation in which you are driving around and arrive at the belief that there is a barn in the field. Suppose that after a few miles there is a sign by the road saying that for the purposes of shooting a movie a lot of mere facades of buildings can be found in this area. Maybe the sign mentions the specifics of the movie and advertises the variety of barn facades as a tourist attraction. If you consider this sign a defeater for the belief that there is a barn in the field, then you expected or assumed the luck-precluding condition that the facades are parts of real building.

The idea is thus the following. Consider your evidence for a particular belief. You expect there to be a certain luck-precluding condition if and only if evidence that the luck-precluding condition does not obtain would be regarded by you as a defeater for that particular belief. In the example, you have a sensory experience as of a part of a barn. You believe that there is a real barn in the field. You expect or assume that the façade, which is all that is visible from the road, is part of this building in the sense that evidence that it is not part of a real building is a defeater for the belief that there is a real barn.

Maybe this expectation or assumption is a dispositional belief – I am not sure what kind of a thing exactly a dispositional belief is. However, there is a sense in which this expectation or assumption is present in the epistemic agent and necessary for the belief that, for example, there is a barn in the field. This sense depends on what the epistemic

agent would consider a defeater for the belief that there is a barn in the field. If evidence that a specific luck-precluding condition does not obtain would be a defeater for the belief that there is a barn in the field, then the epistemic agent expects or assumes that the luck-precluding condition obtains. And this expectation or assumption is something that the epistemic agent would regard as necessary for adopting the belief that there is a barn in the field.

We can also see that this expectation or assumption allows us to at all move from our limited evidence to the beliefs that go beyond this evidence. Again, the situation is one where we have little bits of evidence, glimpses, snatches, casual touches, and so on. On the basis of this limited evidence we adopt beliefs that outstrip this evidence. We believe that there is a barn, house, tree, road, rather than a façade, something of a tree-like shape, and so on. In order to move from this evidence to these beliefs we need to assume or expect that the luck-precluding conditions obtain. In most cases we will never think about the nature of these luck-precluding conditions. We do not reason from the evidence to the belief. There is no inference involved here. However, we do assume or expect that this evidence is not misleading. And this expectation can be brought to consciousness when we think about what we would consider a defeater for the belief in question.

Summary Box 7: Expected luck-precluding conditions

The mere obtaining of a luck-precluding condition is clearly not sufficient for avoiding Gettier cases. The luck-precluding condition as an external requirement must be expected

or assumed to obtain, where this is an internal requirement. This expectation or assumption is not an occurring belief. It is however present and necessary for the belief in question in the following sense: An epistemic agent assumes that a specific luck-precluding condition for a particular evidence/belief pair obtains if and only if any evidence that the luck-precluding condition does not obtain would be a defeater for the belief adopted on the basis of this evidence.

3.4.2 Gettier cases revisited

In this section I will discuss again Gettier cases with the goal of showing that if the luck-precluding condition is expected or assumed to obtain in the sense defined above and if it indeed obtains, then the Gettier case is no longer a Gettier case; the Gettier intuition would no longer be present for that case. The idea is thus that if we have evidence for a true belief and the luck-precluding condition obtains and is expected to obtain, then this de-Gettierizes the case in question and we would have knowledge in such a case.

Let us begin by thinking about the two Gettier cases discussed again in the preceding section. Suppose that my evidence consists in a sensory experience as of a sheep.

Suppose furthermore that I assume that even if I mistake a sheepdog for a sheep, there is an unrecognizable sheep in the distance. So we suppose that I have a sensory experience as of a sheep, and I recognize the possibility that it might be a sheepdog. However, I am also aware of the regularity between sheepdogs and sheep in the distance. I expect this regularity to obtain in the sense that if I were to learn of evidence that it does not obtain I

would no longer believe that there is this sheep in the field. This is now a really weird case, but it seems to me that we would no longer have a Gettier intuition about it.

The assumption or expectation of the luck-precluding condition is present, as we suppose, in the sense that if I were to learn that no sheep ever follows a sheepdog at any distance, maybe because the sheep are scared of the sheepdogs, I would no longer believe that there is a sheep in the field on the basis of this sensory experience. Suppose that in this sense the luck-precluding condition is expected to obtain and it does indeed obtain. Then it seems to me that by mistaking a sheepdog for a sheep I can still know that there is a sheep in the field.

Let us look at the other Gettier case discussed in the preceding section. Suppose my evidence is that I have seen Jones drive a Ford, refer to “his Ford,” and that I also have seen some documentation supporting that Jones is the owner of a Ford. Furthermore, suppose that I recognize that there is a regularity that connects Jones’ driving a Ford to Brown’s being in Barcelona. So we suppose that the luck-precluding condition obtains and is expected to obtain. The expectation is present in the sense that if I were to learn that Brown never lets anyone drive his Ford, I would no longer believe the disjunction. So the expectation of the luck-precluding condition is present in the sense that the belief in the disjunction can be defeated by evidence that speaks against the luck-precluding condition.

We thus have a situation in which I expect the luck-precluding condition to obtain and it does obtain. I assume that there is a regularity between Jones' driving a Ford and Brown's being in Barcelona. Given this assumption and the fact that the regularity indeed obtains, it seems that I can know on the basis of the evidence the disjunction. I take it that the Gettier intuition is no longer present in this case. It is, however, obvious, that these are now really weird cases, and our intuitions about them might not be very clear. We need to discuss some simpler examples.

Let us look at a case about which we hopefully have clearer intuitions than about these really weird de-Gettierized Gettier cases. To begin, let us describe a Gettier case.

Consider again the case described at the beginning of this chapter, involving an elephant-caused accident on I-10. Lucky truly believes that an escaped circus elephant caused a major accident on I-10. She has good evidence for the belief: a report in a very reliable and trustworthy newspaper. However, the luck-precluding condition is absent because this particular report was merely called in as a hoax by Sokal. The world is not such so that a fictitious report by Sokal will lead us to a true belief except by pure chance, as in this case. This, then, is clearly a Gettier case. Lucky has a true belief and good evidence for it, but she lacks knowledge because of the absence of a luck-precluding condition.

Now suppose that there is such a luck-precluding condition. Suppose that Sokal suffers from a weird case of partial memory loss. In every day there are two hours from which he can later on remember nothing, nothing about what he did or learned or what happened to

him. Out of each day, two hours are missing from his memory. However, Sokal is unaware of this partial memory loss, and moreover, although if asked what happened during those two hours he has no answer, when we comes up with fictitious stories, somehow part of the memory is restored. However, he regards these fictitious stories as mere fiction, not as memory. So somehow he is able to tap into the information about these two lost hours when coming up with fictitious stories, however without realizing that he is drawing in some mysterious way on memory rather than imagination. Suppose that all fictitious stories Sokal ever comes up with draw on lost memories from these two hours. We suppose that Sokal has otherwise normal epistemic abilities, and so his “fictitious” stories are not fiction, but memories although not recognizable as such by himself.

We thus suppose that when Sokal calls in his fictitious report about the elephant-caused accident, he is speaking from memory although he does not realize that. We thus suppose that even Sokal’s fictions are good evidence for true beliefs. Suppose then that this luck-precluding condition obtains. Lucky now has a true belief for which she has good evidence and there is a luck-precluding condition. I take it that she still has no knowledge. This is because she does not expect or assume that the luck-precluding condition obtains. If she were to learn that Sokal’s memory loss has progressed and that he is no longer able to recall these two hours in fictitious stories, she would not regard this as a defeater for the belief in the elephant-caused accident on I-10. The belief that this accident occurred does not in any way depend on an assumption about the nature of

Sokal's memory loss. In that sense the luck-precluding condition is not expected, and because of that even if it does obtain there is no knowledge in this case.

However, suppose now that Lucky has this expectation or assumption of the luck-precluding condition. This means that she believes that the elephant-caused accident occurred in part because she recognizes that Sokal's fictions are based on memories. She would respond to a defeater to this luck-precluding condition by withdrawing the belief in the elephant-caused accident. She now has a true belief: an escaped circus elephant caused a major accident on I-10. She has the evidence of the newspaper report to support this true belief. If confronted with evidence that the nature of Sokal's memory loss has changed, she would no longer believe that this accident occurred. Furthermore, this report that Sokal regards as fiction is instead based on a reliable although inaccessible memory. I take it that we would now no longer regard this as a Gettier case, one in which Lucky lacks knowledge. Admittedly I am not sure whether this is a simpler case about which we have a clearer intuition.

Maybe the simplest case would be one where an epistemic agent has been visiting a movie set in which there are only facades and no real building. Now he emerges from the studio grounds. When walking around among the facades on the movie sets, we assume that he has a variety of sensory experiences as of parts of buildings. When emerging from the movie set he still has very similar sensory experiences as of parts of buildings. However, while on the movie set, he does not assume that these facades are parts of real

structures. When emerging from the movie set grounds, he now does assume that the facades he sees are parts of real buildings. In the movie sets the epistemic agent does not expect there to be the luck-precluding condition. When emerging from the movie set he now does expect and assume that the facades are parts of buildings rather than mere facades. This is the assumption of the luck-precluding condition that I claim to be necessary for perceptual knowledge.

The claim is furthermore that if we have good evidence for a true belief, the luck-precluding condition obtains and is expected to obtain, then no Gettier case can be constructed. If all of these conditions obtain, we simply do not have a Gettier intuition about this case. This is what we have to argue for now.

3.4.3 Can there be a Gettier counterexample to the proposed account of perceptual knowledge?

In this section I will consider the question of whether there can be a Gettier counterexample to the proposed account of perceptual knowledge. My claim is that no Gettier case can be constructed for a case in which we have good evidence for a true belief and the luck-precluding condition obtains and is expected to obtain. In the fifth chapter an argument will be given for the claim that the various necessary conditions are jointly sufficient for perceptual knowledge. The discussion in this section is more limited, involving a response to a proposal of such a Gettier counterexample.

In the preceding section I take Gettier cases and try to supply those conditions that would de-Gettierize them. I propose that if the luck-precluding condition obtains and is expected to obtain, then we no longer have a Gettier intuition about such a case. Now what we need to see is that if these conditions all obtain, we have evidence for a true belief and the luck-precluding condition obtains and is expected to obtain, then we cannot somehow tweak this case so that we would have a Gettier intuition about it. Put as a question, we might ask whether it is possible to construct a case in which these four conditions are met and still render it a Gettier case. My claim is that this is not possible.

In this section I will argue for this claim by looking at a proposal of such a case, as presented by John Pollock. Suppose that an epistemic agent enters a somewhat strange room in which some strange apparatus interferes with normal vision so that in the normal sense the epistemic agent does not see any of the objects in the room. Suppose furthermore that this epistemic agent has no idea that this room that he is entering is different from any other room that he has entered. However, by some unusual way this apparatus causes in the epistemic agent a sensory experience as of a red cube if and only if the only object in the room is a red cylinder. The question is whether this epistemic agent knows that the object is red, or whether this is a Gettier case.

We can see that the epistemic agent truly believes that there is a red object in the room and has a sensory experience as of a red object as evidence for this belief. Furthermore, there is a causal connection between this red object and the sensory experience as of the

red object, although a somewhat unusual one. So it seems that the luck-precluding condition obtains. Furthermore, it seems that the epistemic agent expects such a causal connection to obtain, although the details would surely be somewhat surprising. So it seems that given the account of perceptual knowledge proposed here, I would have to say that the epistemic agent knows that there is a red object in the room.

It seems to me that there are two elements in this case that might give rise to a problem for saying that there is knowledge in this case. First, the causal connection between the red object and the sensory experience as of a red object is very unusual and its details would be extremely surprising to the epistemic agent. The epistemic agent is not perceiving in the normal and natural way; instead his sensory experience is caused by a strange apparatus. Second, there is a distortion involved in the sensory experience created by this apparatus, in that the epistemic agent sees the object as a cube when it is a cylinder. We now need to investigate whether these two points in isolation or together would require us to say that there is no knowledge in this case.

Let us consider first the second point in isolation. So let us examine whether a partial distortion would undermine our knowledge of the non-distorted elements. To do so, let us take a case of basic sensory knowledge. We begin with a true belief. In this case let us suppose that it is true that a specific person is typing on a computer. Next we need evidence. So we suppose that the person has normal sensory experiences as of typing on a computer. Furthermore, we need a luck-precluding condition. So we suppose that there is

both in general as well as in this particular instance a regularity of a computer causing in this person the sensory experiences as of typing on a computer. Let us now apply some distortion to this case and see whether this would give us a Gettier intuition about it.

Suppose that you have sensory experiences as of typing on a sleek new bright red laptop with a really big screen and the latest software when in fact you are typing on an old bright green desktop computer with a tiny monitor and 10 year old software. Would not such evidence undermine knowledge that you are typing on a computer? Even more radically we might imagine that you have sensory experiences as of typing on a normal computer when in fact you are sitting at your kitchen table typing on the plain wood surface in which a computer keyboard is embedded that sends your key strokes to a computer.

Suppose that in these cases we have evidence for a true belief and there is an expected luck-precluding condition that links the evidence to the truth. You truly believe that you are typing on a computer. You have sensory evidence for this belief. Furthermore, there is a luck-precluding condition that links your sensory experiences as of typing on a computer to the truth, namely that you are indeed typing on a computer. However, your sensory experience is in some, more or less radical way, misleading, though not in so far as it concerns your typing on a computer. Would we want to say that the person in question knows that he is typing on a computer?

It might seem problematic to say that the person knows that he is typing on a computer because your evidence leads to just this one true belief and to many false beliefs. You would be right about typing on a computer, but you would be more or less radically wrong about what kind of a computer it is, what its color is, whether its keyboard is embedded in a kitchen table, and so on. So it might seem that it is just lucky that your evidence leads to this one true belief amidst all these false beliefs that your evidence also leads to.

However, I am inclined to think that one can have such “isolated” knowledge, knowledge that arises in the midst of a large number of false beliefs. Suppose for example that you witness a car accident, and it all happens really fast, so that you might have no beliefs or false beliefs about many details of what you witnessed. Nevertheless it seems that you can know some things about this accident, how many cars were involved, which car crashed into which other car, and so on. You might be hazy or wrong about whether the cars tried to avoid the accident by breaking or steering away, or which car honked the horn, and so on. But it seems clear that in the midst of ignorance and false beliefs you can still know some things, those that are particularly obvious to you, those to which you pay particular attention. We can render this case more parallel to the computer case by thinking about the reports such witnesses might give on the basis of memory. The same memory would be evidence for some false and some true beliefs.

Let us see in what way the computer case might be parallel to the car accident case. We have a sensory experience as of the car accident, but it leads to some false beliefs. The sensory experience and then the memory of it lead to a number of false beliefs. But it also leads to some true beliefs, and these true beliefs can be knowledge despite that they arise amidst of a number of false beliefs. If something similar is true in the case of the computer we would have a way of answering a part of the objection just given. You have a sensory experience as of typing on a computer, and some of your beliefs about the computer arrived at on the basis of sensory evidence are false. But some are true, and if all the other conditions are met, it seems that these are instances of knowledge.

There is some temptation to say that if some evidence leads to both true and false beliefs, then its leading to the true beliefs is just too lucky to be knowledge. This is a version of the problem of the gap. When I talk about the problem of the gap I usually consider cases where we have evidence that might hypothetically lead to a false belief. So the problem of the gap arises because one and the same bit of evidence can in some possible scenarios lead to a true, but in others to a false belief. We have now another version of the problem of the gap, one where some evidence leads in one and the same scenario to some true and some false beliefs. Now, I accept the claim that our evidence cannot by itself give us knowledge. I accept that there is such a gap, and I try to argue that it is compatible with knowledge. However, accepting the gap means accepting that some evidence might lead to a false belief, or even if it leads to a true belief, it might also lead to false beliefs. So there is an element of luck involved in knowledge that is just unavoidable. However, the

picture of knowledge I paint is realistic. Cases like those of witnessing a car accident are cases where we generally grant a witness some knowledge even though we know that he might be mistaken in some of his beliefs.

What we have said so far is only a partial answer to the objection because we have only talked about cases of misperception. A misperception is a case where a sensory experience that leads in this case or in very similar cases to true beliefs gives rise in this instance to a false belief. However, some of the cases described in the objection are more radical than misperceptions; they involve deceptions. A deception case is one of hallucinations rather than a mistaken perception. A hallucination is not in some way caused by what the hallucination is about. In the case of a misperception of a cylinder as a cube, we still have a causal relationship between the object in question and the sensory experience as of the object, although it involves some distortion. In the case of a hallucination there is no causal relationship between what the hallucination is about and the hallucination itself. Let us now consider a case where there is not misperception but hallucination going on.

It seems to me that in the case of a very radical deception or hallucination we would hesitate to say that we have the kind of justification for the true belief we would need for knowledge. If it is not a case of partial perception or partial misperception that leads to false beliefs, but somehow a vivid hallucination that leads to mostly false beliefs one of which happens to be true, we would ordinarily hesitate to say that we have the kind of

justification for this belief needed for knowledge. It seems that this is what is going on in the embedded keyboard case. This would be a case where someone who has the sensory experiences as of typing on a normal computer, is in fact sitting at a kitchen table in which a keyboard is embedded that sends his typing to a computer – of all of which he has no inkling. We would probably not want to say of such a case that the person has the type of justification needed for knowledge. If such a radical deception is going on, we would hesitate to say that the person has the kind of justification needed for knowledge. However, this is something we will have to look into more carefully in the next chapter. The question tackled in the next chapter is whether someone who is radically deceived can still have the kind of justification required for perceptual knowledge.

Let us now consider the first element in Pollock's case discussed above, that the details of the luck-precluding condition would be unexpected and surprising. Consider the case where someone undergoes the sensory experiences as of typing on a computer when he is in fact sitting at the kitchen table and tapping out a pattern, which is picked up by the embedded keyboard. What makes it true that the person is typing on a computer is very different from what he would expect on the basis of the evidence. A sensory experience as of typing on a normal computer leads to expectations that are only partially realized by the computer on which one types given that a keyboard is embedded in a kitchen table. The situation is similar to expecting a wave that is good for surfing and being caught in a Tsunami wave. Surely, the expectation is satisfied at some level, but one might also say that what realizes the expectation is so far removed from what one expected that it would

be misleading to say that things were happening as expected. In a deception case we should thus say that the luck-precluding condition is not really expected because its details are so surprising and unexpected given the evidence. However, what about cases of misperception that involve unexpected and unusual details of the luck-precluding condition?

The problem is that for most of us today and all of us historically the details of the luck-precluding conditions are surprising and unexpected even for normal vision or any other normal sensory experience. Most people do not know how vision works, or how any other sensory perception works. The details might well be very surprising to us. It seems, however, that this is irrelevant to whether we know that there is a barn in the field. I do not need to know how perception works in order to know that there is a barn in the field. Even if the details are very surprising, even if I have mistaken views about what these details are, it seems to me that the knowledge that there is a barn in the field cannot and should not depend on having a good theory of how vision works. Furthermore, whether vision works in the usual and natural way or in some technical and artificial way does not seem to undercut our knowledge gained on the basis of such perception. No matter how my senses work, no matter whether they work in the natural or an artificial way, it still seems to me that I can gain knowledge on the basis of such sensory experiences.

Let us now return to Pollock's case. I have considered two objections to the claim that the epistemic agent in this case has knowledge that there is a red object in the room. The first

objection is that the details of the luck-precluding condition are surprising and that they involve an artificial apparatus. As it seems to me, for most of us the details of natural perception are surprising also, and there is no reason to say that we can gain perceptual knowledge only through natural sense perception but not with some artificial apparatus. The second objection is that there is a distortion involved. Again, the problem is that in most of our perceptions, for example when witnessing a car accident, there is some distortion involved, in that we will also end up with a variety of false beliefs and that any knowledge that we gain would arise amidst such false beliefs. However, it does not seem to me that this distortion by itself undermines our knowledge of the non-distorted elements. We rely on witnesses to car accidents even though we know that they might be wrong about many things.

As should be clear, John Pollock has given the most forceful and pressing objection to the account of perceptual knowledge developed in my dissertation. I am not sure whether what I have just said answers this objection. It might be that some would have the intuition that these two points are not answered by the discussion I have just presented. However, if there is a Gettier counterexample to the account developed here, I am sure it is one along the lines of the case presented by Pollock.

Another point that needs to be mentioned is that a different response is required for cases of misperception (which I take is involved in Pollock's case) and a case of radical deception or hallucination. For a case of misperception I am prepared to claim that there

is knowledge in these cases. If someone has sensory evidence for the true belief that the red car crashed into the side of the blue car, and the luck-precluding condition obtains and is expected to obtain, then we can know this fact about the accident. And it does not matter that maybe the witness has a false belief as to which car honked the horn, or about whether the light was still yellow or whether it had just turned red. To demand that a sensory experience or a memory is exact in all respects would be to demand too much. However, in the case of radical deception or hallucination I would not be prepared to say that there is knowledge even if the luck-precluding condition obtains. In such a case I would try to argue that the epistemic agent does not have the kind of justification for his true belief that is required for perceptual knowledge. This is something I will develop in detail in the fourth chapter of my dissertation.

One might ask whether the distinction between misperception and radical deception is vague in that we can create cases of more and more radical misperception that then at some point turn into cases of deception. There thus might be cases where we are not sure whether they are instances of misperception or deception. Let us consider such a case.

Suppose that you are typing on a computer and you have the sensory experiences as of typing on a computer. Your sensory experiences as of the typing appearing on the screen are not misleading in any way. However, you have sensory experiences as of typing on an ordinary keyboard when in fact you are just tapping a table in which a keyboard is embedded. Furthermore, the signals of your striking keys are not sent to the box you take

to be your desktop, but instead to a computer on the other side of the globe that then sends it back to your actual screen. I take it that this is a borderline case where we are not sure whether we want to say that the person is misperceiving things or whether he is deceived in a radical way.

Now, if you are just misperceiving some things you would have knowledge that you are typing on a computer. However, if you are deceived to the extent that your belief is not justified, in that case you would not have knowledge that you are typing on a computer. So we have a borderline case of where we are not sure whether to say that the person in question knows that he is typing on a computer. However, this borderline case is benign, because it is a case where we are not sure whether all the four conditions for knowledge are really met. If they were met, if we somehow could be sure that the person has a justified true belief plus the appropriate luck-precluding conditions, we would have a case of knowledge. So this weird case is a borderline case for whether or not all of the four conditions are met. It is not a case in which we are sure that all the four conditions are met, but we have a Gettier intuition about it.

The question I have examined in this section is whether there can be a Gettier counterexample to the four conditions proposed for perceptual knowledge. The strongest case I am aware of is the one presented by John Pollock. I claim that such a case of misperception might still be argued to allow knowledge of the non-distorted elements. I am aware that this idea that we can have knowledge amidst false beliefs might be

disputed. However, it does seem to me that even those who think that there can be no such knowledge would not have as strong a Gettier intuition about this case as one usually has for the original cases. So if this is, as I believe, the strongest possible Gettier proposal for the account of perceptual knowledge developed in my dissertation, at least I could say that it does not lead to as clear and decisive a Gettier intuition as do the original and many other Gettier cases discussed in the literature.

Summary Box 8: A Gettier counterexample?

We can imagine a case of partial misperception that involves some distortion but also some true beliefs. Suppose that for the true beliefs all the conditions for perceptual knowledge proposed in the dissertation obtain. Would they constitute knowledge? I claim that we often rely on knowledge that arises amidst false beliefs, for example in an account by a witness to a car accident. If the misperception is more radical, rather a hallucination or radical deception, I would dispute that the epistemic agent in question has the kind of justification required for perceptual knowledge. It, thus, seems to me to be an open question whether such a case would generate a Gettier intuition, and I suspect that such a Gettier intuition would be less strong than the one generated by the traditional Gettier cases as discussed in the literature.

3.5 Comparison with other views

In the introduction to this chapter I mention the work of Linda Zagzebski, Peter Klein, and Keith Lehrer, which explores similar positions to the one presented in this chapter. Let us now see how their work gets developed further in this chapter.

In her paper “The inescapability of Gettier problems,” Zagzebski (1994) argues that if the justification requirement for knowledge does not entail that the truth requirement is also satisfied, then a Gettier case can be constructed. Furthermore, since it seems implausible to render the justification requirement so strong that it entails truth, we, according to her, need to accept that knowledge contains some luck. It thus seems that Zagzebski considers the kind of epistemic luck that gives rise to Gettier cases as unavoidable.

I agree with Zagzebski’s that Gettier cases arise because of epistemic luck. However, since we have the intuition that there is no knowledge in a Gettier case we need to find a way to define knowledge so as to exclude this kind of epistemic luck that leads to a Gettier case. I also agree with Zagzebski that the Gettier case type of epistemic luck arises because there is what she calls “independence” between the justification and the truth requirement for knowledge. This means that one might have a justified but false belief. However, rendering the justification requirement so strong that it entails truth is only one way of tying together the justification and truth requirement. This chapter explores a different way of dealing with this independence problem. Let us explain this.

The proposal in this chapter has been to assume that our evidence is compatible with the falsity of the belief for which we have evidence. Because of this gap, we add another condition, an external requirement, namely the luck-precluding condition, which makes sure that the world is such that the evidence will lead us to the truth. There is no entailment here between evidence and truth. However, the evidence plus the luck-precluding condition will lead us to a true belief. As I have tried to show, if we combine the evidence requirement with a luck-precluding condition that both obtains and is expected to obtain, then the resulting situation does not allow us to form a Gettier intuition about it.

I agree with many of the elements of Zagzebski's analysis. She says that what gives rise to the Gettier problem is that there is no guarantee that the evidence will lead us to a true belief. The luck-precluding condition is designed to provide such a guarantee. Furthermore, she says that we should not render the justification requirement so demanding that skeptical worries will arise even if the so understood justification will then guarantee the true belief. In the account developed in this chapter the evidence requirement is not rendered so demanding as to guarantee by itself the true belief for which we have the justification. Furthermore, she says that there is an unavoidable element of luck involved in knowledge. This can refer to skeptical epistemic luck, which I regard as unavoidable. However, as I try to show, it is not the kind of epistemic luck that gives rise to Gettier cases.

Peter Klein (1971) in his paper “A Definition of Propositional Knowledge” develops an analysis of knowledge to which, as he proposes, there can be no Gettier example. There are several similarities between his analysis and response to the Gettier problem and the one proposed in this chapter. He begins by explaining that knowledge is incompatible with a felicitous coincidence, which he defines as a situation in which there are facts about this situation that defeat the evidence that the person in this situation has for her true belief. So the idea is that Gettier problems arise if someone is in a situation in which he has a belief that is supported by a certain body of evidence, however, there are facts about this situation that would defeat this body of evidence. This means that given the body of evidence plus the defeating facts there can be no reasonable expectation that the belief for which one has the original body of evidence is true.

The effects of this no-defeating-facts condition are similar to the luck-precluding condition. The requirement that there are no facts that would defeat the evidence makes sure, as one might argue, that the world is in fact such that through the evidence one will arrive at a true belief. So it seems to me that the idea of requiring luck-precluding conditions is similar to the idea to require the absence of defeating facts. However, there are also some differences. Klein proposes a condition that in effect says that there should be nothing in the situation that speaks against the evidence. The luck-precluding condition is rather concerned with facts about the situation that allow us through the evidence to arrive at the truth. In a way Klein’s proposal is more demanding than the proposal of the luck-precluding condition. Let us explain this.

There might be situations in which someone has evidence and there is a luck-precluding condition, but there are also defeating facts as understood by Klein. In such a case Klein would be committed to saying that there is no knowledge, while according to the proposal in this chapter I can say that there would be knowledge. There seems to me to be only a small number of cases in which these two proposals would come apart, namely when there is a misleading defeating fact in the situation. Let me first explain Klein's analysis of knowledge and then describe a case in which the two proposals come apart.

Klein proposes an analysis of knowledge that responds to this diagnosis that Gettier cases always involve a felicitous coincidence. The analysis of knowledge goes as follows: "S knows that p at t1 if and only if (i) p is true; (ii) S believes p at t1; (iii) p is evident to S at t1; (iv) there is no true proposition such that if it became evident to S at t1, p would no longer be evident to S." (Klein 1971, p. 475) So the idea is that for knowledge we need evidence for a true belief and there cannot be anything in the situation that would count as a defeating fact for the evidence we have for the true belief. Let us now describe a case where there is a misleading defeater.

Suppose that you visit a small village in a foreign country, and you meet the mayor of this village. He tells you that recently a movie was shot in the village, involving in particular one of the streets. For this movie it became necessary to build five barn facades. However, the locals have come to be attached to the look of the street with these

structures, and decided to keep them by building them into barns without changing the facades. The mayor furthermore tells you that he as the head of the committee in charge of this project has told the work crew that converts the barn facades into barns to start with number 5. Moreover, yesterday this first step of the work was reported by the work crew to be completed. You find this interesting and you decide to drive by this street and to look at these various barn facades. As you drive by you say to yourself: “So, these are four barn facades and one barn.” Suppose that the belief is true. Suppose however, that the mayor’s instructions were misunderstood by those who convert the barn facades into barns. The only way in which the instructions were misunderstood is so as to result into converting the barn façade number 1 into a barn rather than number 5.

We thus suppose that you believe truly that you are looking at one barn and four mere barn facades. However, there is now a defeating fact. Suppose that you get out of the car and you decide to open the door of barn façade number 5, expecting that there is a barn, however finding a mere façade. You now would have a true proposition that is evident to you, namely that barn number 5 is not a barn but a mere façade. This defeats your evidence from the testimony for supposing that one of the barn facades was converted into a barn. So it turns out on Klein’s proposal that you do not know that you are looking at four barn facades and one barn. Now I take it that we agree that you would not know once you become aware of the defeating fact, by for example getting out of the car and opening the door of structure number 5. The question is whether you know that there are four barn facades and one barn if this defeating fact does not come to light. The problem

is that the defeating fact while defeating your evidence is misleading. It is misleading because you are indeed looking at one barn and four barn facades, and you are indeed relying on good evidence for this, the testimony of the mayor. It is just that there is one mistake, you think that structure number 5 is the barn when it is in fact number 1.

I do not think that this is a decisive example against Klein's definition. It just raises the question what happens if there is some defeating fact that would actually mislead rather than lead to the truth. So the question is what happens if the additional defeating fact worsens the epistemic situation with respect to the belief in question rather than improves upon it. Probably for many of our true and evident beliefs there might be some facts that would lead us astray rather than lead us to the truth. However, even facts that would lead us astray can become evident and be defeaters to other evident true beliefs that we have.

Again, I do not claim that this is a counter-example to Klein's proposal. However, it underscores one of the motivations for thinking about luck-precluding conditions in response to Gettier problems. The idea behind luck-precluding conditions is that they are regularities in the world that make sure that through the evidence we will arrive at a true belief. So what is required for knowledge is that the world is such that through the evidence we will arrive at a true belief. This is compatible with there being misleading facts around. The only exception to this will be developed in chapter four, when we will see that if a defeater is easily available then one is not justified in one's belief. The example just described might be a borderline case. In any event it would be interesting to

think more about misleading defeaters and whether they even if they are undiscovered can undermine our knowledge.

In his “Coherence and the Truth Connection,” Keith Lehrer (2006) explains that the Gettier problem arises because of the possibility that there is no truth connection, no connection between the justification and the truth of the belief for which we have the justification. He furthermore explains that what we need is not just any such connection, but one that results into its being the case that the justification for the belief explains why this belief is also true. He then goes on to show that a coherence theory of justification can provide a truth connection that results into the justification being an explanation of why the belief for which we have the justification is true.

As should be obvious from this brief statement of his position, I agree with Lehrer on the analysis of why Gettier type luck is incompatible with knowledge. The basic idea is that Gettier type luck results into there being no sufficiently tight connection between what renders the belief in question justified and what renders it true. This is exactly the reason why I suggest the luck-precluding condition, a condition that connects through a regularity the evidence to the truth. The problem of the gap involves that the evidence is compatible with the falsity of the belief for which we have the evidence. So we need something to connect the evidence to the truth of the belief for which we have the evidence. And this is what the luck-precluding condition is supposed to do.

Furthermore, Lehrer shows that the truth connection must be sufficiently connected to what is available to the epistemic agent in question from the inside. The truth connection cannot just be an external condition that has no relationship to what is available to the person from the inside. The truth connection should result into the person in question having available to him an explanation of why his belief is true. I agree with all of this. Even though the luck-precluding condition is an external condition, it has to be expected by the person in question, which is an internal requirement. This is something that I will explain more fully in the fourth chapter, but the luck-precluding condition is so to speak mirrored in justification in that the justification leads to the expectation of the luck-precluding condition. This is similar to how Lehrer shows that actual trustworthiness is mirrored in and supported by certain principles of trustworthiness that form part of the coherence system.

Even in the account of justification provided in the fourth chapter, I will leave it open whether the details of the justification should be understood on a coherentist or foundationalist picture. It seems to me that both would be candidates for fleshing out the general account of justification given in the fourth chapter. Furthermore, it seems to me that Lehrer proposes a way of how such justification could lead to an expectation of the truth connection provided by the luck-precluding condition.

Before concluding this section there are two questions that I would like to discuss. The first was raised by Richard Fumerton, and it concerns whether the proposed analysis and

response to the Gettier problem is different from the false belief analysis and response.

The second question was raised by Shaun Nichols and Gerald Gaus, and it concerns empirical data on the Gettier intuition and some methodological implications of it.

Richard Fumerton suggests that according to the analysis of the Gettier problem given in my dissertation there is a belief or belief-like element involved, namely the expectation or assumption that the luck-precluding condition is present. In a Gettier case according to my proposal the luck-precluding condition is absent. This means that this expectation or assumption is false, and if it is a belief or belief-like thing, this means that in a Gettier case there is a false belief. However, the proposal that Gettier cases depend on a false belief has been discussed for a long time. The question thus arises what distinguishes the analysis and response to the Gettier problem that is developed in my dissertation from the false belief response.

Now, in a Gettier case according to the discussion in this chapter the luck-precluding condition is not only absent but not expected or assumed to obtain. For example, in the sheep case the luck-precluding condition would involve a regularity where sheepdogs are followed by sheep from a distance. However, in a Gettier case not only is this luck-precluding condition absent, but it is typically something that the epistemic agent is not aware of. In the sheep case we do not assume that the epistemic agent has any beliefs about such a regularity or assumes such a regularity to obtain. Thus while it is correct that for someone to have perceptual knowledge, the external requirement of a luck-precluding

condition has to be supplemented by an internal requirement, this expectation or assumption that the luck-precluding condition obtains. However, in a Gettier case we have typically not just the absence of the luck-precluding condition, but also a situation in which this luck-precluding condition is not expected or assumed to obtain by the epistemic agent. So it need not on my account be the case that someone has a false belief in a Gettier situation that would then lead to the Gettier intuition.

Furthermore, the assumption or expectation that the gap-closing condition obtains is not a belief in the sense of a premise that we need to reason to the belief about the physical world. It is not present in the sense that it would be required for an inference. In most cases this assumption will be present only in the dispositional sense, where this is linked to what the epistemic agent would regard as a defeater. In the false-belief analysis of the Gettier problem it is usually assumed that the false belief is a step in a reasoning process that then leads to a justified true belief that is not knowledge. No assumption of any reasoning or inference is involved in the account of perceptual knowledge developed in my dissertation.

Even further, in those cases in which the luck-precluding condition is absent but for some reason the epistemic agent expects it to obtain, we would indeed have a false dispositional belief, and this might well lead to a Gettier intuition. In these cases the only difference between the false-belief and the analysis developed in my dissertation would be that I identify which false belief leads to a Gettier intuition. However, this is a

significant difference because in a Gettier situation according to my analysis the flaw does not lie in what the epistemic agent believes, rather in how given his beliefs the world is structured. The problem that leads to a Gettier intuition is primarily that the evidence does not in this instance indicate the truth, that the evidence is in some way misleading. The evidence is misleading because the world is not such so that this evidence will other than by chance lead us to the truth. The Gettier intuition arises because of this mismatch of evidence and truth, not somehow because of a false belief, or so I claim.

Shaun Nichols and Gerald Gaus have raised some questions concerning some empirical data that the Gettier intuition is not shared by all survey participants, and that there might even be cultural differences as to whether someone has such a Gettier intuition about a case. In this dissertation I have assumed that we have the Gettier intuition about the standard Gettier cases as discussed in the literature, and also about some of those I have imagined here. How can this be justified? Furthermore, what exactly is the methodology with which I am working here?

Admittedly I do not have a fully worked out methodology with which I operate here. However, I generally do not look at intuitions as data in the sense that they are somehow independently of the theory justified. Often intuitions are hunches or pointers that signal a direction in which to go. Furthermore, I would doubt a philosophical intuition unless it can at some point be incorporated, understood and explained in a philosophical theory

that is supported by arguments. In that sense philosophical intuitions seem to me to be different from linguistic intuitions. I am not sure exactly how I should respond to someone who after reading this dissertation claims that he has no Gettier intuition about any of the cases discussed in this chapter. However, I recognize that I do not really have an answer to these methodological questions.

3.6 The problem and the proposal revisited

This chapter examines Gettier cases. It proposes that the type of luck present in a Gettier case is non-skeptical epistemic luck. This means that a luck-precluding condition is absent in a Gettier case. Furthermore, the intuition that we have no knowledge in a Gettier case can be explained by means of this absence of a luck-precluding condition. This also means that in order to avoid Gettier cases we need a luck-precluding condition.

However, luck-precluding conditions are strictly external conditions. Even if they were present in a Gettier case, there still would be no knowledge unless the epistemic agent expects or assumes that the luck-precluding condition obtains. I have tried to show that if we have evidence for a true belief and the luck-precluding condition obtains and is expected to obtain, we no longer have a Gettier intuition about such a case. This is a step towards sufficient conditions for perceptual knowledge.

In the second chapter I suggest that the mere presence of skeptical epistemic luck is in the absence of non-skeptical epistemic luck compatible with knowledge. In this chapter we

have seen again that non-skeptical epistemic luck is incompatible with knowledge, in that it leads to a Gettier case. However, if non-skeptical epistemic luck is absent and certain further conditions obtain, then even though there is still skeptical epistemic luck, we do not end up in a Gettier situation. This suggests that skeptical epistemic luck is in the absence of non-skeptical epistemic luck compatible with knowledge.

The next step in order to arrive at a set of necessary and sufficient conditions for perceptual knowledge is to think about what kind of justification is required for perceptual knowledge. I call this type of justification that is required for perceptual knowledge “evidence-based justification.” Based on the discussion so far, we already suppose that it should contain an evidence requirement and an expectation or assumption that the luck-precluding condition is present. The next chapter develops an account of such evidence-based justification.

4. EPISTEMIC AGENCY AND JUSTIFICATION

4.1 The problem and the proposal: What kind of justification is necessary for perceptual knowledge?

Suppose that a person, named “Lucky,” has all the sensory experiences and beliefs of a normal person leading a normal life in a normal world. However, Lucky is a brain in a vat, although she has no inkling of that. Right now Lucky has the sensory experiences and beliefs as of driving to the grocery store. Let us suppose that she has the same internal states that constitute our evidence for the belief that we are driving to the grocery store. She has the same sensory experiences and beliefs we have when we actually drive to the grocery store. Of course it is not true that Lucky is in fact driving to the grocery store – as a brain in a vat she cannot drive or buy groceries. However, the question is whether Lucky is justified in this false belief. It seems that Lucky is a responsible epistemic agent: she does not jump to conclusions, she does not believe things for no good reason at all. Lucky does the best she can. In this instance, Lucky is just unlucky that her belief is false.

This is a case of what is called the new evil demon problem.⁷ The intuitive idea is that we would want to say that Lucky could act in more or less epistemically responsible ways. However, no matter how careful she is in adopting beliefs, they will all be false. The question raised by the new evil demon problem is thus whether and how Lucky could

⁷ The new evil demon problem is developed in Stewart Cohen (1984).

gain a justified false belief. Often the intuition is that a brain in a vat can have a justified false belief. Now, just as there is not just one notion of epistemic luck, there is not just one type of justification. Justification comes in different kinds and different degrees. The question I am concerned with in this chapter is what kind of justification is necessary for perceptual knowledge, and whether a brain in a vat could possibly have this type of justification.

The proposal is that what I call “evidence-based justification” is what we need for perceptual knowledge. Furthermore, we will see that a brain in a vat does not have this type of justification. This is because even though we need evidence in order to possess evidence-based justification, further external constraints have to be in place for us to have evidence-based justification. There is no evidence-based justification without evidence. However, two epistemic agent might have the same evidence, and yet one of them might have evidence-based justification when the other lacks this type of justification.

In previous chapters I have merely spoken of evidence rather than justification. As we will see in this chapter, the type of justification necessary for perceptual knowledge is not identical to an evidence requirement. However, in accordance with the discussion in previous chapters, evidence-based justification involves both an evidence requirement and the expectation or assumption that the luck-precluding conditions obtain.

The account of evidence-based justification developed in this chapter is rooted in an understanding of the nature of our epistemic agency. By epistemic agency I mean those of our capacities that are relevant to whether or not we can gain a justified belief and knowledge. In the first main part of this chapter I will explain these capacities, the nature of our epistemic agency in so far as it concerns perceptual knowledge. In the second part I will develop the notion of evidence-based justification on the basis of this understanding of our epistemic agency. The third part explains how evidence-based justification is tailored to the other main requirement for perceptual knowledge, the luck-precluding condition. In the next chapter an argument is given that evidence-based justification is necessary for perceptual knowledge.

4.2 Epistemic agency

By epistemic agency I mean the set of capacities that we have and that are relevant to whether or not we can gain a justified belief and knowledge. All those capacities that are required or helpful for us to gain a justified belief or knowledge are part of our epistemic agency. Epistemic agency does not require practical agency. Even someone who is completely paralyzed can have epistemic agency. It is not required for epistemic agency that we can go to the library or engage in other practical activities in the search of knowledge.

Epistemic agency consists in certain epistemic capacities that we have.⁸ I am in this dissertation only concerned with perceptual knowledge, and relevant to such perceptual knowledge are three capacities that we have. There is the capacity to receive sensory experiences, the capacity to read these sensory experiences, and the capacity to evaluate the readings of the sensory experiences. These are the three capacities that I will describe in this section on epistemic agency. Furthermore, I will capture these capacities from the inside. This means that I will talk about these capacities in terms of how they present themselves to us from the inside. This is important since I ultimately aim at explicating a notion of justification that depends on evidence, where such evidence is understood to consist in internal states. I will, however, find it necessary to capture one of these capacities also from the outside, rather than merely from the inside.

4.2.1 The capacity to receive sensory experiences

In this section I will explain the capacity to receive sensory experiences as input. We each have certain internal states that come as input. By input I mean that they come to us as though from something independent of our mind.⁹ Right now I have a visual sensory experience as of writing appearing on a computer screen. Let us compare the sensory experience as of writing appearing on a computer screen and my belief that writing is

⁸ A similar general understanding of epistemic agency can also be found in Reed (2001).

⁹ In this discussion of sensory experience as input I do not mean to enter the discussion of the myth of the given. The idea here is merely that we can distinguish between sensory experiences and beliefs according to three criteria, and that sensory experiences behave as though coming from something external. These properties of sensory experiences are so to speak the properties that make it seem that sensory experience is something given. This is compatible with their being no robust sense in which the content of a sensory experience is something given. This is also compatible with a non-foundationalist picture of justification. For a brief discussion of the myth of the given see Roderick M. Chisholm (1964).

appearing on a computer screen. Both are about something external. My belief is about a computer screen, something that exists independently of me. My sensory experience as of writing appearing on a computer screen is similarly about a computer screen that exists independently of me. However, only the sensory experience is input. Let us explain this in detail.

Let us compare sensory experiences to beliefs in three respects: First, sensory experiences are beyond our control or will. Beliefs on the other hand are to a lesser degree and in a different way beyond our control and will. Second, sensory experiences are transparent in ways in which beliefs are not. Third, sensory experiences have a different kind of unity and depth than beliefs. It is because of these three differences that sensory experiences constitute input, while beliefs are not input.

First, sensory experience is not controlled by us or willed into existence. Right now I am sitting at my computer and I cannot will myself to have the sensory experiences as of sitting on the beach and basking in the sun listening to the sounds of the ocean. Sensory experiences are input in the sense that they are not controllable or willed into existence by me. I cannot make myself have the sensory experiences as of lying on the beach when I have sensory experiences as of typing on a computer. Of course, I can try to put myself in a situation where I would have sensory experiences as of lying on a beach rather than typing on a computer in my apartment. I can also close my eyes and plug my ears. But it

still seems clear that sensory experiences are not in my control; they are not something I can will into existence.

It is often said that beliefs are involuntary, that they are not in our control or subject to our will. There is some truth in this, in that we cannot will ourselves to believe some random thing. For beliefs about the physical world, those that would be candidates for perceptual knowledge, we need sensory evidence. However, we can be mistaken about the merits of the evidence. We can look for evidence or refuse to believe on the basis of what is in fact strong evidence. As I will explain somewhat more when talking about the third capacity that defines our epistemic agency, beliefs are not involuntary in the same way and to the same degree to which sensory experiences are involuntary. We do not have full voluntary control over our beliefs, however, our beliefs do not seem to depend on our situation in the way in which beliefs are. I will later on explain that beliefs adopted on the basis of defeasible evidence require a decision on part of the epistemic agent. No such decision is needed for receiving sensory experiences.

Second, sensory experiences have a characteristic that is often called “transparency.”¹⁰

The idea is that I do not notice the internal state in virtue of which I learn about external states. Sensory experience produces an internal state, but this internal state is transparent, we look through it into the world just as we can look through glasses or a window at what is going on in front of us. The internal states are transparent so that we usually do not

notice their presence. Only if something goes wrong, if we have reason to suspect that the internal state is misleading in some way, is it that we become aware of the internal state. Our sensory experiences suggest to us that we live in a world full of various objects with various properties, and we are part of this world. Our sensory experiences place us in a world full of objects and their properties without it being obvious that our epistemic access to the world is mediated by internal states.

Beliefs are not transparent in the same way in which sensory experiences are transparent. My beliefs do not place me in the world in the way sensory experiences place me in a world. Beliefs can be disconnected from my actual situation. My beliefs are about the world, but it is clear that they are internal states that represent external states. It is not the case that the internal state, the belief, somehow disappears and I only notice what the belief is about. While the internal states in sensory experiences are similar to looking through glasses or a window, the internal states in beliefs are similar to writing in a book. We see what the writing in a book is about, but we also see the representation, the writing.

Third, sensory experiences allow us to focus our attention on various aspects of them. We can look more closely at an object. For example, we can get out of the car and walk around the barn, knock on its walls, open its door, and so on. Our sensory experiences have a certain unity and depth in that we can look more carefully, shift attention to

¹⁰ This term is used in particular in the philosophy of mind. I do not mean to endorse any particular

various aspects of them, seek further sensory experience of the same object. In doing so we can gain new knowledge of the world around us.

Beliefs do not have the depth and unity of sensory experiences. I can shift attention from one belief to the next, and I can see connections between them. Maybe some beliefs entail other beliefs, or they are inconsistent, or they support each other in various ways. Given one belief I can also emphasize various aspects of it. However, that is the extent to which beliefs have unity and depth. It is not the case that I can learn something new and connected about the physical world by focussing on various aspects of a belief, or by seeing logical connections between beliefs. At least I cannot thus learn more about the world in the same way in which a shift of attention and further related sensory experiences can tell us more about the physical world.

Let us explain this somewhat more. Right now I have sensory experiences as of typing on a computer. I can look more closely at various details of the computer, I can learn new things about the computer and other objects in the physical world through shifting attention and gaining more sensory experiences. Nothing like that will ever happen when I examine my belief that I am typing on a computer. Suppose that later on I sit down at my kitchen table and examine the belief that I was typing on a computer earlier on. It just is not the case that I will learn more about the computer or other objects on the desk just by examining my beliefs. The difference seems to be that different sensory experiences as

of typing on a computer are somehow all connected by the external object, the computer. However, different beliefs about typing on a computer have at best various logical relationships among each other. So in the case of beliefs the depth results from connections between the internal states, the representations. In the case of sensory experiences their depth results from each of them telling us something about one object, the computer. So the connections between different sensory experiences are not a matter of connections between the different representations, but are due to connections between what is represented.

Let us now see in what sense these three characteristics render sensory experiences input while beliefs are not input. These three differences between sensory experiences and beliefs make no mention of a causal relationship to something external; they are characterized from the inside, and even a brain in a vat could have internal states with these characteristics. It is not clear what one usually means when one speaks of input. Here the idea is just that input is an internal state that comes to us as though from something external, rather than as though being the product of some mental activity. If an internal state comes to us as if not controlled by us or willed into existence, it seems that it comes as though from some source independently of us. Similarly if an internal state is transparent so that it seems to us as though we are directly aware of things outside of our mind, it seems again that these internal states come to us from outside. If an internal state comes to us as though it has this kind of unity and depth that is created through the object represented rather than through logical relationships between the representations, it seems

that this internal state comes to us as though from something external. In that sense then is it that sensory experiences are input.

What does it mean to have the capacity for receiving sensory experiences? It means that we are capable of sensory experiences as input. It does not mean that this input has a specific causal origin, or that it is somehow accurate. It just means that among our internal states we can distinguish some of them as being sensory experiences rather than beliefs. I take it that the three characteristics of sensory experiences are familiar to all of us. Furthermore, even a brain in a vat can have this capacity.

4.2.2 The capacity to read sensory experiences

Let us now consider the second capacity, to read our sensory experiences. By reading a sensory experience I mean seeing which beliefs it suggests as true. By “seeing which beliefs it suggests as true” I do not mean that we somehow need a belief as to which belief is suggested as true by the sensory experience to read it. I also do not mean that this is somehow an additional sensory experience the object of which is somehow what the sensory experiences indicate. In some respects reading a sensory experience is similar to reading a book. When we read a book we see what the words mean. We have acquired the ability to match the lines in the book with the content they contain. Similarly when we read a sensory experience we take it and match it with the beliefs that it suggests as true. If, for example, I describe a sensory experience as a sensory experience as of typing on a computer, I read the sensory experience. The idea is that each sensory experience

has a natural relationship to some beliefs in that it suggests that these certain beliefs are true. The sensory experience I have right now suggests that I am typing on a computer. It does not suggest that I am sitting on the beach basking in the sun and listening to the sounds of the ocean.

A reading of a sensory experience could be expressed as the claim that it seems as though I am typing on a computer or sitting on the beach. Even if we immerse a stick partly in water it will be the case that the resulting sensory experience suggests that the stick is bent. My reading of the sensory experience could be expressed as the claim that it seems as though the stick is bent. The sensory experiences have a certain character that allows me to read it, to associate it with those beliefs that it naturally suggests. I might not believe that the stick is indeed bent, but it still seems as though it is bent. So the sensory experience suggests by itself that the belief that there is a bent stick is true. If I know that this is an illusion I might conclude that the stick is indeed not bent at all. But this belief that the stick is not bent is not suggested by the sensory experience as of the stick partly immersed in the water. I cannot read off of the sensory experience that the stick is not bent.

What exactly allows us to thus read our sensory experiences, to associate the beliefs that it naturally suggests with the sensory experience is mysterious. It seems that in part it is learnt through learning a language and the appropriate concepts. However, even animals seem to be able to thus read sensory experiences, at least in some cases. I am not sure at

this point how to make the idea more precise that a sensory experience naturally suggests that certain beliefs are true. This is a really difficult question, and I will not pursue it further in my dissertation. In what follows I will assume that we are capable of reading our sensory experiences, of seeing which beliefs a sensory experience naturally suggests to be true.

4.2.3 The capacity to evaluate readings of sensory experiences

Let us now consider the third capacity, to evaluate the readings of sensory experiences. The first capacity concerns our ability to receive input. The second capacity allows us to get from the input to the beliefs that the input naturally suggests as true. The third capacity concerns how on the basis of this reading of the input we form beliefs about the physical world. So we have input, we then read this input, which results into “seemings,” and then we adopt beliefs about the physical world on that basis. However, these are three steps and corresponding capacities that we distinguish only in a theory, not in normal belief formation. Normally we will go instantly to the belief that there is a barn in front of us, rather than making these three steps one after the other. We will not first somehow receive input, then read it, and then adopt a belief on its basis. Nevertheless these three steps are present in the sense that if something goes wrong, if for example we encounter an illusion or a defeater, we will retreat from the third step to the second step, or even reconsider the second step.

Let us explain this a little bit more. Suppose again that I say that there is a barn in front of me. Suppose now that I learn that I am driving through barn façade county. In that case I will retreat from the third step to the second, and say that it merely seems as though there is a barn out there. Now suppose that you tell me that it does not seem like a barn but like a cottage. This will not mean that I retreat to the first step, because the first step is just the basis for the second step, not itself something that I can express. So I will reconsider the second step and arrive at a different second step when I say that it indeed seems more like a cottage. If I form any beliefs, if I participate in a conversation I need to make the second step. The first step, the input, is just the basis for the second step.

Let us now consider this third step. We want to see how we go about adopting beliefs on the basis of reading our input. The most important characteristic of this capacity is that we do not need to believe whatever our sensory experiences suggest to us. We read a sensory experience as suggesting that the stick, partly immersed in water, is bent. This reading however does not force us to believe that it is indeed bent. We have a choice of either accepting what our reading suggests or not. This does not mean that we have complete voluntary control over our beliefs. Rather, this means that there is no automatic relationship between sensory experiences and beliefs. Sensory experiences do not trigger automatic belief in what they suggest. It might be that in some animals certain sensory experiences mandate belief in that the animal cannot form a belief that rejects the sensory experience. However, for us it seems clear that we need not believe whatever our sensory experiences suggest to us.

Now we need to ask how then belief is formed on the basis of sensory experience. It is often said that belief is not voluntary, that we cannot make ourselves believe some random claim. There is truth in that in so far as we need evidence or the semblance of evidence to form a belief that is a candidate for perceptual knowledge. We form beliefs about our immediate physical surroundings, and probably about many other things also, on the basis of sensory evidence. Sometimes we form beliefs on the basis of not very good evidence. It still seems that for most of our beliefs we need some kind of evidence for them. We cannot form the belief in some random claim, for example that the number of stars in the universe is even, in the absence of any evidence for it. In the case of possible perceptual knowledge, we seem to be unable to form a belief, for example that there is a table in this room, in the absence of any sensory evidence for it. If perceptual knowledge can be described as seeing or hearing or touching, and so on, an object, then the belief that is hopefully an instance of perceptual knowledge requires sensory evidence for it.

Now in the case of perceptual knowledge this evidence is defeasible.¹¹ This is already contained in the claim that there is a gap between our evidence for our perceptual beliefs and their truth. Again, the gap means that it is possible to have this evidence for the belief when the belief is false. This also means that there is no deductive relationship between a description of our evidence and the belief it indicates as true.

In the case of a deductive inference, we can add anything we want to the premises, if the original premises entail the conclusion, so do the premises plus the additional stuff. So if a set of premises *P* entails deductively a conclusion *C*, then so does the conjunction of *P* and any addition at all. However, in the case of evidence that is not deductively related to the belief for which we have the evidence, the same does not hold. Defeasible evidence means that additions to the evidence can have the result that this new set of evidence no longer supports the original conclusion. Let us illustrate this.

Suppose that I believe both that all sheep have four legs and that this specific animal in front of me is a sheep. Then this would deductively entail that this animal in front of me has four legs. Suppose now that I add various claims to the premises. My new set of premises now looks as follows: All sheep have four legs, and this specific animal in front of me is a sheep, and all sheepdogs have three legs, and this specific animal in front of me might be a sheepdog. From this set of premises it still follows that the animal in front of me has four legs.

Suppose now that I have the sensory experiences as of seeing a sheep at a short distance in front of me. This is evidence for the belief that there is a sheep at a short distance in front of me. However, this evidence is defeasible. This can be seen if we add the following claims to the evidence. Our new set of evidence now reads as follows: I have the sensory experience as of seeing a sheep at a short distance in front of me. A local tells

¹¹ The characterization of defeasible reasons can be found in John Pollock and Joseph Cruz (1999). See in

me that there are no sheep in this area but a lot of sheepdogs that look exactly like sheep from a short distance. This new set of evidence no longer supports the conclusion that there is a sheep in front of me.

This additional evidence is a defeater. According to John Pollock there are two kinds of defeaters.¹² Given an evidence-belief pair, a rebutting defeater is evidence that suggests that the belief is false. An undercutting defeater is evidence that suggests that the original evidence is not good evidence. If we think about the evidence-belief relationship as an argument, a rebutting defeater would show the conclusion to be false, and an undercutting defeater would attack either a premise or the inference. In the example above we have given both kinds of defeaters. That there are no sheep in this area is a rebutting defeater. That there are sheepdogs that look exactly like sheep from a distance and that run around in this area is an undercutting defeater.

Now let us return to characterizing this capacity of evaluating our readings of sensory experiences. Basically this capacity is the capacity to appreciate defeaters and to look for defeaters. Our ability to withhold belief in something that a sensory experience suggests requires evidence, namely a defeater or our belief that there might be such a defeater. Now often the defeaters are not obvious, we have to look for them. Thus in accepting what a sensory experience suggests to us and to believe it requires a decision, namely not to look further for defeaters. Let us explain this furthermore.

Suppose again that you have the sensory experience as of a sheep in front of you at a short distance. Suppose furthermore that you are not aware of any defeaters. This still does not trigger belief unless we decide – often not conscious – not to further look for defeaters. Belief on the basis of defeasible evidence requires a decision not to look for a defeater. This decision, though in most cases not a conscious decision, is then what allows us to form the belief. We will explain in a little bit in what respects this decision is similar to and different from practical decisions. Before doing so let us explain somewhat more and illustrate this idea of a decision that is required for forming a belief on the basis of defeasible evidence.

The capacity to evaluate our sensory experience allows us to make or refuse to make this decision to go look for defeaters. And this is the reason for calling these capacities “epistemic agency.” Even as epistemic agents we have to make certain decisions. In the case of perceptual knowledge we need to decide whether or not to look for further defeaters. It is only this decision that then triggers belief.

Let us look at a case where this decision is conscious. Suppose that a detective investigates who caused a certain property damage. He will try to gain evidence as to who caused the property damage. There will be witnesses to interview, there will be physical evidence to collect and investigate, and so on. All of this evidence is defeasible.

¹² See in particular: Pollock and Cruz (1999), p. 196

It is possible that there is further evidence, for example that a witness is lying or that some physical evidence is misleading, and so on. So at some point the detective will have to decide to take the evidence he has to show that a certain person caused the property damage.

This decision is a genuine decision. In some cases it might be epistemically permissible to make the decision, but also to go look for further evidence and defeaters. In some cases the decision will not be made consciously; we will just proceed on the basis of such a decision. This is what renders us epistemic agents: We are able to reject our defeasible evidence and to look for further evidence and defeaters. If we make the decision not to look for further evidence, then this will trigger belief in what our sensory experiences suggest. Belief on the basis of defeasible evidence is thus linked to this decision, and thus to our epistemic agency.

Let us consider an objection to the claim that forming a belief on the basis of defeasible evidence requires a decision. It might be said that talk of a decision suggests that beliefs are somehow formed voluntarily. In particular if we consider basic perceptual beliefs, as for example the belief that I am sitting in my apartment, typing on a computer, how could it possibly be true that such beliefs result from a decision?

Let us consider the way practical actions depend on decisions. When we think about decisions we often consider only those that are the result of deliberation, where we find

ourselves pulled in different directions. So the decisions that we remember are those that are in some respect difficult, where we consider reasons for different options and then somehow decide between these different options. However, there are many actions that depend on a decision where no such deliberation takes place, because they are a matter of habit or standing preference. When we drive to the grocery store we will in a sense decide to stop at the stop sign, but not in the sense that we deliberate between stopping and not stopping at the stop sign. When we get up in the morning we will in a sense decide to turn on the coffee maker, but not in the sense that we deliberate between having coffee and not having coffee.¹³ Nevertheless there is a decision involved, one that renders the resulting action voluntary in a sense.

Something similar goes on in adopting many perceptual beliefs. Just as we might have a standing preference for coffee, so might we be inclined to form perceptual beliefs about familiar surroundings on the basis of a certain kind of sensory experience. We often do not even notice the objects in our familiar surroundings particularly, but just form the appropriate perceptual beliefs on the basis of a certain habit. The decision involved is not one that is made consciously and as a result of deliberation.

However, there is also a difference between the way we make decisions as practical and as epistemic agents. As practical agents we can form a considered judgment and still fail to act on it. So even if we perceive that we have good reasons to engage in a certain

¹³ Matthias Steup uses similar examples to establish a similarity between actions and beliefs. His claim is

action we might not decide to do so. As practical agents our decisions are sometimes only loosely based on the reasons that we see for making a specific decision. As epistemic agents, our decisions are more tightly connected to the reasons we see. It would be very strange for someone to say that he has much better reasons to believe that p rather than that q , but that he still believes that q . It seems that as practical agents we can have very good reasons to make a given decision and still refuse to make it. It is not clear that something similar is really possible for us as epistemic agents. If we have very good reasons for not looking for further defeaters, it is not clear that we really can refuse to believe what we have the reasons for. We might still keep an open mind in a sense or decide to investigate the matter further. However, in so far as the decision leads to belief, if we have very good evidence for a given belief, it seems puzzling how we could decide not to believe on the basis of this very good evidence. We might be mistaken about the strength of reasons, we might believe for bad reasons. However, if we see that we have very good reasons for believing something, then it seems that we make the decision to believe even if we somehow keep an open mind and want to investigate further.

There is thus a difference between the decisions we make as practical agents and those we make as epistemic agents. As epistemic agents our decisions that trigger belief are more closely linked to our perception of the reasons. Nevertheless, there is still a decision involved in that we can reject our defeasible evidence and decide to go look for further evidence and defeaters. Furthermore, there will be many cases in which the evidence

that if our actions are free in a compatibilist sense, then the same compatibilist freedom is enjoyed by our

does not mandate a particular decision to accept the evidence or to go look for further evidence. In such cases we might deliberate about this kind of a decision in a way similar to deliberations about practical decisions. Many other decisions will be a matter of standing disposition to accept certain evidence similarly to practical decisions based on habits and standing preferences. Thus even though the way epistemic decisions depend on evidence is different from the way practical decisions depend on reasons, there are still sufficient similarities to call both “decisions.”

I am not sure how psychologically and empirically accurate the picture is that I am painting here. One might thus object that people sometimes believe something and still look for defeaters, or that they might not believe something without having found a defeater, and so on. My dissertation is not and cannot be an empirical investigation of how people adopt beliefs. It does seem to me, however, that these three capacities are familiar to all of us from the inside, that we indeed have these capacities. In later sections I will try to show that epistemic agency thus understood is crucial for our ability to gain the kind of justification necessary for perceptual knowledge.

Summary Box 9: The nature of our epistemic agency

In the case of perceptual knowledge there are three epistemic capacities that characterize our epistemic agency. The first capacity consists in our ability to receive sensory experiences as input. The second capacity allows us to get from the input to the beliefs

beliefs. See Steup (2006).

that the input naturally suggests as true through reading our sensory experiences. The third capacity concerns how on the basis of this reading of the sensory experiences we form beliefs about the physical world. So we have input, we then read this input which results into “seemings,” and then we adopt beliefs about the physical world on that basis. These are three steps and corresponding capacities that we distinguish only in a theory, not in normal belief formation.

It is because of the third capacity that epistemic “agency” is an appropriate choice of word. When we adopt a belief on the basis of a reading of a sensory experience, we do so on the basis of defeasible evidence. Belief on the basis of defeasible evidence requires a decision not to look for a defeater. This decision, though in most cases not a conscious decision, is then what allows us to form the belief.

4.2.4 Comparison with other views

Greco (2000) proposes what he calls a set of “elements of perception,” that look similar to the three capacities just sketched. He distinguishes between the following three elements: “i. The uninterpreted *qualia* of sensations. (Sensory appearances characterized as lacking conceptual content.) ii. Interpreted *experience*. (Sensory appearances characterized as involving conceptual content.) iii. *Beliefs* about objects in the world. (Greco 2000, p. 98) There are some similarities here. Interpreted sensory experience seems to look a lot like readings of sensory experience. However, Greco does not think of these three elements as three steps, each of which is accomplished through a different capacity.

Furthermore, reading a sensory experience is in my dissertation understood to mean matching the sensory experiences with those beliefs about the physical world they naturally suggest as true. Interpreting a sensory experience is according to Greco not obviously linked to the beliefs about the physical world the sensory experiences suggest as true. Greco's idea seems to be that we have these appearances, for example of a red triangle, and arrive through them at beliefs about the physical world. The step of interpreting the sensory experience does not result, so to speak, into possible beliefs about the physical world, but would rather result into a description of the experience. Greco seems to think that appearances have content that is not identical to the beliefs about the physical world the appearances suggest to be true. This is very different from what I propose in the dissertation. Reading a sensory experience results into proposals of the beliefs about the physical world that the sensory experience suggests as true, not into some kind of a description of how things appear.

4.3 Evidence-based justification

4.3.1 The proposal

In this section I will develop an account of the type of justification required for perceptual knowledge. I call this type of justification "evidence-based justification." The account of evidence-based justification is based on the understanding of our epistemic agency developed in the preceding section. Let us consider again what the picture is that I have painted of our epistemic agency. In the case of perceptual knowledge there are three

steps and three corresponding capacities to make these steps. First we receive input, our sensory experiences. Then we read these sensory experiences, which means that we match them with the beliefs that the sensory experiences naturally suggest as true. The result of this second step without the third step would be a belief about how things seem to be. The third step then evaluates whether how things seem to be is also how they are, in the light of other evidence we have or reasons we might have to look for further evidence. Only if in this third step we approve, so to speak, of the suggestion we have read off of the sensory experience do we end up with a perceptual belief.

Let us work through an example. Again I am driving in a car and upon looking out of the car window I arrive at the belief that there is a barn out there. I do not consciously run through these three steps, but they can be brought to my attention through challenges to this belief that there is a barn out there. The belief might be challenged through introducing a defeater, for example that I am driving in barn façade county. If I accept this defeater I will retreat from the claim that there is a barn out there to the claim that it seems as though there is a barn out there. This is a step from believing what I read off of my sensory experience, to rejecting it and retreating to the weaker claim that I have a sensory experience such that I can read off of it that there is a barn out there. So this defeater makes me go from step three to step two. I retreat from a claim about the physical world to a claim about my sensory experience, namely that it naturally suggests that there is a barn out there.

Suppose now that someone comes along and claims that the building looks more like a cottage than a barn. This would be a challenge to step two. The challenge focuses on what can be read off of the sensory experience. This means that we begin to consider whether our sensory experience should be read as a sensory experience as of a barn or as of a cottage. This is not a retreat from step two to step one. We cannot retreat from step two to step one, because step one does not furnish us with a belief or something we can express in a sentence. Only what we read off of sensory experience results into a belief and something that we can express in a sentence. However, this second challenge brings out that there is indeed a step two that is based on step one, the input. So when we want to settle whether the building looks more like a barn or more like a cottage, we will consider what we exactly mean by a barn and by a cottage and how we can distinguish sensory experiences as suggesting either. So this challenge results into a move from a specific reading of a sensory experience to a different reading. In doing so, we become aware that there is such a step involved in gaining perceptual knowledge.

Again, I do not claim that these three steps are taken consciously when we gain a perceptual belief. We will often run immediately through these three steps and become conscious only of having formed the perceptual belief on the basis of a sensory experience. These steps are nevertheless part of our forming a perceptual belief in the sense that there are distinct and familiar challenges to our perceptual beliefs that raise these steps to consciousness. I will now briefly sketch the proposal of evidence-based

justification, and then develop a definition of this type of justification by considering a variety of different brain in vat situations.

The main idea is that evidence-based justification requires evidence, where this is understood to consist in internal states, the sensory experiences and beliefs, plus some external constraints. Evidence-based justification thus requires an internal condition, that we have evidence for our beliefs. However, there are further constraints on evidence-based justification, constraints that are external and about which we can be mistaken as to whether they obtain in a given case. The intuitive idea here is that the third capacity of evaluating the readings of the input and making the decision of whether to accept the evidence or whether to look for further defeaters must be free from deception and manipulation. For that to be possible the other two capacities also cannot yield input and readings of the sensory experiences that are completely and fully subject to deception and manipulation. This means that in some scenarios of deception and manipulation we will not have evidence-based justification because we are not able to make the decision as to whether to accept evidence or to look for defeaters in a reasonable way. The evidence-based justification requires an internal condition in the sense that the evidence for a given belief, the other beliefs and sensory experiences, is something internal. However, there are some external constraints on evidence-based justification. This is what we will have to explain now in detail. In the next section we will give a reason for defining evidence-based justification in this way by showing that this type of justification is perfectly tailored to the luck-precluding conditions.

4.3.2 Evidence

Let us begin by distinguishing between evidence and evidence-based justification. I will use these two concepts in a somewhat technical sense here. By evidence I mean those internal states that suggest that a given belief is true. In the perceptual case the evidence would consist in a sensory experience together with other relevant beliefs. For example, my sensory experience as of a part of a barn might constitute my evidence that there is a barn out there. Often other beliefs will be part of the evidence. For example, evidence consisting in my sensory experience as of a sheep can be strengthened by the belief that there are no sheepdogs in the area. By evidence I thus mean those internal states that suggest that a specific belief is true.¹⁴

Let us consider two objections to thinking of evidence as internal states. First, when asked what our evidence is for a given belief, we will usually point to something external rather than mentioning internal states. We thus need to address the objection that these answers to questions as to what our evidence is show that our evidence consists, not in internal states but something external. Second, one might think that if our evidence consists in internal states then some form of introspection is the way of understanding what our evidence is. These are the two most pressing objections to the claim that our evidence consists in internal states. Answering these objections does not fully justify

¹⁴ The details of how evidence supports the belief in question could be worked out so as to be compatible with foundationalism or coherentism.

thinking of evidence as internal states. In this section I will merely provide answers to these two objections.

Let us now consider the first objection. When asked what my evidence is that there is a barn I will usually say something like the following: “I have seen the barn.” “A local has told me that there is a barn.” “The map says that there is a barn.” All of these statements contain reference to something external. If my evidence is that I saw a barn, then part of this evidence is a causal relationship between me and an external object. If my evidence is that someone told me that there is a barn, then part of this evidence is some other external object, this other person, that said something. If my evidence is that a map says that there is a barn, then part of this evidence is an external object, the map, and what is shown on this external object. So how can one defend the claim that our evidence is nevertheless something internal?

That we make such claims when asked for evidence seems to me obvious. We do not ordinarily talk about internal states as evidence. Instead we talk about something external as evidence. Moreover, this is due to an important property of evidence. Evidence is not something idiosyncratic or private. There is no such thing as evidence that is only evidence for me. Instead evidence is something that can in principle convince anyone. So when we say these things about external states as evidence, we do so because of an important property of evidence: evidence is something publicly accessible, something that can convince other people as well. The idea is of course not that every bit of

evidence needs to be as a matter of fact available to other people as well. A lot of evidence is a fleeting matter, where one needs to be in a specific place at a specific time to gain this evidence. However, the idea is that evidence is something that can in principle convince anyone who as a matter of fact possesses this evidence. It needs to be possible that evidence would convince other people also if they had the same evidence. It cannot be that evidence can only convince one person.

Now beliefs and sensory experiences of one person are private and idiosyncratic. Nobody can share the exact beliefs and sensory experiences of another person. One might have similar sensory experiences and beliefs in the same proposition. But there is a sense in which nobody other than the person himself can have his beliefs and sensory experiences. However, evidence is something public, and so one of the things we do when responding to questions about evidence is explaining how one might acquire the sensory experiences and beliefs that then would constitute the same evidence for the same belief.

Let us explain this furthermore. If someone asks what my evidence is that there is a barn, I might say that I have seen the barn. In saying this, one of the things I do is that I give a recipe as to how to acquire the evidence that supports my belief that there is a barn. I tell this other person that one can see the barn in the appropriate circumstances. Similarly I might say that a local has told me that there is a barn. This person can also tell other people, and thus again I give a recipe as to how to acquire the evidence, the testimony, for the belief that there is a barn. If I say that the map shows that there is a barn, I again

give a recipe as to how to acquire the evidence that there indeed is a barn. Let us be more precise.

Let us take one example and discuss it in greater detail. Suppose that I say that there is barn in such and such a field. You ask what my evidence for this claim is. I respond that the map shows that there is a barn in that particular field. In saying this I say all of the following: “I have good evidence that there is such a barn due to my looking at a map. By looking at a map I acquired this evidence. Other people can acquire this evidence through looking at the map. I am confident that my evidence is good for this belief.” Suppose now that you respond: “The map is outdated. Half of the buildings on it no longer exist.” I would now respond in the following way: “From reading this map it seemed to me that there is a barn, but this is not good evidence.” So I will no longer regard the reading of the map as good evidence that there is this barn.

When we thus give as evidence something external, we do two things: We both suggest how one can acquire this evidence and express our confidence that it is indeed good evidence. If a defeater is introduced we will retreat to a claim that we thought that something was good evidence. It is then that we might talk about our beliefs or sensory experiences. Talking about our beliefs or sensory experiences by saying that it seems to me that there is a barn, or I believed that the map is current seems to imply some doubts about whether this is indeed good evidence. When I say: “It looks like a barn, it seems like a barn, I thought that the map is current,” I suggest that my evidence is doubtful.

Talking about our internal states in this way always suggests that they might not be good evidence or that there is something doubtful about it.

We have now suggested an answer to the question of why our responses to demands for evidence usually involve claims about external objects, and that statements about our internal states are usually not made in response to questions about our evidence. Evidence is not something private or idiosyncratic. Evidence is essentially public, something that is capable of convincing anyone. Thus when we think that we have good evidence for a claim, we will talk about some external objects, expressing the confidence that others can acquire the same evidence and become convinced of the same belief. It is only when we think that we do not have good evidence, if there is something doubtful about it, that we speak of the internal states that we take or took to provide us with good evidence. Such a statement always carries with it the qualification that there is something doubtful about this evidence.¹⁵

Given that evidence is something public and that good evidence can convince anyone, we will talk about it in terms of external objects when we are confident about it, but retreat to something internal if we are doubtful about it. This practice is a result of the predicament we find ourselves in. Our access to the external world is mediated through internal states that provide the evidence. These internal states are private and cannot be fully shared. However, evidence is something public and something that can be shared. Thus we talk

about how one might acquire this evidence, or about those features of it that can be shared rather than referring to these internal states directly as evidence. The idea is also that my having a certain belief or a certain sensory experience is not by itself evidence capable of convincing other people, unless other people can acquire similar sensory experiences and beliefs with the same content. That I have a certain belief is no reason for you to be convinced of anything. So when talking about evidence, we somehow have to talk about something publicly available rather than something that just happens to go on in our minds.

Now one might object that there is a conflict between what we have said so far. Have we not said that internal states constitute evidence, that internal states are private, but that evidence is essentially public? Are these claims not inconsistent with one another?

Internal states are private in the sense that they are states of one particular person and in that they as particular states cannot be shared by another person. However, internal states are evidence in virtue of features that can be shared by another person. So we have a token sensory evidence that might reflect something idiosyncratic about that person.

Suppose that someone has a sensory experience as of a barn. This particular sensory experience might have some features that are idiosyncratic to the person: maybe his color vision is in some way different from that of many other people. Or maybe he has lost the ability to perceive objects in a three-dimensional way. There might be many ways in which if we could compare this person and some other people's sensory experience as of

¹⁵ I am ignoring here contexts of philosophical and epistemological discussion. Often in such a discussion

the barn they would differ. However, we then read off of this sensory experience that it is the sensory experience as of a barn. And that is a feature that many different sensory experiences can share. So even if many different people have slightly different sensory experiences as of a barn, as long as they can all read them to be sensory experiences as of a barn, they can understand them to be evidence for the belief that there is a barn.

In other words those features that render a sensory experience private and idiosyncratic are not those that render them evidence. When we read off of a sensory experience that it is a sensory experience as of a barn, we abstract away from those features that cannot be shared by another person. Many different sensory experiences can be sensory experiences as of a barn. There is no need to have identical sensory experiences for them each to be evidence that there is a barn. When we read a sensory experience, we abstract away from many of the particular details of it. Instead we focus on something that can be shared, namely that it is a sensory experience that naturally suggests that it is true that there is a barn out there. So it is compatible to say that internal states are idiosyncratic and private but also evidence, if they are evidence in virtue of abstracting away from what renders them idiosyncratic and private.

Let us now consider the second objection. It might be objected that if our evidence consists in internal states, then we should become aware of this evidence through introspection. However, it is obvious that we do not read our sensory experiences through

statements that it seems that *p* have a special defined meaning. The contexts I talk about would be those of

introspection. It is not the case that we can close our eyes and introspect our minds and thereby read our sensory experiences. Clearly, one cannot see anything in the literal sense with one's eyes closed. So the internal states that sensory experience produces are not mental states that we can somehow introspect. If we close our eyes and try to examine visual sensory experiences, there is nothing to introspect. That sensory experiences are internal states does not mean that they are somehow mental states that we examine through introspection.

Being aware of a sensory experience is also not being aware of a representation that is perceived to be different from what is represented. Reading our sensory experiences is not like looking at a photograph or a painting. If we look at a photograph or a painting, we will usually realize that we are not looking directly at what it shows but just at a representation. In the case of sensory experiences we are not looking at images or representations in that sense. Looking at the external world and looking at the sensory experience that is produced by the external world are indistinguishable to us. Looking at sensory experience is not like looking into your mind, it is like looking at the world and its objects.

What we have said so far does not fully answer the question of why we should think of evidence as internal states, what exactly these internal states are, and how we can see the world through these internal states. It seems to me that these three questions are

normal situations in which we, for example, come to know that there is a barn in the field.

interconnected and cannot be answered by epistemology alone. Rather we would need to tackle them in the philosophy of mind that is informed by epistemology and cognitive science. In the following I will assume that evidence consists in internal states rather than something external. Again, I do not investigate here fully what exactly goes on when we read a sensory experience, and what exactly the nature of the capacity is that allows us to thus read a sensory experience. This would be a further project the outcome of which could support or serve as an objection to the account of evidence-based justification developed here.

4.3.3 Comparison with other views

The position that evidence consists in internal states is not the same position as the sense-data theory criticized by Bouwsma, or the position about appearances developed by Chisholm. Let me try to distinguish the position developed here from these two positions.

In his essay “Moore’s Theory of Sense-Data,” Bouwsma (1942) analyzes and criticizes a theory of sense-data that he ascribes to Moore. The main idea of this sense-data theory is explained most clearly in one of the last sentences of the essay: “But I am not moved by the suggestion that whenever I look at my hand an image of my hand may be interposed between my hand and my eye.” (Bouwsma 1942, p. 20) The view that Bouwsma criticizes is thus the view that in perception we perceive only images, but not the physical thing. Furthermore, Bouwsma thinks of these images as similar to sounds or tastes, entities that exist and that can be described separately from their causes. In his own

example we can distinguish the sound that a bird makes from the bird. Visual sense-data are like sounds in that they exist and are supposed to be separate from the physical objects that causes them. He then shows that there is no such “image of my hand that may be interposed between my hand and my eye.”

Now the view developed in this dissertation does not claim that there is an image between the hand and the eye. Rather it claims that in the mind there is an image of the hand. The image is not an entity out there. It is an entity “in here,” meaning in the mind. The claim is thus that perception yields internal states, states of the mind, and it is through these internal states that we learn about external states. There is no claim that we are not in direct causal contact with hands and birds or other objects. Rather that we can learn about these external objects only through internal states that are hopefully caused by the external object.

It is revealing that in his criticism of sense-data Bouwsma talks a lot about how various linguistic phrases work, but among them is not “seems.” Now even in non-perceptual statements that involve “seems,” for example “it seemed to me to be a mathematical proof of the theorem,” the “seems” statements is about something in the mind rather than about something in the world. Suppose that you confidently present some mathematical reasoning as a proof. Now someone points out an error. You respond by saying that it seemed to you as though there was this mathematical proof. On a natural reading of this statement, you say that you thought that you had a proof, that you considered based on

some evidence that you had a proof. All of this is a statement about what went on in your mind. So a statement involving “seems” even in non-perceptual cases is a statement about mental states, one’s beliefs, thoughts, evidence. As I have tried to suggest, in a perceptual case a statement involving a “seems” refers to the sensory evidence you have, the sensory experiences understood as mental states.

Bouwsma criticizes the view that there is some kind of an image in the case of visual perception that is between the external object and the eye. The position explained here claims that there is an internal state in the mind through which we learn about external objects. In the acoustic perception, according to Bouwsma the sense-data would be the sound waves. However, according to the position developed here, there is an acoustic internal state that mediates our perception of the sound. The location of the sense-data is according to Bouwsma out there in the world. The internal state is according to the view developed here in the mind, not in the world.

In a selection, titled “Myth of the Given,” from his book “Philosophy,” Roderick Chisholm isolates a specific perceptual sense of the word “appears.” He says that when someone says “I am appeared to redly,” he states something about a perceptual experience he has, where there is no commitment as to whether there really exists something red in his physical surroundings. He states that this claim is self-justifying and can serve as part of the foundation on which all of our knowledge of the physical world depends.

There are some similarities between this idea of self-justifying appearance statements and the position developed here. The main similarity is that I also identify a class of statements, those that involve “seems,” as statements, not about the world but about our sensory experiences, where they are understood as internal states. However, I would hesitate to think of the statements that describe these sensory experiences as containing unfamiliar terms like “appeared to redly.” In order to render these appearance statements immune from any source of error, Chisholm furthermore suggest that when we say that we are appeared to redly, we do not mean to say that we are appeared to in a way similar to other occasions in the past or similar to what other people experience. So the idea seems to be that when I say that I am appeared to redly, I describe a sensory experience, however, without comparing it to anything else or claiming that it is similar to anything else. It is also not a hypothesis, so to speak, to the effect that there is something red in my immediate surroundings.

As will be obvious, the problem is that it is now not clear how this statement can serve as evidence for anything else. It would now be immune from error, but I could as well have said about the same sensory experience that I am appeared to bluely, or whitely. The point is that such a statement is no longer capable to serve as evidence for a claim like that there is a red apple on the table or something similar. It has been stripped of any content that would allow it to serve as evidence. It is probably immune from error, but just because it is devoid of content.

In contrast the position developed here concedes that one might be mistaken about whether a structure seems to be a cottage or rather a barn. Even when we talk about our sensory experiences we can make some mistakes in certain cases about them. The claim is however that we are usually able to thus see what our sensory experiences indicate as being true. We have the capacity to read our sensory experiences. This capacity can go wrong in some instances. What we gain in allowing for some possibility of error is a claim that can serve as evidence for beliefs about our immediate physical surroundings.

4.3.4 Distinguishing evidence-based justification from evidence

Now that we have seen what evidence is, we need to turn to evidence-based justification. Evidence-based justification requires evidence, but it also requires something else.¹⁶ Let me first distinguish evidence-based justification from evidence. Throughout this section I will mean evidence-based justification even if I just speak of justification.

An important difference between evidence and evidence-based justification consists in the following. One might have evidence for and against a given belief. One might have evidence of various types: sensory experience, beliefs through testimony, memory.

¹⁶ Often in epistemology the distinction between evidence and justification is made in the following way: Evidence is tied to the reasons someone can give for his beliefs. Justification is tied to whatever it is that renders the belief in question likely to be true. Justification can be external according to this distinction. Evidence will be something internal. Evidence would be related to giving reasons, while justification would be something that one “has” rather than “gives,” and that renders the belief in question likely to be true. This is not the distinction I have in mind. The notion of justification I develop is evidence-based justification, which means that it requires evidence. However, it also requires something more, which can prevent evidence from justifying a belief.

Evidence-based justification depends on an assessment of our total evidence. To see this, suppose that you are looking at an animal at a short distance and you have a sensory experience as of a sheep. Suppose furthermore that a local told you that there are in the vicinity a lot of sheep but also a lot of sheepdogs that look very much like sheep. So you have evidence both for and against the belief that there is a sheep in front of you. You have to weigh what evidence speaks for and what evidence speaks against a belief. If you consider the evidence that there is a sheep in front of you stronger, you will consider yourself justified in believing that there is a sheep. If you think that the evidence that renders it doubtful that there is a sheep is stronger, you will consider yourself justified in believing that it is either a sheep or a sheepdog, where you cannot tell which it is. You might of course also not consider yourself justified in any of these beliefs, because you do not think that conclusive evidence is available to you.

To use an analogy, evidence-based justification involves something similar to the verdict that a jury might reach after considering all the evidence both for and against that it is aware of. The cases for and against a given verdict as presented by the two opposing sides in the court are similar to your evidence for and against a given belief. Evidence-based justification results from an evaluation of your total evidence, just as a verdict results from the jury considering what both sides said. So by evidence I mean whatever speaks for or against a given belief. Evidence-based justification involves a pronouncement on the basis of the total evidence, whether this total evidence speaks for or against a given belief.

Let us see how evidence-based justification in the case of perceptual knowledge can be understood as resulting from our epistemic agency. Let us again consider the following case. You have sensory experience as of a sheep at a short distance, but a local told you that even though there are a lot of sheep, there also are a lot of sheepdogs running around in this area. In this example you have a sensory experience. You read it as a sensory experience as of a sheep. This is evidence that speaks in favor of their being a sheep in front of you. Then you have a further belief acquired through testimony. This belief is to the effect that there are both sheep and sheepdogs in this area. This is evidence that it might well not be a sheep that you are seeing. Now if we combine these two parts of the evidence, we should conclude that we should not read the sensory experience as one as of a sheep, but rather as one as of a sheep or a sheepdog. So this combination of evidence no longer supports the belief that there is a sheep. When we evaluate our total set of evidence to reach a verdict as to whether or not we are justified, we might come to reevaluate some evidence.

We can link this distinction between evidence and evidence-based justification more closely to the account of epistemic agency given in the previous section. A reading of a sensory experience will be evidence by itself, although if this reading is combined with other beliefs we might come to reconsider what it is evidence for. However, what a reading of a sensory experience by itself does is tell us for which belief this sensory experience is evidence. When we evaluate the reading, we will consider all of the

evidence available to us, and then, if we conclude that this evidence indicates the truth of a certain belief, this means that we regard this belief as justified. An evaluation of a reading of a sensory experience thus aims at moving from the evidence to evidence-based justification. When the evaluation is positive, we will regard ourselves justified in the belief on the basis of the sensory experience.

Let us consider a case of an illusion. Suppose that we immerse a stick partly in water. We will then have a sensory experience that we will read as a sensory experience as of a bent stick. This reading tells us what this sensory experience is evidence for, namely that the stick is bent. However, we have other evidence that it is not bent at all. When we move from this total set of evidence to evidence-based justification we will say that we are justified in believing that the stick is not bent. So the illusion arises because there is evidence that the stick is bent. However, when we look at our total set of evidence, we will consider ourselves justified in the belief that the stick is not bent. The reading of a sensory experience tells us what the sensory experience is evidence for. The evaluation of the reading of a sensory experience tells us whether we are justified on the basis of this evidence.

Let us now consider the role that defeaters play in whether or not we are justified in a given belief. We have already seen that when we evaluate our total set of evidence we will often take note of how some evidence is a defeater for some other evidence. Thus in the case in which we believe on the basis of testimony that there are a lot of sheep and

sheepdogs around, this defeats our evidence on the basis of the sensory experience as of a sheep. In the case of a rebutting defeater we will have evidence both for and against a given belief. In the case of an undercutting defeater, the defeating evidence will undermine other evidence we have for the belief. This is the case with the belief that there are both sheep and sheepdogs in this area. This belief changes our reading of the sensory experience as of a sheep to a reading of a sensory experience as of either a sheep or a sheepdog.

What we have said so far is not the only role that defeaters play in this evaluation that moves us from evidence to evidence-based justification. If we accept that our knowledge of the physical world is in general defeasible, that it is characterized by the problem of the gap, then this means that for any given set of evidence we might have, there might be further evidence that could change our verdict involved in evidence-based justification. Any given set of evidence we might have for a given belief about the physical world is such that it is possible that some further evidence could come up. This further evidence might be such that given the new set of evidence we would no longer regard ourselves justified in the belief that we took the old set of evidence to justify.

This raises the question of when we can reasonably make the decision to not look for defeaters or further evidence, but to stick with the set of evidence we have at the moment. If for any set of evidence there might be further evidence that would lead to a different verdict on evidence-based justification, then the decision to look for defeaters or not to

look for defeaters is an important element of our evidence-based justification. Let us explain this by looking at an example.

Suppose that someone has never heard of or experienced the bent stick illusion.

Furthermore suppose that this person has very restricted experience with objects that are immersed in water. Maybe this person was raised in the desert and never left that desert, and he has never experienced with immersing objects in water. Furthermore, this person has only a very hazy idea of physics. Suppose now that this person sees for the first time a stick that is partly immersed in water. For this person it will be a possibility that sticks are bent upon being partly immersed in water. However, we would want to say that for this person to come to a justified belief, he should look for defeaters to the belief that immersing a stick in water results into bending the stick. There is something to be investigated, defeaters to be uncovered. And if the person does not try to find these defeaters, to gain further evidence, we would not want to say that he has evidence-based justification that the stick is indeed bent on the basis of the sensory experience.

The idea is thus that evidence-based justification is not the result of evaluating any set of evidence. Similarly in the jury case, a just verdict is not the result of any set of evidence presented to the jury. So in the case of evidence-based justification if someone has a very limited or biased set of evidence then he is not justified on the basis of evaluating this evidence. Someone does not have evidence-based justification just because he happened to consider only very limited or biased evidence. To be justified requires the decision to

not look for further defeaters, and if this decision is not reasonable then the person in question does not have evidence-based justification.¹⁷ We will need to explain in some detail what it would mean for this decision to be reasonable. The term “reasonable” will be used in a technical sense.

In the preceding section I said that we would have to capture one of the three capacities that define our epistemic agency from the outside also, rather than merely as it presents itself to us from the inside. What we are trying to see now is that the third capacity, of evaluating our readings of sensory experiences needs to be captured partly by introducing something external. Someone might not be aware of the need to look for defeaters, and thus make the decision to not look for defeaters when this does not result into evidence-based justification. That the decision to not look for further defeaters is reasonable is not something that will be completely available to every epistemic agent from the inside. Again, what it means for this decision to be reasonable will be explained later on.

Summary Box 10: Distinguishing evidence from evidence-based justification

By evidence I mean those internal states that suggest that a given belief is true. In the case of perceptual knowledge our evidence consists in sensory experiences and beliefs.

¹⁷ This intuition is also contained in an example discussed in epistemology. Suppose that a newspaper prints a true report. However, for some reason the newspaper in the next edition says, falsely, that what it has earlier reported is not true. Suppose now that in a society everybody but one person reads the second edition. So they all have a defeater for what would be a true belief in the initial report in the first edition. However, there is one person who has not read the second edition and still believes the true report on the basis of reading the first edition. According to the account of evidence-based justification developed here this person who has not read the second edition would not be justified in the true belief because his

When I read a sensory experience I match it with those beliefs for which it is evidence. Evidence-based justification requires such evidence, but it is not identical to an evidence requirement. Evidence not only comes in different types and strengths, but we often also have both evidence for and against a given belief. Evidence-based justification depends on an assessment of our total evidence. It is the result of the third capacity, to evaluate our readings of sensory experiences. In moving from evidence to evidence-based justification we need to consider what our total set of evidence indicates.

However, evidence-based justification is not the result of evaluating any set of evidence.

If someone has a very limited or biased set of evidence then he cannot get evidence-based justification on the basis of evaluating such a flawed set of evidence. This holds even if the epistemic agent is not aware of the fact that his evidence is limited and biased.

Evidence-based justification requires the decision to not look for further defeaters, and if this decision is not reasonable – in a defined sense - then the person in question does not have this type of justification.

4.3.5 Different brain in vat situations

In this section I will explain in some detail what renders the decision to not look for a defeater reasonable or unreasonable. That is, I will explain the external constraints that have to be met in addition to an evidence-requirement for us to have evidence-based justification. The idea is thus that evidence, although necessary, is not by itself sufficient for evidence-based justification. There is no evidence-based justification without

evidence is limited and biased in a significant way. Strangely enough a more inclusive set of evidence

evidence, however, two epistemic agents could have the very same evidence and yet one might have while the other lacks evidence-based justification. Further external constraints have to be met for us to gain evidence-based justification. These external constraints make sure that the third capacity, of evaluating evidence, is free from deception and manipulation. In order to see this we will look at a variety of different brain in vat scenarios. The main question will be the one raised by the new evil demon problem, whether an envatted epistemic agent could have a justified false belief. The claim will be that evidence-based justification is not available to a brain in a vat.

In all of the following brain in vat situations I assume the following. The brain in vat, I will call it “envatted epistemic agent,” has sensory experiences that are not different from the inside in any way from those that normal epistemic agents have. In particular, the sensory experiences are input in the sense explained above. However, this input of the envatted epistemic agent is controlled by a computer or scientist, I will call whatever or whoever is in charge of the input “evil demon,” such that what input the envatted epistemic agent receives is determined by the evil demon. The question we will ask is whether and how such an envatted epistemic agent could gain a justified belief.

The first brain in vat scenario is characterized by the following. The evil demon causes input, the sensory experiences, in the envatted epistemic agent that are systematically misleading in that those beliefs that the envatted epistemic agent reads off of the sensory

would distract from the truth rather than lead to the truth. This will be discussed later on in more detail.

experiences are all false. Thus when it seems to the envatted epistemic agent that she is driving to the grocery store, she is in fact doing no such thing. Furthermore, the evil demon tampers with the ability to evaluate the readings of the sensory experiences so that the envatted epistemic agent will take what we would consider epistemic shortcuts. She will inductively generalize on very small samples without thinking about whether the sample could possibly be representative. She will believe or not believe everything someone says on the basis of a single successful or unsuccessful testimony. Sometimes she will believe something just because she wishes it to be true without seeking any evidence that it is true. In short, the demon causes the epistemic agent to respond to evidence in ways that seem highly problematic.

It seems to me that if the evil demon thus tampers with the ability to evaluate sensory experience or evidence more broadly then we have the intuition that the envatted epistemic agent is not justified in her beliefs. She does not search for evidence and evaluate evidence in the appropriate way, or so it seems. Thus she is not justified. Notice that this is through no fault of her own. It is the evil demon that causes the capacity to evaluate evidence to take these epistemic shortcuts. There is no sense in which the epistemic agent could avoid evaluating evidence in this way. In a sense she does the best she can, it is just that she is manipulated into taking these epistemic shortcuts.

Now consider the second brain in vat scenario. In this scenario again all the input, the sensory experiences that the envatted agent receives are systematically misleading in that

the beliefs that can be read off of the sensory experiences are all false. Again it might seem to the envatted epistemic agent that she is driving to the grocery store when she is in fact doing no such thing. Now the evil demon does not directly tamper with the capacity to evaluate the readings of the sensory experiences. The envatted epistemic agent can evaluate the readings of the sensory experiences in any way she wants.

However, the evil demon has created a perfect virtual world in which for some pet beliefs of the evil demon there is no shred of a defeater anywhere available to the envatted agent. In particular, the evil demon wants the envatted epistemic agent to believe that all swans are purple. No matter whether the envatted epistemic agent decides to accept her readings of sensory experiences as of purple swans, or whether she tries to investigate the matter as thoroughly as possible, looking for millions of different swans, reading books about swans, discussing swans with other people, there is nothing she could possibly do that would uncover a defeater for the belief that all swans are purple. The virtual world that the evil demon has designed is perfect in that respect in that there is nothing at all that the envatted agent could do that would give her evidence that swans are not purple.

As it seems to me, the two brain in vat scenarios amount to the same justification-wise. Through no fault of her own, the envatted epistemic agent cannot discover defeaters for her beliefs. In the first brain in vat situation the evil demon is directly manipulating her ability to evaluate the evidence. In the second brain in vat situation the evil demon is indirectly manipulating her ability to evaluate the evidence. In this second situation the envatted epistemic agent is in principle prevented from discovering defeaters for her false

belief that all swans are purple. No matter how careful she is, there is no way that she can somehow get evidence that swans are not purple.

The situation of the envatted epistemic agent in the second scenario is similar to the following: Suppose that we take a car and we suspend it from the road so that it sits on some blocks or is in some other way removed from the ground. Now we can start its engine, it might be running perfectly fine, the wheels will be spinning. However, it still is not driving. Similarly the envatted epistemic agent has her third capacity of evaluating evidence running perfectly fine, and spinning idly. However, the result is not justification because this third capacity is in principle prevented from uncovering a defeater to its false beliefs.

The envatted epistemic agent in the second scenario can only irrationally reject the false belief. All her evidence points towards the false belief, and no matter how careful she is she cannot uncover any evidence that would render it rational to not believe that all swans are purple. If she turns a skeptic and refuses to believe that all swans are purple this would be completely irrational. There is no evidence whatever to support such skepticism. So she is rationally compelled to believe what is false.

The point that I am trying to make is that justification does not just depend on evidence. An envatted epistemic agent can have the same set of evidence for a particular belief that I have. It also depends on whether or not the third capacity to evaluate evidence is free

from manipulation and deception. This manipulation can be direct, or indirect through a perfect virtual world as imagined in the second brain in a vat scenario. I here understand manipulation and deception to mean basically the same thing, as illustrated in the various brain in vat situations described here. We need to be able to freely evaluate our evidence. For our decision to seek out further evidence and defeaters to be reasonable, it seems that it must at least in principle be possible that we will discover such a defeater. Let us explain this somewhat more.

The first and the second brain in a vat scenario concern manipulation at different steps of forming a possibly justified belief. In step 2 we have evidence. Every individual piece of evidence indicates either that a given belief is true or false. In step 3.1, we aggregate the total set of evidence and see whether it indicates that a certain belief is true or false. In step 3.2, we decide to either accept or reject a given set of evidence and in consequence either form or refuse to form the belief in question. The first brain in a vat scenario manipulates the step 3.1. The envatted agent in that scenario will judge that a total set of evidence indicates that a certain belief is true when it does not, for example because the evidence is biased and limited. In the first brain in a vat scenario the brain will misjudge what a given set of evidence indicates and how strong the evidence is. In consequence, the step 3.2 will also be misplaced. The envatted agent will decide to accept evidence when she should not.

In the second brain in a vat scenario we have a manipulation that targets more directly the step 3.2, in that the conditions in which the envatted agent finds himself are not such so as to allow him to make the decision as to whether to accept some evidence in a reasonable way. One might object that this is not really a case of manipulation. After all, the envatted epistemic agent is free to make up his mind in any way he likes. However, whether we judge that there is more to be investigated, that we need to look for defeaters, will often depend on our experiences with similar cases in the past. We will often ask ourselves whether when judging about something similar in the past we later on encountered or were made to notice any defeaters. The judgment that there is more to investigate will often be prompted by past experiences. Now in the second brain in a vat scenario there is for the evil demon's pet beliefs no past experience that would suggest the need to look for a defeater. Furthermore there is no evidence at all that indicates the need to investigate further. So how can the envatted epistemic agent become conscious of a need to investigate further?

Let us consider an everyday case to support the claim that even if someone is not conscious of the need to investigate this further, he might lack justification because he has not investigated the matter further. Suppose that in an earlier century there is a medical doctor who also does some medical research. In particular, he forms the idea that a cold can be transmitted through touching something that was handled by someone already infected with a cold. This researcher conducts a number of experiments, and all of these confirm this conclusion that colds can spread through handling things that

someone already infected has handled. However, this medical doctor is very skeptical, and so even after confirming this hypothesis in thousands and thousands of cases, after discussing it with other researchers, who have got the same results in similar experiments, after doing everything in his power to confirm this hypothesis, and when there is no shred of a defeater in sight anywhere, this medical researcher refuses to believe what his experiments indicate.

Would we not want to say that there is something exceedingly strange about this medical researcher and that he should believe what he has such good evidence for? However, the same is true for the envatted epistemic agent in the second brain in a vat scenario. His evidence shows as conclusively as possible that all swans are purple. However, in this case his evidence supports a false belief. Except for this, his situation and the situation of the medical researcher are the same. Only on pain of being irrational, of ignoring what one has good evidence to believe, can the envatted epistemic agent refuse to believe that all swans are purple.

The manipulation of the step 3.2 can thus also be achieved just through manipulating the input. If there is nothing in the input that would allow the epistemic agent to be skeptical and to search for defeaters, he must accept what he has such good evidence for. If the input is thus manipulative in this sense, the epistemic agent cannot make the decision whether or not accept some evidence in a reasonable way.

Strangely enough, on the present account the lack of doing something, doing which would render the epistemic agent irrational, can be unreasonable. It should be clear that the terms “rational” and “reasonable” are used here in a technical sense. I use the term “irrational” as an internal notion, where whether it is rational or irrational to search for defeaters will largely depend on whether there is any reason at all to search for defeaters. In the second brain in a vat scenario, there is no possible evidence that would point towards the need to investigate whether all swans are purple. In such a scenario it would be irrational to investigate the matter further. However, the term “unreasonable” is understood to contain an external constraint. If there in fact is a defeater easily available then not searching for this defeater is in this defined sense unreasonable.

One might still object that the new evil demon problem widely generates the intuition that even a brain in a vat can act in more or less epistemically responsible ways, and that it might sometimes achieve justification. As it seems to me whether or not the brain in a vat can be justified depends on the details of the case. Later on we will discuss a fifth brain in vat scenario in which the envatted epistemic agent can have a justified false belief. What I am trying to say for the first through fourth brain in vat scenario is that the epistemic agent is indirectly or directly manipulated to accept false beliefs on the basis of misleading evidence. He cannot fully exercise his epistemic agency, or make the decision to accept some defeasible evidence in a reasonable way. These points are important for evidence-based justification. Again, evidence-based justification is just one notion of justification. It is perfectly legitimate to define different types of justification. In this

chapter I am seeking to explicate the type of justification that is required for perceptual knowledge. This leaves open the possibility to say that a brain in a vat might have some other type of justification.

A different notion of justification might tie justification to epistemic responsibility. One might say that the envatted epistemic agent can still act in more or less responsible ways and that this should be reflected in our attribution of justification to him. One might say that epistemic responsibility centers on step 3.1, seeing which belief is indicated as true by the total set of evidence. The epistemic agent gets a certain input and then judges on the basis of the total input as to what it indicates to be true. One might say that this judgment can be made in more or less responsible ways. If there is no direct manipulation and if the agent nevertheless takes epistemic shortcuts, he does not make this step 3.1 in the appropriate way. So we can say that an agent does the best he can given his input when he sees correctly which beliefs are indicated as true by this total input, and which are not. One might define a type of justification that is linked to this judgment.

The problem is that this is not sufficient for the kind of justification that is necessary for perceptual knowledge. There is also some intuitive support for the idea that justification is not the result of evaluating any set of evidence. Evidence-based justification introduces further external constraints on this evaluation of a set of evidence.

We can preserve the intuition behind the new evil demon problem by saying that a responsible brain in a vat would perform the step 3.1, of seeing what the total set of evidence indicates, correctly. An irresponsible brain in a vat does not, for whatever reason, judge correctly what his evidence indicates. Basically, a responsible epistemic agent would do whatever his circumstances allow towards gaining justification. An irresponsible epistemic agent would fail to act in ways, which would allow him to gain justification. However, no matter how responsible an epistemic agent is in performing this step 3.1, it is not sufficient for evidence-based justification. This is what we have to see more clearly now.

4.3.6 Evidence-based justification

In this section I will, finally, define evidence-based justification. To do so, suppose that we have a set of evidence for a certain belief. Now we can decide to accept the evidence or to seek further evidence and possible defeaters. Suppose that we decide to seek further evidence. If the belief for which we have the evidence is false then either we are capable of uncovering a defeater for this false belief, or we are incapable of uncovering such a defeater. If all possible evidence we could gain would support the false belief, then we are in the second brain in vat situation. In the second brain in a vat situation the decision to seek further evidence and the decision to accept a given set of evidence amount to the same thing because defeaters for the false beliefs are in principle beyond our epistemic grasp.

It is easier to say when we do not have evidence-based justification given a set of evidence for a belief. If our decision to not seek further evidence is a matter of manipulation or deception, then we do not have evidence-based justification for this belief. Now, one might wonder how it is that one thus can ever be justified in a false belief. Suppose that we have a false belief for which we have reasonably strong evidence. How could this belief be justified? Would we not always have to seek out defeaters, and if we fail to do so would we not be unreasonable in accepting the evidence for the false belief? How could the decision to accept evidence for a false belief be ever reasonable?

If we accept that our evidence is defeasible then this means that it is never justified to completely close our mind and to respond to any evidence we come across that speaks against a given belief with denial. Sometimes in the discussion of skepticism it is claimed that knowledge allows us to do just that: to close our minds, to ignore any further evidence.¹⁸ However, there is a difference between ignoring further evidence and deciding not to seek out further evidence. The former is never reasonable given that our knowledge is defeasible. The latter, however, is a part of our justification. If we consider ourselves justified in a belief, we decide not to go and seek further evidence or to look for defeaters. We now need to see when it might be reasonable to make the decision to not look for defeaters.

¹⁸ Compare Unger (1974). Unger claims that knowledge entitles one to absolute certainty and to close one's mind completely with respect to the belief in question.

In the case of the person without any experience of the bent stick illusion as described above, it seems to make a difference whether a defeater looms, so to speak, on the surface, or whether it would require a great amount of epistemic work to uncover it. In the bent stick case all the person would have to do is to get a tactile sensation of the stick to get a defeater to the claim that it is bent. Sometimes defeaters can be discovered easily, while in other cases they are very difficult to uncover. It seems that if a defeater is easily available then to decide not to seek further defeaters is not reasonable. Again, reasonable is here understood in a technical sense.

What we have said so far has at best some intuitive plausibility. What we would need is a criterion according to which we can reasonably decide to not seek further evidence and defeaters. If we mean by a criterion an internally available standard that would always lead to a reasonable decision to not seek further evidence, then there is no such thing: that was the point of the second brain in a vat situation. Sometimes there are circumstances that do not allow us to make this decision in a reasonable way, although this might not be through a fault in the epistemic agent. What is internally available is just the recognition that defeasible knowledge requires particular care in when to accept evidence and when to seek further evidence and defeaters. We have to do our best so that we do not form beliefs on the basis of evidence that a little bit of reflection would uncover to be biased and limited.

Again suppose that we have a set of evidence for a false belief. Two things need to be satisfied to lead to evidence-based justification for the belief. First, it must be in principle possible to find a defeater for this false belief, even if doing so would require a very great amount of epistemic work. Second, there is no defeater available that a modest amount of epistemic work would uncover. The second condition requires that some care and work go into the decision to not seek further evidence. So if someone has a biased and very limited set of evidence, he will not have evidence-based justification for this belief because there are defeaters available that could easily be uncovered. This condition is external in the sense that even an agent who has no internally available reason to seek out a defeater would not satisfy this condition if in fact there is a defeater easily available to him. The first condition also requires something external, that the epistemic agent is capable of uncovering defeaters for his false beliefs. This first condition means also that the agent is “epistemically connected” to his physical environment in that he could come across defeaters for his false beliefs.

Now, it is obvious that these two conditions need to be refined further. For it might seem that they could be easily satisfied in a third brain in vat situation. Suppose that the envatted epistemic agent is in a situation similar to the second brain in a vat situation except for the following. If he were to walk into his neighbor’s garden he would come across some green swans, thus coming across a defeater to the claim that all swans are purple. Would this not mean that the envatted epistemic agent can in principle uncover a defeater, and furthermore that uncovering this defeater would not be a very difficult thing

to do? One can further modify this situation in rendering it very difficult to come across these green swans. Maybe they are on a different continent and hidden in a rainforest so that it is very difficult to get to these green swans. Would we not now have a situation in which both conditions are satisfied? However, if we think that the envatted epistemic agent in the second brain in a vat situation does not have evidence-based justification, how can making such minor changes lead to this type of justification?

It seems clear that even in this third brain in a vat situation the envatted epistemic agent is manipulated into evaluating her evidence in a certain way. Furthermore, after encountering the defeater, the envatted epistemic agent now believes that most swans are purple but some are green, which is not really an improvement from the beliefs she had before encountering the defeater. The first condition is thus to be understood so that it is in principle possible to have evidence for a true belief about the matter at hand, for example whether all swans are purple. Even in the third brain in a vat situation the envatted agent has false beliefs, and it seems that she is in principle incapable of gaining evidence for a true belief on the issue in question. Defeaters are not as such important, but in so far as they are capable to improve our epistemic situation and to lead us to the truth.

However, this proposal is also inadequate, which can be seen by means of the following fourth brain in a vat situation. Suppose that the envatted epistemic agent is in a situation very similar to the second brain in a vat situation except that the evil demon randomly

selects one belief out of every 100 beliefs and makes this belief true. In this case the envatted epistemic agent has some true beliefs. However, their truth does not reflect any merits of the evidence. It is just that one in 100 beliefs is randomly selected to be true. If we have the intuition that in the second brain in a vat scenario the envatted epistemic agent does not have evidence-based justification, how can he have such justification in this fourth brain in a vat situation?

It is clear that we would not want to say that the envatted epistemic agent has the type of justification required for perceptual knowledge. It seems that in addition to having a true belief, we also need evidence that is at least sometimes linked to the truth of the belief for which we have the evidence. It should not be a matter of randomness that a belief for which we have evidence is true. So we should say that an epistemic agent is reasonable in accepting some evidence for a belief if at least sometimes his evidence is linked to the true belief by means of a luck-precluding condition.

There is one last, the fifth brain in a vat situation we need to consider. Suppose again that the envatted epistemic agent is in a situation similar to the second brain in a vat situation except that he also sometimes has evidence for a true belief that is linked to the truth of the belief by the luck-precluding condition. However, some of the details of the luck-precluding condition would be very surprising to the epistemic agent. The evil demon has in this case some pet beliefs that are true and that he wants the envatted epistemic agent to believe. Among them is that no swan is purple. So the evil demon gives the envatted

epistemic agent lots of evidence that no swan is purple and all of this evidence is linked by a luck-precluding condition to the truth that no swan is purple. However, the evil demon has also a false pet belief, namely that all swans are green. And this evidence that all swans are green is in no way linked through a luck-precluding condition to the truth, because it is not true that all swans are green. Could the epistemic agent be justified in the false belief that all swans are green? And would the true belief that no swan is purple be knowledge?

We need to distinguish between two problems to which this fifth brain in a vat scenario gives rise. First, there is the problem that the details of the luck-precluding condition are unexpected. Second, there is the problem of suffering from something like an inverted hallucination case. Let us begin by examining this second problem first.

We usually assume that most of our sensory experiences arise in the course of perception and only some in the course of hallucinations. But what if things are the other way around? What if most of our sensory experiences are hallucinations and only some arise in the course of perception? This would be a situation where we have some true beliefs where the evidence is linked to the truth by means of a luck-precluding condition, but most of our beliefs would be false. I am inclined to think that in this situation we would have a lot of justified false beliefs. We would have only very little knowledge, but most of our beliefs would be justified and false. The claim that we are not justified in the second brain in a vat scenario depends on the idea that in that scenario we are somehow

completely removed from the truth about the physical world in which we live. So the problem was that we could never come across a defeater for our false beliefs, and that we are thus not sufficiently epistemically connected to what is true in the physical world in which we live.

Now, in this fifth scenario the envatted epistemic agent is sufficiently epistemically connected to the truth of the physical world to have some true beliefs about it, the evidence of which is connected to the truth by a luck-precluding condition. One can probably render this fifth scenario a problematic borderline case by supposing that the envatted epistemic agent has only one true belief plus luck-precluding condition while all other beliefs are just false. There is some vagueness as to whether in a given case the epistemic agent is sufficiently connected to the truth or whether he is completely removed from it. However, this vagueness seems to be unavoidable. Barring such borderline cases it seems to me that it is possible to have a little bit of knowledge and a lot of justified false beliefs in the fifth brain in a vat scenario.

Let us now consider the first problem, that the details of the luck-precluding condition are surprising. In order to see the problem let us suppose the following: Suppose that it becomes possible to conserve the brain by itself in a vat and to put this vat in a robot that then walks and talks and sees things and engages in actions. There now would be many details of the luck-precluding condition that would be surprising. The brain in the robot might have no idea that all of its sensory signals come through cameras and other

receptors that are artificially constructed. However, it is not clear that knowledge of the details of the luck-precluding condition is required. After all, most of us do not know the details of how sensory experience works in biological people. And it does not seem that one would have to know such details before being able to know that there is barn in the field or that there is a red triangle painted onto a white wall. Thus it does not seem that surprising details, as for example sensory signals that are mediated by a computer or programmed by a scientist, would by themselves disqualify the epistemic agent for knowledge.

Let us now define evidence-based justification for a false belief. An epistemic agent has evidence-based justification for a false belief only if the following two conditions are met. First, the epistemic agent must in principle be capable of coming across a defeater that then would allow her to have a true belief plus luck-precluding condition about the issue in question. Second, there is no defeater that could easily be uncovered by the epistemic agent. These two external conditions define when it is reasonable to accept a given set of evidence for a false belief. According to the second condition the epistemic agent has been careful and has made a good effort to gain evidence and to come across defeaters. According to the first condition the epistemic agent is in principle “epistemically connected” to the world and is capable of arriving at a true belief plus luck-precluding condition on the basis of defeaters.

What these two conditions do is they render it a matter of “localized” bad luck if we arrive at a lot of evidence for a false belief. According to the first condition we would in principle be capable of arriving at a true belief plus luck-precluding condition about the matter at hand. So there is no “systematic” bad luck due to pervasive manipulation and deception. According to the second condition we are not in a position where we have evidence for a false belief because of a lack of reflection and epistemic care. So it is a matter of genuine bad luck rather than due to something that is in our control.

We can approach this matter also from a different angle, focusing on the idea that we somehow need to be “epistemically connected” to the physical world. This “epistemic connection” can be understood as a luck-precluding condition. So the idea is that one has evidence-based justification for a belief only if one can in principle arrive at evidence that can be linked through a luck-precluding condition to the truth. So the evidence that we need is evidence that is part of the luck-precluding condition. This is what we will explain further in the next section.

Let us now give the general definition of evidence-based justification. An epistemic agent has evidence-based justification for a belief if and only if he reasonably decides that the total set of evidence available to him supports this belief. The decision that a total set of evidence supports a belief is reasonable if and only if the set of evidence available to the epistemic agent indicates that the belief is true, and the following is true of this epistemic agent given his circumstances. The epistemic agent is in principle able to arrive in his

circumstances at a true belief plus luck-precluding condition about this matter, and there is no defeater that could be easily discovered by him.

One might have this type of justification both for a true and a false belief. One might have this type of justification both when one has evidence that is linked by a luck-precluding condition to the truth that this evidence indicates and when there is no such luck-precluding condition. This type of justification depends on the assumption that it is an objective matter whether a given set of evidence indicates that a given belief is true, and what would count as a defeater to this evidence and belief.

This is an account of evidence-based justification in the sense that evidence, understood as internal states, are necessary for this justification. There is, however, a further external condition, that the decision to accept the evidence has to be reasonable, which is not something that is completely internally available to the epistemic agent in question. This external condition on evidence-based justification can be made precise by means of two constraints, as given above. The epistemic agent needs to be in principle capable of arriving at a true belief plus luck-precluding condition about the matter, and there cannot be a defeater easily available to the epistemic agent.

The envatted epistemic agent in the first scenario is manipulated into violating the condition to exercise care and look for defeaters. The envatted epistemic agent in the second, third, and fourth scenario is manipulated into violating the condition that she

must be in principle capable of arriving at a true belief plus luck-precluding condition about the matter at hand. In none of these scenarios is the envatted agent somehow to blame for accepting the evidence she accepts. Given what is internally available to the epistemic agent, she does the best she can. However, the external circumstances do not allow her to gain evidence-based justification. It is only in this external sense that the envatted agent does not make a reasonable decision to accept the evidence she does accept.

Let us now revisit the new evil demon problem. The question is whether an envatted epistemic agent could possibly be justified in a false belief. If we understand justification to mean evidence-based justification in the sense defined above, we can give the following answer. We have looked at four different situations in which the envatted epistemic agent would not have evidence-based justification for her false beliefs. However, we have also looked at the fifth situation, in which it seems that the envatted epistemic agent would have evidence-based justification for her false beliefs. In the case of evidence-based justification, our response to the new evil demon problem depends on the details of the deception in which the envatted epistemic agent is entangled. Again, it is possible to define other notions of justification, and given such other notions one might have to give a different response to the various brain in vat scenarios.

Summary Box 11: The definition of evidence-based justification

An epistemic agent has evidence-based justification for a belief if and only if he reasonably decides that the total set of evidence available to him supports this belief. The decision that a total set of evidence supports a belief is reasonable if and only if the set of evidence available to the epistemic agent indicates that the belief is true, and the following is true of this epistemic agent given his circumstances. The epistemic agent is in principle able to arrive in his circumstances at a true belief plus luck-precluding condition about this matter, and there is no defeater that could be easily discovered by him.

One might have this type of justification both for a true and a false belief. One might have this type of justification both when one has evidence that is linked by a luck-precluding condition to the truth that this evidence indicates and when there is no such luck-precluding condition. This type of justification depends on the assumption that it is an objective matter whether a given set of evidence indicates that a given belief is true, and what would count as a defeater to this evidence and belief.

There is some work left unfinished in this chapter, and in the dissertation. This concerns the notion of evidence, what exactly evidence is and how we arrive at it, for example by reading a sensory experience. This also concerns the idea that a total set of evidence can be seen to support a specific belief rather than an alternative belief in an objective way, and that it is an objective matter what would count as a defeater. This is something that I will not discuss in this dissertation. The discussion in this chapter and in the dissertation

targets in detail only what the external constraints are on deciding to accept a given set of defeasible evidence so that this decision will be reasonable in a defined sense.

We have now defined and explained a notion of evidence-based justification. What remains to be done in this chapter is to see how this type of justification is perfectly tailored to the luck-precluding conditions. In the next chapter I will give an argument that this type of justification is necessary for perceptual knowledge.

4.4 How evidence-based justification is ideally tailored to the luck-precluding condition

In this section I will suggest that evidence-based justification is ideally designed to match and fit with the main other necessary condition for knowledge, the luck-precluding condition. I will explain that this type of justification, evidence-based justification, is the ideal mate for the luck-precluding condition. In the next chapter I will give an argument that because of this match between the two main requirements for perceptual knowledge no counterexample that is based on epistemic luck can be given to the account of perceptual knowledge developed in my dissertation.

In preceding chapters I have argued that the luck-precluding condition is necessary for perceptual knowledge. It is also assumed in those chapters that evidence is necessary for perceptual knowledge. I will give arguments based on epistemic luck in the next chapter for each of the conditions that show them to be necessary for perceptual knowledge. In

this section the goal is more limited. I merely wish to elaborate that evidence-based justification is the type of justification that ideally complements the luck-precluding condition.

To begin our discussion let us consider a simple case of perceptual knowledge, for example the perceptual knowledge that a red triangle is painted on a white wall. Suppose that we have an epistemic agent. Let us consider what needs to be the case for this person to gain perceptual knowledge that there is a red triangle painted on a white wall. First the person needs sensory experiences that can be read as sensory experiences as of a red triangle on a white wall. The epistemic agent will need further evidence, that the lighting is standard and that the triangle is not obviously projected on the wall rather than painted on it. Suppose that the epistemic agent gains this further evidence, and judges that all of his evidence now indicates that there is a red triangle painted on the wall. So, the epistemic agent accepts what his evidence indicates as true. For this to be reasonable, in the defined sense, two further conditions have to be satisfied. The epistemic agent has to be in principle capable of arriving at a true belief plus luck-precluding condition about this matter, and there is no obvious defeater to his belief and evidence that he could easily discover. Let us suppose that both these conditions are met.

So now we have an epistemic agent who has evidence-based justification for the belief that there is a red triangle painted on the wall. Now let us look at what the luck-precluding condition does. The luck-precluding condition basically tells us that if the

belief that would result from the second step is true, so is the belief resulting from the third step. The second step without the third step would tell us that it seems as if there is a red triangle painted on the wall. This is a statement about the evidence to the effect that the evidence indicates that there is a red triangle painted on the wall. The third step results in the belief that there is a red triangle painted on the wall. The idea is that all that the luck-precluding condition does is link what renders the second step true to what renders the third step true. The luck-precluding condition in this case consists in a regularity that links what renders it true that it seems that there is a red triangle painted on the wall and what renders it true that there is a red triangle painted on the wall. The sensory experience as of a red triangle supported by the other evidence is linked to the truth of the belief for which it is evidence.

Let us explain this again to get completely clear about this. We have an epistemic agent who has a variety of evidence to the effect that there is a red triangle painted on a white wall. This evidence consists in sensory experiences and supporting beliefs. If we say that it seems as though there is a red triangle painted on a wall, this is a statement about our evidence, that our evidence indicates that there is a red triangle painted on the wall. On the basis of this evidence the epistemic agent comes to believe that there is a red triangle painted on the wall. What the luck-precluding condition does is it links the evidence to the truth of what it is evidence for. The luck-precluding condition consists in a regularity between what renders it true that it seems as though there is a red triangle painted on a wall and what renders it true that there is a red triangle painted on the wall.

The luck-precluding conditions consist in contingent regularities in the physical world. In this case we have foremost the regularity of the red triangle painted on the wall causing in the epistemic agent a sensory experience as of a red triangle painted on the wall. If this regularity is present then there is the luck-precluding condition that links the evidence to the true belief.

We can now also see in what way a luck-precluding condition has to be expected or assumed to obtain. When we move from step two to step three, from the evidence to the belief that is suggested as true by the evidence, we expect there to be a luck-precluding condition. A luck-precluding condition is thus expected in the sense that we have to assume it to be present in order to move from step two to step three. Evidence-based justification and luck-precluding condition are related in the following way. When we move from our evidence to the belief for which we have the evidence, and thus from what would be a “seemings” claim about our evidence to a belief about the physical world, we assume that the luck-precluding condition is present that links the evidence to the true belief for which we have the evidence.

We can expect there to be a luck-precluding condition when there is none. We can even reasonably expect there to be a luck-precluding condition when there is none. The expectation of a luck-precluding condition is an assumption rather than somehow a premise in an argument or some claim that we need to support. As we have explained, the

luck-precluding condition is the way the world needs to be so that through the evidence we will arrive at a true belief. In moving from the evidence to the belief that the evidence indicates, we assume the luck-precluding condition to be present in the sense that we assume that the world is in fact such that the evidence indicates the truth. There is both an internal and an external dimension to the reasonable expectation of a luck-precluding condition.

There is a further way in which evidence-based justification and luck-precluding conditions are related to one another. As we have seen, an important part of evidence-based justification is a reasonable decision to accept some evidence as telling us the truth. The idea behind the two conditions that define when it is reasonable, in the defined, external sense, to accept evidence is that these conditions explain when it is reasonable, in the defined, external sense, to expect that the luck-precluding condition obtains. If there is a defeater easily available, then it is not reasonable to expect that there is a luck-precluding condition linking the evidence without the defeater to the truth. Furthermore, if an epistemic agent is in principle unable to arrive at a true belief plus luck-precluding condition, then it is not reasonable to expect that there is such a luck-precluding condition. Again, these are external constraints on when it is reasonable to expect that there is a luck-precluding condition.

Let us again sketch the picture of perceptual knowledge we are now painting. There are a lot of contingent regularities in the world. Furthermore, we have evidence through

sensory experiences and supporting beliefs for beliefs about the physical world. When we accept the defeasible evidence, when we decide not to look for further defeaters but to believe what our evidence indicates, we assume that there is a regularity that links our evidence to the truth of what the evidence indicates. We assume that our seeming to see a red triangle is linked through a regularity to the truth that there really is a red triangle. So in moving from evidence to the belief that the evidence indicates as true, we assume that a luck-precluding condition is present. This assumption will be reasonable only if there is no easily accessible defeater available and only if we in principle could have a true belief plus luck-precluding condition about this matter. A defeater would suggest that the luck-precluding condition is not present. And if we are in principle incapable of arriving at a true belief plus luck-precluding condition about this matter then we are in principle not reasonable in the assumption that there is a luck-precluding condition.

Let us work through a more complicated example in order to illustrate this account somewhat further. Again an epistemic agent is driving around and arrives at the belief that there is a barn in a particular field at a short distance. We assume that the epistemic agent has a sensory experience that he reads as indicating that there is a barn in front of him. When he moves from this evidence to the belief that there is a barn in front of him, one of the things he does is he decides to accept the evidence he has and not to look for defeaters. The epistemic agent will have evidence-based justification for the belief that there is this barn only if he is reasonable in making this decision to accept the evidence. He is reasonable in doing so only if he has taken sufficient care and he is in principle

capable of arriving at a true belief plus luck-precluding condition about this matter. This true belief that there is a barn will be perceptual knowledge if and only if the epistemic agent has evidence-based justification for it and the luck-precluding condition that is assumed when moving from evidence to the belief about the barn obtains in this case. In the example the luck-precluding condition he assumes consists in the regularity that there is a barn causing in him the sensory experiences as of a barn.

Let us consider some of the cases in which the epistemic agent would fail to have evidence-based justification. Suppose first that he is driving through barn façade county. It seems that in this case a defeater would be easily available to him, and so unless he has investigated the matter further he does not have evidence-based justification for his belief. It seems also that in this case the epistemic agent has no perceptual knowledge even if one of his barn beliefs should turn out to be true. Suppose second that the epistemic agent is a brain in a vat that is tricked into believing only false things in particular about barns. Since he is now unable to arrive at a true belief plus luck-precluding condition about this matter his belief again will not have evidence-based justification. Furthermore, there is intuitive support for the claim that even if one of his beliefs by chance turns out to be true, it would not be knowledge.

Suppose third that his sensory experience as of a barn is caused by a barn picture that is placed in front of the barn and that from the distance looks exactly like the barn behind

the picture.¹⁹ This raises the question of whether we need to know the details of the luck-precluding condition. It seems that sometimes these details would furnish us with defeaters, while in other cases they would support the assumption that the luck-precluding condition obtains. Suppose that perception always involves, not a causal contact with the object perceived but with a projected image of the object that is placed slightly to the left and in front of the object we are perceiving. Would this be perception? Would we know that there is a barn in front of us? It seems that we could have perceptual knowledge, even though if we were to discover that perception always involves these other projected images, we might count that as a defeater as to whether there really is a barn, until we gain good evidence that every projected image comes with an appropriate physical object.

The problem is that in requiring too many details of the luck-precluding condition to be expected or assumed to obtain, we would require perfect scientific knowledge of the physical world before we can ever gain perceptual knowledge. But that seems highly counterintuitive. I can know many things through perception even if I have no idea as to the causal mechanisms at work in perception. Even those who lived in a pre-scientific time could know that there is a barn in the field or that there is a red triangle painted on a white wall. Perceptual knowledge does not depend on a scientific theory of the causal mechanism underlying it.

¹⁹ A similar example was suggested by Keith Lehrer (2006), and also Justin Fisher.

Nevertheless we need to assume that our evidence is telling us the truth, and that means assuming that a luck-precluding condition obtains. Even if it does obtain there might be details about it that we would consider a defeater, until further investigation shows that there indeed is a luck-precluding condition even though it has some features that are surprising. The idea is thus that we need evidence-based justification that involves an expectation of a luck-precluding condition even if not all the details of the luck-precluding condition might be expected and even if some of them would be considered a defeater.

We should clarify that the evidence-based justification involves an expectation of the luck-precluding condition. However, as the fifth brain in vat scenario shows one might be justified in a false belief. This would also mean that one might be justified in a belief for which there is no luck-precluding condition. So expectation of a luck-precluding condition does not mean that the luck-precluding condition indeed obtains. The expectation is part of the evidence-based justification, not that the luck-precluding condition indeed obtains.

As I have tried to show, the notion of evidence-based justification contains some external requirements. There are external constraints on this evidence-based justification. Nevertheless, speaking of evidence-based justification is appropriate because evidence understood as internal states is required for this type of justification. Furthermore, the external constraints are constraints that limit when the evidence leads to justification. The

external constraints cannot do the job of the evidence or somehow in the absence of the evidence lead to justification. Rather these external constraints limit the cases in which evidence will give us evidence-based justification. Whenever someone has evidence-based justification, he has evidence for the belief that he is justified in. It is, however, not the case that whenever someone has a certain set of evidence for a belief, then he also has evidence-based justification for this belief. It is evidence that leads to evidence-based justification. It is just that there are further external constraints when some evidence can lead to this type of justification. There is thus no external “mechanism” that can in the absence of evidence lead to this type of justification. There is no reliability or probability constraint. There is no constraint here that would be familiar from currently discussed versions of externalism. It seems that this can render it appropriate to call this account of justification “evidence-based justification,” even though there are two external constraints on it.

Summary Box 12: Evidence-based justification and luck-precluding conditions

There are two ways in which evidence-based justification is the perfect mate for the luck-precluding conditions. The first concerns the move from the second step to the third step in the three steps of our epistemic agency. The second step without the third step would result into a belief that, for example, it seems as though there is a barn in the field. This is a statement about our evidence, claiming that this evidence indicates that there is a barn in the field. The third step would be, for example, the belief that there is a barn in the field. This is a statement about the physical world. In moving from the second to the third

step we move from evidence to the belief that this evidence indicates as true. The luck-precluding conditions consist in those regularities that underwrite this move. In the example, the luck-precluding condition links by means of a physical regularity what renders it true that there seems to be a barn to what renders it true that there is a barn. The second concerns the claim that the type of justification required for perceptual knowledge needs to contain an expectation or assumption that the luck-precluding condition obtains. Furthermore, this assumption needs to be reasonable. The two conditions on when it is not reasonable, in the defined sense, to accept a set of evidence as telling us the truth, also explain when it is not reasonable to expect that the luck-precluding condition is present. If a defeater is easily available, then it is not reasonable to expect that the evidence that one has (without the defeater) is linked by a luck-precluding condition to the truth that this evidence indicates. If an epistemic agent is in principle not able to arrive at a true belief plus luck-precluding condition about a given matter, then it is not reasonable to expect that the luck-precluding condition obtains.

4.5 The problem and the proposal revisited

The main questions that I have tried to answer in this chapter are what kind of justification is necessary for perceptual knowledge, and whether a brain in a vat could have this type of justification. The chapter gives an account of evidence-based justification that is based on an understanding of the nature of our epistemic agency. Evidence-based justification contains an evidence requirement and two external constraints. It thus combines an internal and an external component, where, however, the

external component merely restricts when the internal component leads to this type of justification.

In this chapter I have also explained how this type of justification is ideally tailored to the luck-precluding conditions, which have been shown to be a necessary condition for perceptual knowledge in the previous chapter. In previous chapters I have also suggested that the type of justification required for perceptual knowledge should contain both an evidence requirement and make sense of the idea that we somehow expect or assume that the luck-precluding condition obtains. Both of these ideas have been incorporated into the definition of evidence-based justification.

As for the new evil demon problem, I have not really given a general answer to it. Rather I have suggested that evidence-based justification is not available to certain types of situations that involve deception and manipulation. There is not just one notion of justification. Therefore, there is not just one answer to the new evil demon problem. Furthermore, the details of the deception and manipulation matter as to whether a certain type of justification is attainable by a brain in a vat.

In this and preceding chapters I have suggested the following account of perceptual knowledge. A true belief is perceptual knowledge if and only if the epistemic agent has evidence-based justification for it and the luck-precluding condition obtains. The requirement that the luck-precluding condition is expected or assumed to obtain is

contained within the requirement of evidence-based justification. In the next chapter I will give for each of these conditions an argument that it is necessary for perceptual knowledge. Then I will give an argument that they are jointly sufficient for perceptual knowledge. This argument will depend on the idea that the evidence-based justification and the luck-precluding condition are not independent requirements but are intertwined in various ways, which is what we have looked at in the preceding section.

5. EPISTEMIC AGENCY AND SKEPTICISM

5.1 The problem and the proposal: A set of necessary and sufficient conditions for perceptual knowledge and a reply to some skeptical worries

Consider again Lucky from the fourth chapter, who has all the sensory experiences and beliefs of a normal person leading a normal life in a normal world. In fact, Lucky has and always had for her entire life the same sensory experiences and beliefs I have. Lucky and I are completely indistinguishable from the inside. However, Lucky is a brain in a vat and incapable of gaining any knowledge of the physical world. The problem now is whether and how it might be that I have perceptual knowledge even though Lucky is incapable of such knowledge.

If Lucky and I are indistinguishable from the inside, then whatever enables me to have perceptual knowledge must be something external, something that is not available from the inside. The problem now is that from the inside whether or not these external conditions obtain seems to be a matter of epistemic luck. But how can it be that knowledge depends on such epistemic luck? How can it be that of two people with identical sensory experiences and beliefs, one of them can have perceptual knowledge while the other is incapable of such knowledge? Should knowledge not reflect the strength of my justification, and be an achievement due to my evidence rather than a mere matter of epistemic luck?

This problem raises the main question we will ask in this chapter. How exactly need internal and external conditions be combined so as to allow us to gain knowledge rather than end up with a true belief as a mere matter of epistemic luck? The question is thus how external and internal requirements for knowledge need to be blended together, to complement one another, to be tailored to each other, to allow knowledge rather than a true belief that is a mere matter of epistemic luck. This is what we will look into in this chapter.

The proposal is that evidence-based justification contains an expectation of the luck-precluding condition, and the luck-precluding condition, if it obtains, underwrites our adopting the belief in question on the basis of our evidence. Evidence-based justification contains an internal element, and the luck-precluding condition is an external requirement. It is the match between these internal and external requirements that will allow me to argue that the proposed conditions for perceptual knowledge are indeed sufficient.

According to the account proposed in the dissertation, a true belief is an instance of perceptual knowledge if and only if the epistemic agent has evidence-based justification for his belief and the luck-precluding conditions obtain. In the first section I will provide arguments to the effect that each of the conditions for perceptual knowledge developed in the preceding chapters is individually necessary for perceptual knowledge. In section 5.3 I will argue that those necessary conditions are together sufficient for perceptual

knowledge. In section 5.4 I will show that we are capable of such perceptual knowledge and that thus skepticism, in so far as it claims that knowledge is impossible, is false.

5.2 Necessary conditions for perceptual knowledge

In the previous chapters I have suggested that a true belief is perceptual knowledge if and only if the epistemic agent has evidence-based justification for the belief and the luck-precluding conditions obtain. In this section I will give arguments for each of these necessary conditions. We will begin by looking briefly at true belief. Then I will give arguments for the other necessary conditions, drawing in these arguments on epistemic luck. The basic idea for all of these arguments is that if one of the necessary conditions is absent, then the belief is true as a matter of a certain type of epistemic luck. Furthermore, this type of epistemic luck can be shown to be incompatible with knowledge. It follows that the proposed condition is indeed necessary for perceptual knowledge.

The claim is thus that each of these conditions is necessary because they each eliminate a kind of epistemic luck that is obviously incompatible with knowledge. We see again the importance of the notion of epistemic luck. It is epistemic luck that can give us arguments that the proposed conditions are indeed necessary for perceptual knowledge. Epistemic luck is, as should be clear, a crucial notion of defining knowledge. Again we also see that there are more than one form of epistemic luck, and that it has to be understood in its various varieties. Furthermore, particular attention has to be paid to the question of whether a specific type of epistemic luck is incompatible with knowledge.

5.2.1 True belief

The first necessary condition for knowledge is true belief. There is among epistemologists disagreement about what account to give of the belief- and truth-requirement for knowledge. For example, Keith Lehrer argues that knowledge requires acceptance rather than belief, where acceptance seems to be a specific kind of belief, one that reflects the strength of our reasons.²⁰ One might thus disagree about what exactly is involved in the belief-requirement for knowledge. Similarly there can be disagreement about what account of truth to give. For example, pragmatists about truth claim that truth does not reflect mind-independent facts but rather what is ideally rational to believe.²¹ There is thus a substantive issue as to what account to give of the truth- and belief-requirements for knowledge.

In this subsection I will not discuss what account to adopt of the belief- and truth-requirements, because this can best be clarified and argued for by looking at the other necessary conditions. The argument for the luck-precluding condition will clarify and provide an argument for an account of the truth-requirement for knowledge. The arguments for evidence and evidence-based justification will allow us to see in what sense belief is required for knowledge. The key necessary conditions for knowledge are evidence-based justification and the luck-precluding condition. In arguing that these are necessary conditions, we will also see how we need to understand the belief- and the

²⁰ Compare Keith Lehrer (2000). In particular Chapter 2: Truth and Acceptance.

truth-requirement. For now I will just assume that knowledge requires a true belief and bracket the question of what account to give of the belief- and truth-requirement for knowledge. However, we will return to this question at a later point in this chapter.

5.2.2 Evidence

In this section I will argue that evidence is a necessary condition for knowledge. We will see that a true belief for which we have no evidence cannot be knowledge. The argument for this claim involves the idea if we do not have strong evidence for a belief, then its being true is a matter of the type of epistemic luck present in certain cases of a lucky guess. So I will try to capture the type of epistemic luck that is present when we have a true belief for which we lack strong evidence. I will claim that this type of epistemic luck is present in certain cases of a lucky guess. Furthermore, I will give an argument that this type of luck, as involved in a guess made in the absence of strong evidence, is incompatible with knowledge.

My use of the term “lucky guess” is in part technical, but also in part continuous with everyday intuitions and a tradition in the literature that goes back to Plato. Let me begin by discussing how the notion of a lucky guess is continuous with everyday intuitions. In order to do so, suppose that someone is driving to a grocery store to which he has never been before. Suppose furthermore that he has only a very hazy sense as to where this grocery store is located. Now he is approaching an intersection, and his guess is that he

²¹ Compare Roderick Chisholm (1966), who attributes this position to William James.

needs to make a right turn so as to shortly reach his destination. His guess turns out to be correct.

Different things might go on when someone makes such a lucky guess, in this and other cases. The person making the lucky guess might once have had knowledge of the location of the grocery store even though he cannot consciously remember what he once knew. So it might be that someone who makes a lucky guess taps into some lost knowledge in some inexplicable way. Alternatively, it might be that the person making the lucky guess responds to very weak evidence that would not by itself license belief. It might be that she exaggerates the strength of some evidence, or is aware of the limits of this evidence but somehow acts on an intuition somewhat supported by the weak evidence.

Alternatively, the person making the lucky guess might act for non-epistemic reasons. Such non-epistemic reasons that might prompt belief include habits, recognition of sentimental value, superstition, and random procedures like a coin toss. It seems that all of these are familiar and reasonable explanations of a lucky guess.

It seems that in each of these cases the lucky guess that results into a true belief is characterized by the absence of strong evidence. If someone somehow taps into some lost knowledge, then, since the knowledge is lost, evidence in the sense of consciously available beliefs or sensory experiences that support the lucky guess are not available. If some evidence, some knowledge is not remembered, then it is not available as evidence to the person making the lucky guess. If someone acts on clearly insufficient evidence by

somehow relying on an intuition, it seems again clear that the person making the lucky guess arrives at a true belief without strong evidence for it. Whatever little evidence the person might have cannot rationally license belief, and so, again, if the person arrives at a true belief it is not because of strong evidence available to him. Finally, if someone makes a lucky guess for non-epistemic reasons it is clear that he does not believe because of strong epistemic evidence.

It seems to me that these are all familiar situations. When we make a lucky guess, we will often make it in this way. Furthermore, I take it that intuitively we would support the idea that such a lucky guess is a guess and lucky precisely because we do not have strong evidence for it. It is the absence of strong epistemic evidence that seems intuitively to characterize such a lucky guess.

There are some objections that we need to consider. One might say that a guess is not something that results into a confident belief. Furthermore, one might have a confident belief even in the absence of strong evidence for it. According to this objection, a guess is characterized by its feeling like a guess, which means that the person in question lacks confidence; he lacks what it would take to have a confident belief in the result of his guess.

In response, there is some truth to this in that something can feel like a guess, and something might not feel like a guess even though the person has no strong evidence for

the belief. There probably are different kinds of guesses, only some of which will feel like a guess. In the grocery store example, it seems to me that we would intuitively say that this is a guess even if the person for some reason feels very confident in making that turn. It does not seem to me that a lack of confidence, its feeling like a guess, is the defining characteristic of a guess. There seem to be forms of guessing that would result into a strong and confident belief in the absence of good evidence for this belief.

I am thus more interested in those lucky guesses where we somehow have a confident belief, but where we cannot point to any strong evidence for this belief. In those cases there is nothing consciously available to us that would count as strong evidence for this belief. This includes also cases where one might be mistaken about whether something counts as evidence for a belief, and how strong such evidence might be. It seems to me that there are still everyday cases that would have these features and be familiar situations. The case where one does not consciously remember the location of the grocery store, but one once knew where the grocery store is located, and is because of that confident in making the right turn would be one such instance.

Furthermore, in the literature that goes back to Plato, a lucky guess is usually understood as a situation in which we lack strong evidence for the result of this guess, and nevertheless believe it and act on it. Often in this literature we find the idea that a true belief is not sufficient for knowledge because it might be the result of a lucky guess.²² So

²² This idea goes back to Plato in the *Meno*. Compare Peter Klein (1971).

the notion of a lucky guess is linked by this tradition to a true belief for which we lack evidence.

The notion of a lucky guess as discussed in this subsection might thus not be fully continuous with everyday intuitions, in particular in that it does not demand that something should somehow feel like a guess. However, there are clear cases that we would normally count as lucky guesses and which are all characterized by the absence of strong evidence. Furthermore, if something feels like a guess, we would also expect the absence of strong evidence – why else should something feel like a guess?

One might object that one cannot draw on an intuition that a lucky guess is not knowledge unless one understands a lucky guess in the usual way. If we define a lucky guess in an unusual way, surely our intuitions as to whether a lucky guess could be knowledge should not be appealed to. Fortunately, I will not appeal to such an intuition, but give an argument that a lucky guess as understood here, characterized by the absence of strong epistemic evidence, is an instance of a form of epistemic luck that is incompatible with knowledge.

What I have so far suggested is that a certain type of a lucky guess is characterized by the absence of strong epistemic evidence for the belief in question. Now, this is supposed to be an argument that evidence, understood as internal states, is necessary for perceptual knowledge. An externalist will of course disagree with this claim. An externalist will

deny that we need any evidence for our true beliefs for them to possibly count as perceptual knowledge. We thus need to address this externalist challenge to the ideas presented in the dissertation.

Externalists about justification and knowledge claim that a belief might be justified and might be knowledge even if it is arrived at in the absence of any evidence for it at all. For example, reliabilism claims that a belief is justified and might be knowledge if it is arrived at through a reliable cognitive process, even in the absence of any evidence for the belief at all.²³ What I want to show is that if we arrive at a true belief in the absence of any evidence for it, the type of epistemic luck is present that characterizes a lucky guess. However, it must immediately be said that I do not mean to say that the all true beliefs arrived at due to a reliable cognitive process would be lucky guesses. That would be much too strong a claim. The idea is rather that if we lack evidence for a belief, then if the belief is true this is a matter of the same type of epistemic luck that would also characterize the guess made in the example of whether to turn right to get to the grocery store – even if the belief is caused by a reliable cognitive process.

Many examples that externalists give for their position actually involve a form of evidence for the belief in question. This would be evidence that the epistemic agent has a certain reliable cognitive process. However, the externalist I imagine objecting to the ideas presented here has to claim that even in the absence of any evidence at all one

²³ Compare Alvin Goldman (1976)

might gain perceptual knowledge due to a reliable cognitive process. This is the type of position I will examine now.

Let me illustrate that often externalists draw on a certain type of evidence, even though they claim that it is the reliable cognitive process that is doing all the heavy lifting. An example that is often used in the externalist literature is the following. There is a need to sort chicks into male and female chicks. Now someone can learn how to thus sort the chicks. However, this learning does not involve consciously recognizing any indicators of the sex of the chick. The person learning this skill is not told to look or feel or listen for something specific that then would allow him to tell apart which chicks are male and which are female. Instead the person is just given feedback by someone who is experienced in this matter, and after some time he will master this skill without seeking out consciously any evidence that would allow him to do this job.²⁴

As it seems to me, only evidence of a certain kind is absent in this case. It is true that the person who sorts the chicks does not sort them on the basis of a consciously available sensory experience or belief that would indicate that the chick is male. There is no evidence that speaks differentially for the belief that the chick is male. However, the person acquiring the skill is given feedback, and after a while he will have evidence that he possesses a certain reliable capacity. He has evidence that his judgments that a certain chick is male are arrived at through a reliable capacity, and thus he has evidence that

supports that whatever his judgment is, if it is made on the basis of this reliable capacity it is likely to be correct. So this is not a case in which no evidence at all is available to the person in question. The person has evidence that he has a certain reliable cognitive capacity, and this evidence can support individual judgments about the sex of a chick.

Let us thus construct a case in which even this evidence, that the person has a certain reliable cognitive capacity, is absent. Suppose that a person has a certain reliable cognitive capacity, namely the perfect pitch. He has no training in music, but one of his childhood toys included a little flute, and so he knows the different tones that this flute can play. Let us furthermore suppose that he gave them names, some combinations of numbers. Now suppose that this person, maybe at a young age, hears for the first time a piano played. He immediately recognizes some of the tones that the piano plays and he judges them to be the same as the ones he played on his little flute. Let us consider a specific judgment, namely that a specific tone is a “12.” Let us stipulate that he has no evidence at all for this claim. In particular, he has no evidence that would indicate that he has the perfect pitch. We assume that he just spontaneously forms the judgment that this specific tone is a “12,” which is his name for a specific tone played on his toy flute.²⁵

²⁴ This example is discussed extensively in the epistemological literature, including Duncan Pritchard (2005),

²⁵ Depending on how perfect pitch works there might in actual cases be some evidence available to the person. The sounds he hears might be acoustically different according to whether they are a “12” or some other tone. It might be that someone with perfect pitch just hears the tones differently than someone without a perfect pitch. This would be similar to someone who is colorblind and someone who has color vision: there are differences in what they see. However, for the purposes of this discussion I will assume that there is no consciously available acoustic difference that would count as evidence for saying that a tone is a “12.”

I grant that this will not feel to the person with the perfect pitch like a guess. He might be strongly convinced that this tone that he is now hearing is a “12.” I also grant that this person has a capacity that in principle would allow him to know a certain range of truths, including what tone is now being played on the piano. However, the problem is that in the absence of any evidence that he has this reliable capacity, he is epistemically in the same situation as anyone in whose mind suddenly a belief pops.

Suppose that the person with the perfect pitch has a little brother. Suppose that the little brother is also raised without musical training, that he also plays on the little flute, giving the tones the same names. Suppose the two brothers hear together the tone played on the piano. Both brothers form spontaneously the belief that the tone is a “12.” However, the little brother does not have the perfect pitch. However, the belief in question is just as subjectively convincing to the little brother. What distinguishes the two epistemically? What is that would allow one of them to rely on this spontaneous belief, to have knowledge, while the other one should not rely on this spontaneous belief, and he would not have knowledge?

Let us consider a different example to answer these questions. Suppose someone is very superstitious and he thinks that the number 19 and the color red are extremely unlucky. Suppose that he plays roulette on a regular basis and he never bets on 19 or on red. His betting behavior is based on this superstition, and he has no evidence whatever for the belief that this number and this color are unlucky. Could it possibly matter whether he is

right or wrong, in one instance, or on the whole, for us to say that whatever happens will be a lucky or unlucky guess? Even if most of his bets that are based on this superstition turn out to be successful, it still does not seem that we would want to attribute this success to anything but lucky guesses. Merely getting it right, even systematically, does not by itself allow us to gain knowledge.

The problem for the externalist is to say how the situation of the person with the perfect pitch is different from the situation of his little brother and the superstitious roulette player. There is no evidence available to either of them that would allow us to differentiate between them. There is no difference in subjective conviction. We can even stipulate that there is no difference over a certain period of time in their getting it right. What then is the epistemically significant difference between them? Why should the person with the perfect pitch rely on his judgment when his little brother and the superstitious roulette player should not? It is clear that the little brother and the superstitious roulette player are merely guessing, and that they are merely lucky if they get it right. It seems to me that in the absence of any evidence that one has a certain reliable cognitive capacity, one should not rely on one's beliefs produced by this reliable cognitive process. In the absence of any evidence that one has this reliable cognitive process, one's beliefs produced by this process are from an epistemic point of view similar to the lucky guesses I have described. Furthermore, it is the absence of any such evidence that renders them similar. If we investigate the person with the perfect pitch and his little brother, finding out that one of them has the perfect pitch, while the little brother

lacks it, we will have acquired evidence for the reliable capacity. And given this evidence, we can gain knowledge on the basis of this reliable capacity. However, in the absence of any evidence that one has a certain reliable cognitive capacity, it seems to me that from an epistemic point of view there is no difference between the person with the perfect pitch in the example and his little brother and the superstitious roulette player.

I am not sure that this will be very convincing to an externalist. However, the challenge would be to say exactly why the person with the perfect pitch should rely on his judgment about the sound, while his little brother should not, given that they are both right and have no evidence of any relevant difference between them.

The main contribution that I see the externalist making is twofold. First, he introduces an important kind of evidence – strangely enough – namely evidence that one has a certain reliable capacity, a capacity that is likely to get certain things right. This type of evidence is different from, for example, a sensory experience as of a table, in that it does not speak differentially for a specific belief, but rather generally for a certain type of beliefs. This is a very important kind of evidence. Second, many of our beliefs are probably due to intuitions, hunches or gut feelings, where we lack strong evidence. The externalist seems to say that such intuitions, hunches or gut feelings could in fact be the result of reliable cognitive processes that do not rise to the surface of consciousness. Evidence, understood as consciously available internal states, is only one way of forming beliefs. Many of our beliefs probably result from a variety of processes that do not lead us by way of evidence.

Given that many of our decisions are made in partial ignorance of their consequences, given that we depend on many beliefs which cannot possibly be knowledge, often because they involve the future, this is an important point to note.

I take it that we have now seen that in the absence of evidence the type of luck is present that characterizes certain kinds of lucky guesses. Now we need to argue that if a true belief is a lucky guess in the sense that it is arrived at without strong evidence for it, then it is not knowledge. So we need to argue that this type of epistemic luck is incompatible with knowledge.

Suppose that someone makes the lucky guess that he needs to turn right at this intersection so as to shortly reach the grocery store. We have seen that in that case this true belief is arrived at without strong evidence. This means that the person could rationally have chosen a different guess without its being the case that anything in his epistemic situation would have had to change. In the example, the person might have on the basis of the same evidence, in the same epistemic situation rationally made the guess that he needs to turn left to reach the grocery store. Nothing in his epistemic situation would need to change, in particular he would not have to discover new evidence or defeaters, for it to be rational to make this alternative guess.

Now, this is very different from knowledge. If a true belief is knowledge then something in the epistemic situation would have to change for it to be rational to believe something

different. Further evidence would have to be uncovered, a defeater would have to be discovered for it to be rational to change a belief that amounts to knowledge.

We might also say that from an epistemic point of view a belief as a result of a guess is at best optional and possible, but never mandated by the evidence. Usually we guess because a practical situation requires a decision or choice or belief in the absence of proper evidence. A belief as a result of a guess is never one that is required or even advisable given the evidence the person in question has. Rather it is an option, a possibility given the evidence the person has. Now it seems that a belief that does not come as somehow required or advisable given the evidence is not knowledge. A belief as a result of a guess is optional. The person in question might given the same epistemic circumstances have believed something different. This is not a situation of knowledge. If someone knows something, then this belief that is knowledge is not optional. It just is not the case that a person who knows something might as well have believed something different. Again, we are comparing cases in which nothing about the epistemic situation changes. Let us get more precise.

A belief as a result of a guess is such that the following is true of it. Nothing needs to change about the epistemic situation of the person making the guess, in particular nothing needs to change about the evidence, for it to be rational to make a different guess. In the case of a guess we can have the same epistemic circumstances, but a different belief as a result of a guess. A belief that is knowledge is such that the following is true: It is not

rational to believe something different rather than the belief that constitutes knowledge if nothing in the epistemic situation of the knower changes. In the case of knowledge we cannot have the same epistemic circumstances, but a different belief than the one that is knowledge.

The argument given in this section can be stated as follows. Consider a belief for which we lack strong evidence. The type of epistemic luck present if the belief is true is the same that is present in certain types of lucky guesses. If a belief is the result of a guess, again in the sense that it is arrived at without strong evidence, then the belief is not knowledge. It follows that if a belief is arrived at in the absence of strong evidence then it is not knowledge. It follows that a belief is knowledge only if it is arrived at through strong evidence. Thus strong evidence is a necessary condition for knowledge. It should be clear that I do not really spell out in my dissertation what would count as strong evidence.

Again, the argument just given does not depend on a questionable intuition about whether a lucky guess could possibly be knowledge. The argument makes use only of the definition of the type of lucky guess I am interested in, one that results into a true belief in the absence of strong evidence. There is some overlap between what we would in everyday life intuitively call a lucky guess and a true belief for which we lack evidence, which is the reason for my using the term “lucky guess” here. I do not mean to try to argue for something controversial, that knowledge requires evidence, through something

fairly uncontroversial, that a lucky guess is not knowledge. Nothing in the argument just given depends on the intuition that a lucky guess is not knowledge.

Summary Box 13: Evidence is necessary for perceptual knowledge

By evidence I mean those internal states that support a belief. If such evidence is absent, then the belief is optional in the sense that it would be rational to adopt a different belief even if nothing about our epistemic circumstances changes. A belief that is knowledge is never optional in this sense. It just is not the case that if nothing changes in our epistemic circumstances, it would still be rational to change a belief that amounts to knowledge. Thus evidence is necessary for perceptual knowledge.

5.2.3 Luck-precluding conditions

We now need to see that in addition to a true belief for which we have evidence, a luck-precluding condition is necessary for knowledge. In the preceding subsection we compare knowledge to a lucky guess, and this allows us to argue that evidence is required for knowledge. In this subsection we will compare knowledge to Gettier cases. This will allow us to present an argument to the effect that the luck-precluding condition is required for knowledge. Again we will begin by trying to understand what characterizes a Gettier case, and then argue that a Gettier situation is incompatible with knowledge.

The following is an example of a Gettier case that targets perceptual knowledge. Suppose that you are looking out on a field and you arrive at the belief that there is a sheep in the

field. Your evidence consists in sensory experiences as of a sheep. However, this sensory experience is caused by a sheepdog, not a sheep. In the far distance, however, there is a sheep in the field even though it is too far away to be recognizable by you. It seems that you have a true belief for which you have sensory evidence. Nevertheless it seems clear that you do not have knowledge that there is a sheep in the field.²⁶

In the third chapter I have already presented the position that Gettier cases are characterized by the absence of a luck-precluding condition. Let us now try to get clearer about this. A luck-precluding condition consists in a regularity between evidence and the true belief for which we have the evidence. In the example such a luck-precluding condition could be stated as a regularity between your sensory experiences as of a sheep and there really being a sheep in the field. Now, if a luck-precluding condition is absent then this means that the evidence and the true belief are mismatched; they are not properly paired with one another. This is what we need to explain in detail now.

In the example of a Gettier case just given it might seem that there is no such mismatch and that the luck-precluding condition is present. We have evidence in form of a sensory experience as of a sheep. And the true belief for which we have the evidence is that there is a sheep in the field. How could this possibly be a mismatch? It seems that in the realm of perception there could not be a better match. The evidence directly indicates the true belief. So how can one possibly claim that there is a mismatch here?

²⁶ This example is given in Roderick Chisholm (1966).

In the case of perceptual knowledge, evidence is the result of reading a sensory experience. We have a certain sensory experience that we then read, for example, as a sensory experience as of a sheep in the field. Now, when we read a sensory experience we associate the experience with the belief for which it is evidence. In the example we take a specific visual sensory experience and we read it as suggesting that it is true that there is a sheep in the field. In thus reading a sensory experience we need to abstract away from many of the details of the sensory experience and the situation in which we are. This can best be seen by reflecting on the variety of situations that could obtain that would cause a sensory experience as of a sheep in the field: We might indeed be looking at a sheep. We might be looking at a sheepdog. We might be looking at a projection of a sheep. We might be looking at sheep made from clay. And so on and so forth. When we read a sensory experience as an experience as of a sheep, many different things might be going on in the physical world that all could result into this sensory experience.

Something similar goes on in the case of a true belief. The belief that there is a sheep in the field can be true due to many different factual situations: There might be a sheep in the middle of the field. There might be one at the far end of it that had its wool dyed red. There might be a sheep that is chased by a sheepdog and running from one end of the field to the other end. And so on and so forth. As it should be clear both the sensory experience and the true belief can be made true by a variety of different scenarios in the physical world.

Now, why are we talking about this? The reason for talking about this is that evidence as a reading of a sensory experience and a true belief can be perfectly matched if we just consider them as propositions, but completely mismatched if we consider what it is in the physical world that makes them true. In the example of a Gettier case we can express the reading of the sensory experience as “I have a sensory experience as of a sheep,” and the true belief as “there is a sheep in the field.” Thus expressed in propositions they seem to be matched. However, let us now consider what it is that renders these propositions true. The sensory experience as of a sheep is rendered true in virtue of a sheepdog causing a sensory experience as of a sheep. The true belief that there is a sheep in the field is rendered true in virtue of an unrecognizable sheep in the far distance. Thus the truth-makers for these two propositions are completely mismatched. There is no connection between a sheepdog that causes in you a sensory experience as of a sheep and there being an unrecognizable sheep in the distance.

What does this have to do with the luck-precluding condition? What we have seen is that the propositions summarizing the evidence and the true belief can be made true by a wide array of different situations in the physical world. This means that sometimes even though the evidence and the true belief are not mismatched as propositions, there is no connection between the truth-maker for the evidence and the truth-maker for the belief. If that is the case then the evidence and the true belief are mismatched. And that allows a Gettier intuition to arise.

Now, this mismatch arises precisely because of the absence of a luck-precluding condition, and the presence of a luck-precluding condition eliminates the possibility of such a mismatch. In the third chapter we saw that a luck-precluding condition depends on a regularity in the physical world between what renders the evidence true and what renders the belief for which we have the evidence true. Let us illustrate this by looking at a normal case. Suppose that you are looking out on a field and you have a sensory experience as of a sheep that is caused by a sheep. In that case there is a link between what renders your evidence true and what renders the belief for which we have the evidence true. The evidence is rendered true by a sheep causing the sensory experience as of a sheep. The belief for which we have the evidence is rendered true by the sheep that causes the sensory experience as of a sheep. We have a perfect match here, not just on the level of propositions but also on the level of the facts in the physical world that render the propositions true. The luck-precluding condition would consist in the regularity of sheep causing sensory experiences as of a sheep. This regularity would be present both in general and in this particular instance.

Now let us consider again the Gettier case. Here there is no regularity between what renders your evidence true and what renders the belief for which you have the evidence true. There is no regularity that links unrecognized sheep in the distance to sheepdogs that are mistaken for sheep. Even though the evidence is matched with the true belief on

the level of propositions, they are mismatched on the level of the facts in the physical world that render these propositions true.

We have thus suggested that a Gettier case is characterized by the absence of the luck-precluding condition. More precisely, we have shown that a Gettier case, as illustrated by the sheep example, involves the absence of a luck-precluding condition. Furthermore it seems that if a luck-precluding condition is absent this generates a Gettier intuition. More details and more example of this characterization of the Gettier cases can be found in the third chapter. With this reference to further discussion in the third chapter, I will for now assume that if someone is in a Gettier case, there is no luck-precluding condition, and if there is no luck-precluding condition then we have a Gettier case.

If we can also show that Gettier situations are incompatible with knowledge we will have argued for the luck-precluding condition as a necessary condition for knowledge. What that would show is that knowledge requires that the evidence and the true belief for which we have the evidence need to be linked by a regularity on the level of the facts that render the evidence and the belief true. Such knowledge requires that there is not merely support for a true belief, but that what renders the evidence true and what renders the belief true need to be tightly connected through a regularity in the physical world.

If we now want to argue that a Gettier situation is incompatible with knowledge, we need to reflect on how the internal and the external conditions are combined in a Gettier case.

In a Gettier case we have evidence for a true belief. We have here a combination of an internal and an external condition. There is evidence, understood as internal states, for example sensory experiences and beliefs. Then there is truth, understood as an external condition, one that reflects what occurs in the physical world. Now we have already seen that in a Gettier case a luck-precluding condition is absent. This means that the evidence and the true belief are mismatched. This in turn means that there is no regularity in the physical world that links the truth-maker of the evidence to the truth-maker of the belief for which we have the evidence.

The internal and the external conditions are thus not connected with one another. They are combined without being in any way related to each other. It is a little bit like making a cake without stirring and blending the ingredients. There is no connection between them, they just happen to sit in the same bowl. Similarly in the case of a Gettier case the internal and the external condition just happen to occur together, rather than as a result of a tight connection between them.

Now it seems that this is at the core of what justifies the intuition that a Gettier situation is incompatible with knowledge. If there is no tight connection between the external and the internal conditions for knowledge then it is a matter of coincidence or a fluke that they indeed obtain in the Gettier situation. This kind of epistemic luck, one where there is no tight connection in form of a luck-precluding condition between the internal and the external requirements for knowledge, is incompatible with knowledge.

Generally, epistemologists share the intuition that there is no knowledge in a Gettier case. What I have said right now explains why that should be the case. The internal and the external conditions need to be connected with one another, rather than that they just happen to occur together in an epistemic situation. This seems to me to be the main lesson to learn from this analysis of the Gettier situations. It is not really an argument that would convince someone who denies that a Gettier situation is incompatible with knowledge. However it is support that there are important reasons as to why people have the Gettier intuition.

The argument I have attempted to give in this subsection goes as follows: If we are in a Gettier situation then we do not have knowledge. We are in a Gettier situation if and only if we have evidence for a true belief but the luck-precluding condition is absent. It follows that if we have evidence for a true belief but the luck-precluding condition is absent, then we do not have knowledge. Thus we have knowledge only if the luck-precluding condition is present. Thus the luck-precluding condition is a necessary condition for knowledge.

Summary Box 14: The luck-precluding condition is necessary for perceptual knowledge

A Gettier situation is characterized by a mismatch of the evidence and the true belief that we adopt on the basis of this evidence. This mismatch occurs on the basis of the facts that

render the evidence and the belief true. In the example, there is a mismatch between the sheepdog responsible for your sensory experience as of a sheep, and the unrecognizable sheep in the far distance that renders your belief true. A luck-precluding condition consists in a regularity that links what renders your evidence and the belief true. In the example such a regularity might be a sheep causing in you the sensory experience as of a sheep. This regularity, and thus the luck-precluding condition is absent in the Gettier case. Furthermore, in the absence of a tight match between evidence and true belief there is no knowledge. The internal and the external requirements for knowledge need to be matched. This suggests that luck-precluding conditions are necessary for perceptual knowledge.

5.2.4 Evidence-based justification

In this subsection we will see what the evidence-based justification requirement adds to a true belief for which we have evidence and for which the luck-precluding condition obtains. The requirement of evidence-based justification over and above evidence eliminates a further element of epistemic luck and contributes to blending together the internal and the external conditions for knowledge.

Consider the luck-precluding condition. It links evidence to true belief through a physical regularity in the world. We now need to think which evidence has to be linked to a true belief through the luck-precluding condition. So far we have assumed in this chapter that it is just some evidence that is part of the luck-precluding condition. However it is

possible to have evidence both for and against a given belief. We also might have evidence of different kinds and different strengths. If any such evidence can fill the slot in the luck-precluding condition, then this means that the epistemic situation is characterized by a kind of luck that is incompatible with knowledge. Let us consider an example to see this.

Suppose that you live in Austria close to the border to Germany. Suppose now that for several days it is reported in the newspapers and on TV and on the radio that a bear has come to live in this very area where you live and that he is attacking all kinds of animals and people.²⁷ You have heard from friends who have seen the bear and witnessed his attacks. So we suppose that you have extremely good evidence for the existence of this bear in this very area in which you live. Now suppose that one day as you are taking your daily walk in the evening, you suddenly have sensory experiences as of a bear running towards you. However, the color of the bear's coat seems to you to be somewhat strange. On the basis of this single bit of evidence, you say to yourself that this bear that seems to be running towards you is not real, that either you are hallucinating or someone or something is playing a trick on you. Instead of running away or trying to ward off any attack, you just calmly stay where you are, telling yourself that this bear is not real. Just as the bear is about to attack you, as it seems to you, the hallucination vanishes. As it turns out the bear was not real.

Now in this instance you have a true belief that the bear is not real. You have some evidence for this belief, the funny looking coat of the bear. We can now assume that indeed there is a luck-precluding condition in that there is a regularity between sensory experiences as of such a funny looking bear and its being a hallucination. So we have a true belief for which we have a little bit of evidence, where this evidence is linked to the true belief by means of a luck-precluding condition. But it surely would not be knowledge, because you ignored all the excellent evidence you have for the existence and reality of the bear. It is a matter of too much luck that there is such a luck-precluding condition between this evidence and the true belief.

This is why evidence-based justification over and above evidence is required for knowledge. Evidence-based justification depends on a judgment as to what your total set of evidence indicates. And this, the evidence that is supported by such a judgment about what the total set of evidence indicates, enters into the first slot of the luck-precluding condition.

Let us try to see more clearly what the evidence-based justification requirement adds to the picture. We often have a potpourri of evidence, some speaking for a given belief, some speaking against it. In moving from such evidence to evidence-based justification, we form a general assessment as to what the total set of evidence indicates. For perceptual knowledge this total assessment would have to support a sensory experience

²⁷ This example is loosely based on an actual story in the German news about a bear that lived in that area

as indicating the truth. So it is not false to say that the first slot is filled by evidence, it is just that this evidence is selected through a general assessment of the total set of your evidence. In moving from evidence to evidence-based justification we will, so to speak, take sides with one of the various elements of evidence that we have, and we will single it out as indicating the truth. Knowledge is possible only if this evidence that is supported by the general assessment is linked by the luck-precluding condition to the true belief for which we have this evidence.

If we think again about the bear example, it seems that the reason why there is no knowledge in this case is that it is just lucky that despite the very strong evidence for the existence and reality of the bear it turns out to be a hallucination. There is luck involved that this little bit of evidence on which the belief that the bear is not real is based turns out to be linked to the truth. The person about to be attacked by the bear is just very lucky that the bit of evidence on which his belief is based is part of a luck-precluding condition.

We might also say that even though the person encountering the bear is expecting the luck-precluding condition, he is not doing so reasonably. In the preceding chapters we have often said that the luck-precluding condition as an external condition needs to be expected through your justification. What the bear example shows is that this expectation cannot be whimsical. To have knowledge it cannot be that someone just happens to expect there to be a luck-precluding condition even though he has overwhelming

evidence that would defeat this expectation. So we can see that this expectation of the luck-precluding condition has to be somehow reasonable.

As we have explained in the fourth chapter, the expectation of a luck-precluding condition is also not reasonable in either of the following two cases. The person in question is in general not able to gain a true belief plus luck-precluding condition about this matter. This means that there is never evidence that is linked through a luck-precluding condition to the truth, as for example in some brain in vat situations.

Alternatively, the person could easily become aware of defeating information of which she is currently not aware. This means that she is missing crucial and easily available elements in her total set of evidence.

The bear example already gives us the idea that the expectation of a luck-precluding condition needs to be reasonable and cannot be due to just ignoring some crucial evidence. Now we can modify the bear example to support these two further constraints on the reasonableness of the expectation of the luck-precluding condition.

Suppose again that you live in this very area of Austria in which according to newspapers, radio and TV and the personal experience of friends of yours a bear has come to live. However, in the place in which you live you are the only person who has not heard of the existence of this bear. Since this is not an area in which bears usually live, there is nothing in your total set of evidence that would indicate that there is a bear

around. Now suppose again that you have sensory experiences as of a bear attacking you. You again tell yourself that this must be a hallucination or a trick of some sort. Again you turn out to be correct. However, I take it that we would not normally think that you know that it is not a bear. For if you had been following the news coverage or even just chatted with your friends you would not have thus easily supposed that it is not a real bear. You are not ignoring evidence you have as in the first bear example. Rather you just for some reason lack evidence that would be easily available to you.

In this second bear example the person does not know because defeaters to her true belief would be easily available. This motivates the claim that evidence-based justification is not achieved if defeaters are easily available. We can strengthen this claim by thinking again about how unreasonable it is for the person in the second bear example to expect that there is this luck-precluding condition for her true belief that the bear is not real. For sure, the person expects there to be this luck-precluding condition and it is as a matter of fact present. However this expectation is due to a set of evidence that does not include important defeaters that would easily be available. The situation of the person in this second bear case is similar to someone who conducts a survey and takes no precautions to ensure that the sample of people he interviews are representative of the whole population. Even if by some chance the survey leads to accurate results, it is clear that this is too lucky to count as knowledge. Similarly it is extremely lucky that the person suffering from the bear hallucination turns out to be correct in her belief that there is no

bear attacking her. The expectation of the luck-precluding condition is just not reasonable.

Let us now consider a third bear example. Suppose that you live in the area in Austria where a bear is, according to the news coverage and friends and neighbors, living and attacking other living things. However, you lead a very strange life. In fact you spend your whole life in a machine that feeds all your senses fictitious information. You live your life inside a hallucination machine that causes you to experience all kinds of things, all of which are completely fictitious. Now those who feed you the fictitious input read the newspapers and they find it funny to let you one day, as you hallucinate going for a walk, experience a bear running towards you. Nothing in your prior hallucinations points towards there being a bear in the area in which you take your walks. Furthermore the coat of the bear looks somewhat strange. So you conclude that you must be hallucinating. You are correct, and there is a luck-precluding condition linking all of your experiences to the truth that they are hallucinations. However, it seems clear that you have no knowledge that it is not a real bear. Let us explain this.

The person in the third bear scenario is in a hallucination machine. His access to the physical world is hopelessly distorted. He is disconnected from his physical surroundings. In this situation it is entirely a matter of luck if for one of his belief evidence pairs there turns out to be a luck-precluding condition. The person in the third bear scenario is in a situation in which the expectation that there is a luck-precluding

condition is not reasonable. This is not due to any fault in the epistemic agent. It is rather that his situation just does not allow him to gain perceptual knowledge. Even if the luck-precluding condition is present in one instance, the general situation of the epistemic agent renders the expectation of a luck-precluding condition unreasonable. No matter how the epistemic agent processes his fictitious input, his situation does not allow any perceptual knowledge.

In this subsection we have given the following argument: Suppose that someone has evidence for a belief and the luck-precluding condition obtains. If it is not reasonable to expect the luck-precluding condition then the belief is not perceptual knowledge.

Someone is not reasonable in expecting a luck-precluding condition if and only if that person does not have evidence-based justification for that belief. It follows that if someone does not have evidence-based justification for a belief the belief is not perceptual knowledge. It follows if a belief is perceptual knowledge then the person in question has evidence-based justification for the belief. Thus evidence-based justification is a necessary condition for perceptual knowledge.

Summary Box 15: Evidence-based justification is necessary for perceptual knowledge

The luck-precluding condition as an external requirement for perceptual knowledge has to obtain and be expected to obtain by the epistemic agent. Furthermore, this expectation of the luck-precluding condition needs to be reasonable. This is ensured by evidence-

based justification. Evidence-based justification involves a judgment on the basis of the total set of evidence. It requires that no defeater that would easily be available is missing from this total set of evidence, and that the epistemic agent is in a situation in which he can at all arrive at a true belief plus luck-precluding condition about this matter. If any of these conditions of evidence-based justification is absent we can construct a case in which it is just lucky that the luck-precluding condition obtains, and where this type of luck is incompatible with perceptual knowledge. This supports that evidence-based justification is necessary for perceptual knowledge.

5.2.5 True belief revisited

Let us now revisit the requirement of a true belief. If we think about the luck-precluding condition we can clarify how to understand the truth-requirement. As we have seen, the luck-precluding condition consists in a regularity in the physical world between what renders the evidence and the belief for which we have the evidence true. We can generate Gettier intuitions about cases in which this regularity is present on the level of propositions but not on the level of the facts that render the propositions true. At least in the case of perceptual knowledge we should thus understand the truth-requirement to mean that there are facts in the physical world that render the belief in question true. The absence of a luck-precluding condition on the level of facts, of regularities in the physical world, generates a Gettier intuition. This suggests strongly that the truth-requirement needs to be grounded in facts of the physical world, rather than in what is ideally rational to belief or something similarly removed from what goes on the physical world.

The first slot in the luck-precluding condition is filled by evidence as supported by a general assessment of the total set of evidence available to the person in question. So we imagine someone taking stock of all of his evidence and then pronouncing reasonably that this total set of evidence indicates that a certain belief is true. This means that the belief as required for knowledge has to follow this reasonable judgment on the basis of the total set of evidence. We can imagine someone who somehow sees that his total set of evidence clearly and strongly indicates a certain truth, but still refuses to believe it. It is thus legitimate to distinguish between beliefs in general and those that respond to the evidence and in particular to the reasonable assessment of the total set of evidence. Only the latter kind of belief, which is called “acceptance” by Keith Lehrer,²⁸ is required for knowledge.

We have now provided arguments for each of the necessary conditions for knowledge. What we have argued is the following: If someone has perceptual knowledge then he has evidence-based justification for a true belief and the luck-precluding condition obtains. Now we need to argue that these necessary conditions are jointly sufficient for knowledge.

5.3 Sufficient conditions for perceptual knowledge

5.3.1 An argument

²⁸ Compare Keith Lehrer (2000). In particular see chapter 2: Truth and Acceptance.

In this section I develop an argument to the effect that the necessary conditions explained in the preceding section are jointly sufficient for perceptual knowledge. The main idea for the argument can be explained as follows. In the preceding section the arguments for the necessary conditions involve the presentation of a case where the proposed necessary condition is absent. In this case it is a matter of epistemic luck if the person in question ends up with a true belief. Furthermore, these types of epistemic luck are incompatible with knowledge. So the idea is that the necessary conditions can be argued for by exhibiting a counterexample that is based on epistemic luck. The main idea for the argument in this section is that if all the four necessary conditions are present it does not seem that a counterexample that is based on epistemic luck can be constructed. In order to argue for this, we will try to capture again what exactly is involved in problematic forms of epistemic luck.

The argument presented in this section goes as follows:

- (1) If an account of perceptual knowledge presents necessary conditions for knowledge that are not jointly sufficient, then there is a counterexample to it that is based on a problematic form of epistemic luck.
- (2) Problematic forms of epistemic luck arise from a mismatch of the internal and external conditions in an account of perceptual knowledge.
- (3) The here proposed account of perceptual knowledge precludes a mismatch of the internal and external conditions.

- (4) Thus there is no counterexample to this account of perceptual knowledge that is based on epistemic luck.
- (5) Thus the here proposed account of perceptual knowledge presents sufficient conditions.

Let us now try to support each of the premises in this argument.

Summary Box 16: An argument that the necessary conditions are jointly sufficient for perceptual knowledge

The argument for the claim that the necessary conditions are jointly sufficient for perceptual knowledge goes as follows.

- (1) If an account of perceptual knowledge presents necessary conditions for knowledge that are not jointly sufficient, then there is a counterexample to it that is based on a problematic form of epistemic luck.
- (2) Problematic forms of epistemic luck arise from a mismatch of the internal and external conditions in an account of perceptual knowledge.
- (3) The here proposed account of perceptual knowledge precludes a mismatch of the internal and external conditions.
- (4) Thus, there is no counterexample to this account of perceptual knowledge that is based on epistemic luck.
- (5) Thus, the here proposed account of perceptual knowledge presents sufficient conditions.

5.3.2 Problematic forms of epistemic luck

By problematic forms of epistemic luck I mean those that are incompatible with knowledge. By unproblematic forms of epistemic luck I mean those that are compatible with knowledge. Let us first look at some unproblematic forms of epistemic luck.

That I have certain evidence might be a matter of luck in that I might easily have not had the evidence. For example, a witness to a car accident has certain evidence because he happened to be in that specific place at that specific time. Similarly, that I have certain beliefs might be a matter of luck. For example, someone might have a belief because he happened to have read a certain book. Furthermore, that a belief is true might be a matter of luck in that the facts that make it true could easily have not obtained. All of these forms of epistemic luck are unproblematic.

What these unproblematic forms of epistemic luck seem to have in common is that they target one of the necessary conditions for knowledge in isolation. They target either the belief requirement or the evidence requirement or the truth requirement or some other necessary condition for knowledge in isolation. So consider a case where someone has evidence-based justification for a true belief and the luck-precluding condition obtains. If we consider each of these requirements in isolation, they might easily have not obtained. Someone might have had a different belief, different or no evidence, the luck-precluding condition might have failed to obtain. However, all of this is unproblematic. It thus seems

that if we look at just one of these necessary conditions in isolation the epistemic luck that targets this necessary condition is unproblematic.

Now let us look at some problematic forms of epistemic luck. Obvious cases are a lucky guess, a Gettier case, or one of those cases where the luck-precluding condition is not expected. If we have a lucky guess, what characterizes it is that one has a belief for which one has no proper evidence but the belief happens to be true nevertheless. It seems that there are three things that characterize this problematic form of epistemic luck: First, it concerns how different requirements for knowledge are related to one another. It does not concern one of the requirements in isolation. In the case of a lucky guess we have belief, truth and the lack of evidence that are all crucially involved. Second, some of the requirements involved are internal and some are external. In a lucky guess there is a belief as an internal requirement, evidence also as an internal requirement, and truth as an external requirement. Third, in a lucky guess the external requirements obtain, but the internal requirements are flawed so that even though “the world does its part,” the result is not knowledge. Even though a lucky guess results into a true belief, the lack of evidence means that the appropriate internal conditions for knowledge do not obtain, thus preventing knowledge.

Let us confirm this diagnosis for Gettier cases. First, involved are the evidence requirement, the belief requirement, and the truth requirement, and also the lack of a luck-precluding condition. Second, among those the evidence and belief requirements are

internal requirements, the truth requirement and the luck-precluding condition are external requirements. Third, even if the external requirement obtains in that the belief is true, the internal conditions are mismatched with the true belief. There is a gap between the internal and the external conditions that is not bridged by a luck-precluding condition.

It thus seems that problematic forms of epistemic luck arise because the internal and the external requirements for knowledge are not sufficiently matched with one another.

Either an internal condition is missing or it is present but it is not the right one for the external conditions that occur in the case at hand. This is what supports the second premise in the argument given above.

5.3.3 Counterexamples based on epistemic luck

Let us assume that this conjecture as to how problematic forms of epistemic luck arise is correct. What follows for how counterexamples based on epistemic luck can be designed? It seems that all such counterexamples need to exploit a possibility of a mismatch between the internal and the external requirements for knowledge. Such a counterexample needs to show that somehow there can arise a gap between the internal and the external requirements that is incompatible with knowledge.

Let us think about what the possibility or impossibility of counterexamples based on epistemic luck means for an account of knowledge. There have been some attempts to try to understand knowledge by analyzing it into purely internal or purely external

conditions. For example, if one thinks that truth has nothing to do with facts in the physical world, but is somehow related to what it is ideally rational to believe, then it is possible to define knowledge as involving only internal requirements. Alternatively, one might think that all the conditions for knowledge involve external elements, that there are no truly internal conditions that are necessary for knowledge. These are examples of attempts to understand knowledge either as purely internal or purely external.

Now in the preceding section we have provided arguments to the effect that there are both internal and external necessary conditions for knowledge. We have attempted to show that an adequate account of knowledge needs to contain both internal and external conditions. If we accept this, then we are faced with the following challenge: The internal and the external conditions for knowledge need to be combined in such a way that there is no mismatch possible. The idea is that the internal and the external conditions need to be combined in such a way that a counterexample that shows that they are mismatched is not possible. And that seems to be the challenge for coming up with sufficient conditions for knowledge. Let us explain this furthermore.

At this point in the chapter we have already argued that both internal and external conditions are necessary for knowledge. Now we ask whether they are jointly sufficient. We assume that they are each necessary. If they are not sufficient then either something is missing or they are not well matched with one another. Now that is precisely how problematic forms of epistemic luck arise. They either point to a missing condition or to a

mismatch of the conditions or both. So it seems to follow that if the conditions are not jointly sufficient, there should be a counterexample based on epistemic luck. This is what suggests the first premise in the argument above to be correct.

5.3.4 Epistemic luck and the necessary conditions

We now need to argue that this mismatch of the internal and the external conditions is not possible in the case of the necessary conditions on which the account of perceptual knowledge developed here is based. Let us discuss how the internal and the external conditions are related in the account of perceptual knowledge developed here.

As a reminder, I am only concerned with perceptual knowledge, where the evidence consists in sensory experiences. Let us begin with the internal conditions and see how they are tailored to the external conditions. Internal conditions are the belief- and evidence-requirement. A belief is understood to be something that occurs in the mind and that has a content that can be expressed in a proposition. To take a simple example, right now I believe that I am typing on a computer. This is a belief about the physical world. As a belief it is something that is internally available to me. Evidence is also understood to be internal states, consisting in beliefs and sensory experiences. In the example the evidence for my belief that I am typing on a computer consists in a variety of sensory experiences as of my typing on a computer. What are internally available are the sensory experience and its reading as one as of typing on a computer. In this simple example I have a belief and sensory evidence that speaks directly for the truth of this belief.

Furthermore, there is no evidence that speaks against it or that could serve as a defeater. So as far as it is internally available my assessment of the total set of evidence available to me indicates that the sensory experiences show that it is true that I am typing on a computer.

Now we need to see how these internal conditions are tailored to the external conditions. There are mainly two such external conditions: first, my assessment of the total set of evidence needs to be reasonable. Second, there needs to be a luck-precluding condition linking my evidence to the true belief for which I have the evidence. None of these external conditions is identical to an internal condition. In fact all the internal conditions might obtain when none of the external conditions obtain. Nevertheless the internal conditions reflect the external conditions. Let us see how this is the case.

In reading my sensory experience I match it with the belief for which it is evidence. All of this is internally available. Now, what the luck-precluding condition does is it links through a regularity this evidence and the belief for which we have the evidence. The luck-precluding condition consists in a regularity in the physical world. This regularity would in the example involve that there is a computer that causes in me the various sensory experiences as of typing on a computer. If the luck-precluding condition is present then there is a regularity, both in general and in this instance, that supports my moving from the sensory experiences to the belief that I am typing on a computer. The

regularity supports my reading of the sensory experience and my adopting a belief about the physical world on the basis of this reading.

Let us try to get more precise about this. Suppose that your evidence for a belief that p consists in a readings that q of sensory experiences. Here is what is internally available: the belief that p , the reading of the sensory experience that q , and the relationship that q is evidence for p . Now the luck-precluding condition means the following: There is a regularity present, both in general and in this instance, between that in the physical world which renders q true and that in the physical world which renders p true. This regularity however is an external condition and it might be absent even if all the internal conditions obtain.

What we have seen now is that the luck-precluding condition is expected by and reflected in moving from a reading of a sensory experience to a belief about the physical world on the basis of this evidence. This expectation can be unreasonable in some cases, as when one's situation precludes any perceptual knowledge or when easily available defeaters are not sought out. However, when we make an assessment of the total evidence and then adopt a belief about the physical world on that basis, the luck-precluding condition linking this evidence to the true belief is expected. So the luck-precluding condition just underwrites through a regularity in the physical world the move from evidence to the belief about the physical world on the basis of this evidence.

So far we have looked at the internal conditions and tried to see how they are tailored to the external conditions. Let us now look at the external conditions and see how they are tailored to the internal conditions. The external conditions are the truth-requirement, the luck-precluding condition, and the requirement of evidence-based justification over and above evidence. The truth-requirement is satisfied automatically when the luck-precluding condition is present. For the luck-precluding condition links the evidence to the true belief. The luck-precluding condition is based on a regularity between what renders the evidence true and what renders the belief for which we have the evidence true. This regularity has to be present not just in general but also in this instance. That means that if the luck-precluding condition is present in this instance then the belief that enters into one of the slots of the luck-precluding condition must also be true. As we have said earlier, the idea is that given the evidence and the luck-precluding condition, the belief for which we have the evidence has to be true.

The truth-requirement is thus satisfied in virtue of the luck-precluding condition. This also changes how we should understand the truth-requirement for knowledge. Both the truth-requirement and the luck-precluding condition are external conditions, but the luck-precluding condition involves a link between the evidence and the true belief. In other words, by moving from just a truth-requirement as an external requirement to the luck-precluding condition we demand that the truth is linked to the evidence, that the true belief is part of a regularity between it and the evidence. So we link the truth-requirement to the evidence requirement and thus to an internal condition.

Similarly, even though evidence-based justification involves an external element, it involves a very strong internal element. Evidence-based justification requires evidence, where evidence is understood as an internal condition. What evidence-based justification adds to the evidence is just the assessment of the total set of evidence and the reasonable expectation that this evidence is linked to the truth. The further external conditions just constrain when someone is justified in virtue of some evidence. They are not conditions that would allow someone to be justified without any evidence at all.

Let us try to see more clearly how the internal and the external conditions are matched. One of the key features of this match is the luck-precluding condition. The luck-precluding condition links the evidence to the true belief through an external condition; however the luck-precluding condition has to be expected, and it merely underwrites moving from the evidence to the belief for which we have the evidence. We can think about it this way: There is evidence on the one hand and a true belief on the other hand. They are linked both through an internal condition and an external condition. We need to understand the evidence to indicate that the belief is true. This is internal. However, the luck-precluding condition provides an external regularity between what renders the evidence true and what renders the belief true. In moving from the evidence to the belief we expect the luck-precluding condition, so we have both an internal connection and an external connection between the evidence and the true belief. Furthermore the internal

connection expects the external connection, and the external connection underwrites the internal connection.

As this somewhat lengthy account is supposed to show, if we are to have both internal and external conditions in an account of knowledge, it seems very hard to see how they could possibly be better intertwined or more tightly connected with one another.

Although this does not prove that there is no possibility of a mismatch, as far as possible it suggests that there is no such possibility of a mismatch between the internal and the external conditions in this account of knowledge. This is what supports the third premise.

Even if the argument given in this section turns out to be wanting in some respect, it seems to me to contain an important insight. This insight is that the game of finding sufficient conditions for knowledge is one of trying to tailor internal and external conditions to each other so that they fit together as seamlessly as possible. Only if that is achieved can we hope to find a set of conditions that are truly sufficient for knowledge.

5.3.5 Comparison with other views

The claim in this chapter has been that we need to combine internal and external conditions for knowledge so that the resulting account of perceptual knowledge is not subject to a counterexample based on epistemic luck. This idea is already present in, for example, Keith Lehrer's account of knowledge, and also in Matthias Steup. Let me briefly explain the similarities and differences between the account developed in the

dissertation and the ones explained in Lehrer's and Steup's work. Afterwards I will briefly address an account of epistemic luck developed by Pritchard, which does not understand epistemic luck as arising from a mismatch of internal and external conditions for knowledge.

In his "Theory of Knowledge" Keith Lehrer (2000) proposes an account of knowledge that combines internal coherentist justification with the external requirement that key beliefs are actually true. The idea is that we have a target belief for which we have an internalist coherentist justification. Now the external requirement for the target belief to be knowledge is not just that this belief is true, but also that certain key beliefs, that the epistemic agent is trustworthy with respect to the matter at hand, are both justified and true. So the idea is that given this account of knowledge it does not turn out to be lucky that the target belief is true, because not only it but also other key elements of the system of beliefs need to be true. If thus not only the target belief but also crucial other beliefs are both internally justified and true, then it is according to Lehrer not lucky that the target belief is true given the internal justification.

The similarities between Lehrer's account and the one developed in the dissertation are more important than the differences. So let me begin by explaining the similarities. We both search for a match of the internal and the external requirements that would render the resulting account of knowledge immune from counterexamples based on epistemic luck. Furthermore, we both regard as the key for doing so that somehow the truth of any

belief must be in some suitable way expected in virtue of having a justification for this belief. Also, we both search for a truth connection for an essentially internalist account of justification as the key for matching the internal and external requirements for knowledge. In addition, the account of evidence-based justification developed in the fourth chapter is compatible with a coherentist account of it. So there could be a coherentist version of the position developed in the dissertation.

There are also some differences. Lehrer sees as crucial for matching the internal and external conditions that there are justified beliefs on a meta-level to the effect that the epistemic agent is trustworthy in the matter at hand. The idea is that this belief about the agent's trustworthiness both supports and is supported by those beliefs in respect to which the agent is trustworthy. If this key belief about the trustworthiness is justified and true, then this, according to Lehrer, renders it non-lucky that the target belief for which we have the justification is true.

In contrast the position developed here sees the truth connection as provided by an external condition, the luck-precluding condition. This external condition has to be expected through the evidence-based justification. The main problem of finding a truth connection is how the evidence and the true belief for which we have the evidence are connected with one another. I propose that they are connected both by an external and an internal condition. The luck-precluding condition links them through a regularity in the

physical world. However, the luck-precluding condition has to be expected in virtue of an evidence-based justification for the true belief in question.

Lehrer combines internal justification with truth of the target belief and other key beliefs. The account of knowledge developed here combines an essentially internal justification with a further external condition, namely that the evidence and the true belief need to be connected by a physical regularity in the world. The external requirements according to the account developed here are thus different from the ones proposed by Lehrer. Also what it means for them to be matched is different. However, it is clear that the present account owes much to ideas developed by Lehrer.

Let us now turn to the view proposed by Matthias Steup (2001) in “Epistemic Duty, Evidence, and Internality.” Steup also searches for a match of internal and external requirements for knowledge. His proposal is to combine internal justification with reliabilism. The idea is that we need for knowledge both internal justification and reliability, the latter in particular as a response to Gettier problems. There are again similarities here. The main similarity is that an internal justification is combined with a further external condition that is not identical to just requiring that the target belief is true.

There are various details that are different, but the most important difference is that reliability is a property of a cognitive capacity in an environment. Reliability is all about whether an epistemic agent is properly equipped with cognitive abilities so as to acquire

through these cognitive capacities mostly true beliefs about his environment. To say that a belief must be reliably produced is in effect to say that the epistemic agent is properly equipped and arrives at mostly true beliefs in his environment.

Now there is some similarity between luck-precluding conditions and reliability in that if the luck-precluding condition is usually present, the epistemic agent who has evidence-based justification will end up with mostly true perceptual beliefs and thus with reliable cognitive capacities. However, it is possible that the regularities on which the luck-precluding conditions depend stop obtaining after some time or that the agent begins to live somewhere where they do not obtain. Luck-precluding conditions depend on regularities in the physical world, and we can imagine scenarios in which these regularities change or stop obtaining or exist only in certain regions of an actual environment. The problem is that the reliabilist then has to say that the same cognitive capacities are no longer reliable even though there might be no change in the capacities of the agent. Let us explain this furthermore.

If an epistemic agent has evidence-based justification and the luck-precluding conditions obtain usually, then it is also true that the cognitive capacities of the epistemic agent are in this environment reliable. However, it is not true that if an agent has evidence-based justification and the luck-precluding condition obtains only sometimes or only in some parts of the environment or only during certain periods of time, his capacities are reliable in this environment. Furthermore, it might be that someone's capacities are generally

reliable but in a given instance the reliably produced belief is false, or true in a Gettier-style way. In these cases we would have a reliably produced belief but the luck-precluding condition does not obtain. So there are clear cases where reliability and luck-precluding conditions come apart.

Furthermore, I take it that we have a good understanding of what regularities in the physical world are. We understand what it means that as a matter of a regularity barns have barn facades, and barn facades cause in an epistemic agent sensory experiences as of barn facades. All of this is much less problematic than trying to understand and delineate cognitive capacities. We understand the distinctions, for example, between perception and deductive inference, between color perception and shape recognition, face recognition and memory. However, what we understand seems to be that there are different kinds of evidence for different kinds of beliefs. We would not try to justify a belief that this is a triangle through color vision. We would not try to justify a perceptual belief through a conceptual deductive argument. However, it is possible that the cognitive processes that allow us to gain evidence for our beliefs actually overlap or serve multiple purposes. It is possible that one and the same cognitive mechanism in the brain gets activated in perception and in memory. In general, if one talks about reliable processes rather than the kind of regularities in which the luck-precluding conditions consist, then one has a lot to explain, including the following. One needs to explain what exactly a cognitive process is supposed to be, how it is related to cognitive mechanisms in the brain, and how they are delineated into different cognitive processes. It is thus not only

that luck-precluding conditions are different from reliability, but also that reliabilists have particular burdens of explication that the position presented in the dissertation does not have.

Let us now consider again Duncan Pritchard's (2005) distinction between veritic and reflective epistemic luck as explained in his book "Epistemic Luck." As I have suggested in the second chapter, the distinction between internal and external justification is what allows Pritchard to draw this distinction. Veritic epistemic luck arises because of the absence of external justification. Reflective epistemic luck arises because of the presence of a gap between the internal justification and the truth that it indicates. In order to avoid veritic epistemic luck we need, according to Pritchard, an external justification. In order to avoid reflective epistemic luck, we would need a gap-closing internal justification.

There are obvious differences between Pritchard's account of epistemic luck and the one presented here. What I would like to do now, is to see whether the two forms of epistemic luck as defined by Pritchard would be considered problematic given the account developed in this dissertation. Veritic epistemic luck arises if and only if a certain type of external justification for a true belief is absent. There is no requirement of evidence or evidence-based justification. The idea is merely that a certain type of epistemic luck can be fully avoided through a certain type of external justification. Suppose that the external justification is present. This external justification need not be reflected in any evidence or any evidence-based justification available to the epistemic agent in question. Given the

account developed here, we would have to say that even if this external justification is present, the resulting situation might still be subject to a problematic form of epistemic luck due to the absence of evidence and evidence-based justification. Given the account developed in the dissertation, avoiding veritic epistemic luck does not mean avoiding all problematic forms of epistemic luck.

Reflective epistemic luck, now, is present if and only if a gap-closing internal justification is absent. Pritchard claims that this form of epistemic luck is unavoidable. What I have tried to show is that even given an essentially internal, but not gap-closing justification we can avoid those forms of epistemic luck that are incompatible with knowledge. Reflective epistemic luck is very similar to skeptical epistemic luck in that both depend on the problem of the gap. However, if non-skeptical epistemic luck is absent, skeptical epistemic luck would still be present, but it would no longer be incompatible with knowledge. If we accept that reflective epistemic luck is similar to skeptical epistemic luck, we get the result that given certain conditions, even though reflective epistemic luck is present, it no longer presents a problem for our knowledge.

Given the account developed here we thus get the following result: Avoiding veritic epistemic luck does not mean to avoid all problematic forms of epistemic luck. In some circumstances when reflective epistemic luck is present, it is not a problematic form of epistemic luck. Again, this makes sense only given the account of problematic forms of epistemic luck developed here.

5.4 A reply to some skeptical worries

In the first two sections of this chapter we have seen which conditions are necessary and sufficient for perceptual knowledge. Now we need to turn to the question of whether these conditions obtain for us in the actual world. I will discuss two versions of this question. First, I will develop an argument to the effect that we are capable of perceptual knowledge. Second, we will see that the account of perceptual knowledge is compatible with our knowing a lot of those things that we ordinarily take ourselves to know. So in the first part of this section I will address the skeptical question of whether we ever have any perceptual knowledge. In the second part we will address the question of whether the account of perceptual knowledge would allow us to know a wide range of things.

5.4.1 Two senses of possibility of perceptual knowledge

The argument I will give has the conclusion that we are capable of perceptual knowledge. Now the skeptic might be taken to claim that we in fact have no perceptual knowledge, or alternatively that perceptual knowledge is impossible. The argument I will develop addresses only the second claim. The aim of the argument is to establish that in an important sense perceptual knowledge is possible. This is compatible with our not actually having any perceptual knowledge.

If the target is to establish that perceptual knowledge is possible, this of course raises the question of what kind of possibility we are talking about. We also have talked at the

beginning of this chapter and the preceding chapters about epistemic agents who are incapable of perceptual knowledge even though from the inside they are identical to us. We talked about a brain in vat who has all the same sensory experiences and beliefs we have, who is however incapable of perceptual knowledge. So we seem to grant the possibility that an epistemic agent that is indistinguishable from the inside from us might be incapable of knowledge. This of course renders all the more pressing the question of what the sense of possibility is given which we are capable of perceptual knowledge.

Perceptual knowledge is knowledge of the physical world on the basis of sensory experience. We have proposed that for this knowledge both internal and external conditions are required. Internal conditions are internal to the mind of the epistemic agent in question. External conditions involve features of the environment in which the epistemic agent lives. Both these conditions need to be combined to yield perceptual knowledge.

This means that there are at least the following two ways in which it might be that perceptual knowledge is impossible. It might be impossible due to a defect in the epistemic agent. It might also be impossible due to a defect in the environment in which the epistemic agent lives. It thus might be that an epistemic agent is because of the nature of his epistemic agency incapable of perceptual knowledge. It might also be that an epistemic agent who is generally capable of perceptual knowledge is trapped in an environment in which no perceptual knowledge is possible.

An example of an epistemic agent who is incapable of perceptual knowledge would be one who does not evaluate any readings of sensory experiences, but believes whatever his sensory experiences suggest. Such an epistemic agent would believe that a stick partly immersed in water is bent, and then straight again when being pulled out of the water. This epistemic agent is incapable of perceptual knowledge because he is incapable of appreciating defeaters. He just believes whatever his sensory experiences indicate. No matter what his surroundings are, no matter what the environment is like in which he lives, he is generally incapable of gaining perceptual knowledge.

An example of an epistemic agent who might be generally capable of perceptual knowledge but who is trapped in circumstances that do not allow such knowledge is the familiar brain in a vat case. If someone is a brain in a vat and he is completely deceived about what his physical surroundings are like, it seems that then he is incapable of perceptual knowledge due to the nature of his surroundings. His circumstances are such so as not to allow any perceptual knowledge, no matter how he responds to the fictitious input.

The argument I want to give is that we know that we are not in the first situation. We know that we are capable of perceptual knowledge. However, it is possible that we are in the second situation. It is possible that I am just a brain in a vat that is deceived about everything in my physical environment. To say that we are capable of perceptual

knowledge is to say that given the kind of epistemic agents we are we can gain perceptual knowledge. However, it might be that we are trapped in circumstances that undermine our epistemic agency in that we cannot gain any knowledge in those circumstances.

We thus need to distinguish between two senses of possibility of knowledge. First there is the sense in which knowledge is possible that requires that we are the kind of epistemic agents who are capable of perceptual knowledge. This is what I try to establish in this chapter. There is also the sense in which knowledge is possible that requires both that we are the kind of epistemic agent who is capable of perceptual knowledge and that we are in fact in an environment in which such knowledge is possible. Lucky, the brain in a vat that we have described at the beginning of this chapter, might be capable of perceptual knowledge in the first sense, but she is in a situation in which knowledge is impossible in the second sense. Even Lucky might be capable of knowledge in the first sense, despite her unfortunate situation in which knowledge is not possible in the second sense.

5.4.2 Knowledge of the nature of our epistemic agency

What would it mean for us to be the kind of epistemic agents who are capable of perceptual knowledge? It seems that it would mean that we are capable of acquiring evidence-based justification for a belief. The further necessary conditions for knowledge are external conditions. Those are the conditions that need to be provided by our environment. Whether we are capable of evidence-based justification on the other hand reflects our nature as epistemic agents. If we know that we are capable of having

evidence-based justification for a belief because of our knowledge of the nature of our epistemic agency, we might be able to argue that given the nature of our epistemic agency we are capable of perceptual knowledge.

What do we know about the nature of our epistemic agency? In the fourth chapter I have distinguished three epistemic capacities that characterize our epistemic agency in so far as perceptual knowledge is concerned. Let us now see that we know that we have each of these three capacities.

The capacity to receive sensory experiences is understood as internally characterized. Even a creature that does not have sensory organs, or receives fictitious input, or hallucinates or has vivid dreams rather than sensory experiences that result from a causal contact with the physical world would have this capacity to receive sensory experiences. Saying that we have sensory experiences is not to say that we have internal states that are caused by the physical world through sensory channels. It merely means that among our internal states we can distinguish some as different in kind from beliefs and that can serve as evidence for beliefs about our immediate physical surroundings. In the fourth chapter I give three characteristics that allow us to distinguish sensory experiences from beliefs from the inside. However, given that no commitment to anything inaccessible from the inside is made when we say that we have the capacity to receive sensory experience, it seems unproblematic to say that we know that we have this capacity.

The second capacity is that of reading our sensory experiences. We have sensory experiences, and we are able to see for which beliefs about the physical world they are evidence. In reading sensory experiences we match them with those beliefs about the physical world that they naturally suggest as true. We can verify that we have this capacity by seeing whether we would understand a specific kind of challenge to the readings of sensory experiences. Suppose thus that you go for a walk with a friend and you see a structure at a little distance. You say that that seems to be a nice barn. Your friend responds that it rather seems like a nice cottage. This would be a discussion of how to read the sensory experience. It seems that we can understand this challenge. Thus we know that we possess the capacity to read sensory experiences. Again this capacity is completely internally understood and characterized. Thus it is possible to possess this capacity and yet make occasional or severe errors in applying it to specific cases.

The third capacity concerns our evaluation of readings of sensory experiences. Here we can verify whether we have this capacity by seeing how we would respond to an illusion. If in response to an illusion we see that the sensory experience should be read as one, for example, of a bent stick partly immersed in water, however that this belief as read off of the sensory experience is not true, we have the capacity to appreciate defeaters. The capacity to appreciate defeaters concerns our ability to withhold belief in something we have read off of a sensory experience. It also concerns whether we are able to see a need for further evidence, to investigate something further before adopting a belief. It seems again clear that we know that we have this capacity.

We thus know that we have each of the three capacities that characterize our epistemic agency in cases of perceptual knowledge. We know the nature of our epistemic agency to be characterized by these three capacities.

5.4.3 We know that we are capable of perceptual knowledge

If we accept that both internal and external conditions need to obtain for us to gain perceptual knowledge, we accept that there is something that we as epistemic agents need to contribute and something that the environment in which we live needs to contribute. To say that given the nature of our epistemic agency we are capable of perceptual knowledge is to say that given the nature of our epistemic agency the internal requirements obtain. Let us explain this now.

We have just seen that we know that we have the three capacities that characterize our epistemic agency for cases of perceptual knowledge. These are the three capacities that allow us to have evidence-based justification for our beliefs. There are external requirements on evidence-based justification, but they again are something that the world needs to contribute. Let us thus examine the internal part of the evidence-based justification. That is, we need to see what the internal basis of perceptual knowledge is, and whether we know that we are capable of achieving this internal basis.

Evidence-based justification in the case of perceptual knowledge requires that we judge based on the total set of evidence available to us that it indicates that a certain belief is true. So we need evidence. This evidence will consist in sensory experiences and relevant beliefs. We know that we have the capacity to receive sensory experience. When we read the sensory experience we match it with the beliefs that it naturally indicates as true, thus understanding it as evidence for a specific belief. This means that we are capable of having sensory evidence for beliefs about the physical world. We also know that we have the capacity to appreciate defeaters and to see the need for further investigation or of reconciling conflicting evidence. This means that we are capable of deciding whether a total set of evidence indicates that a given belief is true. It thus seems that given our knowledge of the nature of our epistemic agency, we know that we are capable of evidence-based justification.

It is possible that the external requirements for evidence-based justification are not met in our actual environment. It is possible that the other external conditions for knowledge do not obtain in our actual environment. It is also possible that we make mistakes in exercising the capacities that allow us to gain evidence-based justification. However, it seems that on the basis of knowing that we have these three capacities we know the nature of our epistemic agency. Furthermore, we know that given this nature of our epistemic agency we are capable of evidence-based justification for a belief. This in turn means that we know that we are capable of realizing the internal requirements for perceptual knowledge. It follows that given favorable external circumstances we will

have knowledge. Given the nature of our epistemic agency we are capable of perceptual knowledge.

The argument just given does not establish that we indeed have perceptual knowledge. It also does not establish that we live in a world in which perceptual knowledge is possible for us. It is possible that we are a brain in a vat, and that we are in circumstances that do not allow us to gain any perceptual knowledge. What the argument just given establishes is merely that given the nature of our epistemic agency we are capable of perceptual knowledge. This means that given favorable external circumstances we will have perceptual knowledge.

5.4.4 Our perceptual knowledge

One might ask whether given the account of perceptual knowledge we could turn out to have a lot of perceptual knowledge. The answer to this question is that given favorable external circumstances we will have a large measure of perceptual knowledge. Assuming that the external conditions for knowledge obtain, there is nothing excessively strict about the internal requirements. For our perceptual knowledge sensory experiences and related beliefs serve as evidence. This is what we would intuitively agree on as the evidence for our beliefs about the physical world. When we understand a sensory experience as evidence we read it as indicating that a specific belief about the physical world is true. Then we will check it against other sensory experiences and other beliefs. This will lead us to adopt a belief about the physical world on the basis of the sensory experience. This

is a somewhat lengthy account of what goes on when we believe that there is a barn in the field or that there is a flower in the garden. However, the account does not require anything that is not present in ordinary situations.

Furthermore, given the account developed here we probably have quite a lot of knowledge. For the luck-precluding conditions are regularities, so if they obtain in one case, they will also obtain in a variety of other cases. So if we gain knowledge in one case, due to the nature of a luck-precluding condition we are likely to have a lot of knowledge of similar things. Furthermore, the luck-precluding conditions make sure that a certain bit of evidence will indicate the truth. This means that similar evidence, if it leads us to knowledge in one case, is very likely to lead us to knowledge also in other cases. So the conditions for knowledge make sure that it is likely that there is in fact a lot of perceptual knowledge if there is some.

Summary Box 17: An argument that we are capable of perceptual knowledge

In what sense of possibility should we try to show that perceptual knowledge is possible?

There are two important senses of this possibility of perceptual knowledge. Perceptual knowledge might be possible for the kind of epistemic agent we are. Alternatively, perceptual knowledge might be possible for us in our actual environment. I aim only at showing that perceptual knowledge is possible in the first sense.

For us to be the kind of epistemic agent for whom perceptual knowledge is possible, we need to have the kind of epistemic agency that allows for us to realize the internal

requirements for perceptual knowledge. Our epistemic agency consists in so far as perceptual knowledge is concerned in three capacities, the capacity to receive sensory experiences, to read them, and to evaluate the readings of the sensory experiences. We know that we have these capacities. Given this, we know that we are capable of realizing the internal requirements for perceptual knowledge. This means that we are capable in virtue of our known epistemic agency of gaining perceptual knowledge.

It is, however, still possible that we are in a situation similar to Lucky, the brain in a vat. Lucky has also, as we suppose, the kind of epistemic agency that would in principle allow her to gain knowledge, but the external requirements are never met.

5.5 The problem and the proposal revisited

One of the most important things that this chapter shows is that the notion of epistemic luck is crucial for defining perceptual knowledge. I have used the notion of epistemic luck to argue that certain conditions are necessary and jointly sufficient for perceptual knowledge. Those arguments depend, as it seems to me, not on problematic intuitions that are somehow taken as data or pre-theoretic support. Instead the arguments depend on intuitions that can be integrated in a whole theory of perceptual knowledge. This theory is not just supported by such intuitions, but it gains much of its justification from attempts to respond to epistemological problems, including the Gettier problem, the new evil demon problem, and skeptical problems.

For the notion of epistemic luck to be crucial in defining knowledge we need to distinguish the various forms of epistemic luck and ask whether such a form of epistemic luck is compatible or incompatible with knowledge. The answer to the question of whether a form of epistemic luck is compatible or incompatible with knowledge should accommodate our intuitions about epistemic luck, but it should also be capable of being integrated into a theory of knowledge.

At the beginning of this chapter I talk about Lucky, the brain in a vat who is from the inside identical to me. How can it be that knowledge depends to such a degree on epistemic luck that although Lucky is incapable of perceptual knowledge, I might have such knowledge? The answer is that Lucky and I are equally capable of perceptual knowledge if we mean by that whether our epistemic agency would allow us to gain evidence-based justification for a belief. Lucky and I are identical as far as the nature of our epistemic agency is concerned. We are both given this nature of our epistemic agency capable of perceptual knowledge. However, Lucky is trapped in circumstances that render it impossible for her to actually gain evidence-based justification and perceptual knowledge. She has the kind of epistemic agency that would allow her to gain perceptual knowledge given favorable circumstances. It is just that Lucky's circumstances are not favorable.

There is no guarantee of perceptual knowledge for anyone of us. We each might be not just from the inside but also from the outside identical to Lucky. However, we can know

that we have the kind of epistemic agency that would given favorable circumstances allow us to gain perceptual knowledge. Although I have argued in the second section of this chapter that there is no counterexample based on epistemic luck to the here proposed account of knowledge, there is an element of luck involved when we actually gain knowledge. We each as epistemic agents depend on circumstances beyond our control, we depend on being in the right environment to permit us to gain perceptual knowledge. This is an element of epistemic luck that cannot be avoided, since no epistemic agent can fashion his own environment. We find ourselves in this world, we find ourselves with certain capacities, and we can only hope that these capacities allow us to gain knowledge of this world.

6. CONCLUSION

The project of the dissertation is to develop an account of perceptual knowledge. In this conclusion I would like to sketch the boundaries of this project. First, we will see whether the account of perceptual knowledge can also be applied to other kinds of knowledge. Second, we will take a look at some related and unfinished work. The aim in both sections will be to clarify the boundaries of the present project, not to tackle entirely new projects. For example, what I say about testimony should not be taken as an attempt to develop an account of this kind of knowledge. The point is rather to allow us to see more clearly the boundaries of the project in the dissertation.

6.1 Other kinds of knowledge, both of the physical world and non-physical things

The focus of the dissertation is on a specific kind of knowledge of the physical world. This kind of knowledge of the physical world requires that the epistemic agent has a sensory experience that he reads as evidence for a specific belief about the physical world. On the basis of the total set of evidence available to the epistemic agent, he then adopts the specific belief about the physical world that he has read off of the sensory experience. So the belief is directly indicated as true by the sensory experience. Examples would be that there is a barn in the field, or that a sheepdog is chasing a sheep from one end of the field to the other.

The question to ask now is whether the account of such perceptual knowledge would also work for other kinds of knowledge. Let us distinguish between knowledge of the physical world and knowledge of non-physical things. By a physical object I mean an object that exists in both space and time. By a physical property I mean those properties that can only be instantiated by a physical object. By non-physical things I mean those objects and their properties that are not physical objects and physical properties.

Other kinds of knowledge of the physical world include inductive or statistical knowledge, theoretical knowledge, knowledge through memory, knowledge through testimony. This is not supposed to be a complete list, just a list of some obvious and important other kinds of knowledge of the physical world. Non-physical knowledge includes knowledge of a language or conceptual scheme, mathematical knowledge, and philosophical knowledge. Again, this is not supposed to be a complete list of such kinds of non-physical knowledge. Let us now discuss the prospects of developing an account of these other kinds of knowledge that is similar to the account of perceptual knowledge developed in the dissertation.

As it seems to me, all the other kinds of knowledge of the physical world just mentioned are also characterized by a gap between evidence and true belief, but it does not seem that a luck-precluding condition as understood in the dissertation can solve the problem. Let us characterize the gap between evidence and true belief for these other kinds of

knowledge of the physical world, and then discuss whether a luck-precluding condition could allow us to solve the problem.

Let us look at a simple case of an inductive generalization. Suppose that we conduct a survey as to preferences of laundry detergent. We interview 500 people and find out that of those 500, 300 prefer brand X and 200 prefer brand Y. Now we want to conclude something about the whole population in that corner of the world. It seems clear that we have a gap between evidence and true belief about the whole population. We might have interviewed only people in a specific subgroup, of which we are unaware and which has very different laundry detergent preferences from the rest of the population. We might screen the 500 we have interviewed in order to decrease the risk of this difference between those interviewed and the whole population, but it does not seem that we can ever fully close this gap.

If we recognize that there is a gap in the case of inductive knowledge of the physical world, that the evidence is compatible with the falsity of the belief for which we have the evidence, the question arises whether a luck-precluding condition could help us out.

Now, a luck-precluding condition consists in a physical regularity, for example a causal regularity or a part-whole regularity, or a regularity of temporal sequence. None of this seems to apply to the case of inductive generalization. How could there be a causal or part-whole regularity between those interviewed and the whole population that would ensure the truth of the inductive generalization? Maybe if we looked at a population in

which there are 500 decision makers about laundry detergent, and we interview all of those, who then make the rest of the population buy laundry detergent according to their preferences, we would have a luck-precluding condition. However, it is clear that nothing like that underlies inductive reasoning, and that whatever renders the gap compatible with knowledge is not a luck-precluding condition that consists in a physical regularity as understood in the dissertation.

Let us now consider theoretical knowledge, for example of elementary particles. As it is usually understood, theoretical knowledge is based on a scientific theory and goes beyond specific observations. The idea is that theoretical knowledge is not perceptual knowledge as understood in the dissertation in that it is not that closely based on sensory experiences. In Quine's image, theoretical knowledge lies more towards the center of our web of knowledge rather than at the periphery, which is more closely connected to sensory experiences. Now it is far from clear what exactly such theoretical knowledge of the physical world is, however, the idea is that such knowledge depends on accepting a larger scientific theory which goes beyond merely recording observations. It seems that such theoretical knowledge is characterized by a gap between evidence and the true belief for which we have the evidence. Consider an example, that now an electron is passing through the cloud chamber. It seems that I could have evidence for this claim in form of a theory that is supported by some observation and experiments, and yet the claim that there is an electron passing through the cloud chamber might be false.

Again, it does not seem that the gap could be rendered compatible with knowledge by a luck-precluding condition as understood in the dissertation. Such a luck-precluding condition would have to link the evidence that includes the theory to the true belief for which we have the evidence. It is not clear how this evidence that includes a scientific theory could be part of a physical regularity that relates it, including the theory, to the true belief for which we have the evidence. If we take the example of the electron passing through the cloud chamber, it does not seem that there is a physical regularity that somehow links the evidence, including the theory, to the true belief that there is an electron passing through the cloud chamber. At least it seems that theoretical knowledge is further removed from the physical regularities that link evidence and true belief in the case of perceptual knowledge.

In the case of knowledge of the physical world based on memory, we again have a gap between the evidence and the true belief for which we have the evidence. For example, someone might have a memory as of buying salt on the last trip to the grocery store. This memory can support the belief that there is salt in the kitchen. This evidence is obviously characterized by a gap, for one might have a memory that is confused as to whether it was the last trip or some other trip to the grocery store on which the salt was bought. One might also think that one remembers buying salt when in fact one was merely dreaming, or buying something similar. So thinking that one remembers buying the salt on the last trip to the grocery store is compatible with its being false that one bought salt on that last trip to the grocery store.

One might think that this gap is very similar to the one arising in perceptual knowledge. In the case of perceptual knowledge we have a sensory experience that indicates something as true, and that is compatible with what it indicates being false. Is it not similarly the case that in memory we have a memory that indicates that something is true, even though what it indicates as true might be false? The situation is in fact not as parallel as it seems. In memory knowledge is so to speak merely preserved, but not generated. Knowledge on the basis of memory is passed on from one time to another and cannot be created. Let us explain this.

It seems that we can only come to know something through memory that we also knew when the memory was created. There are obvious exceptions to that, as when we need to combine two remembered things to gain some new knowledge. However, let me explain the case for which what I have just stated is true. Suppose that you were a witness to a car accident. Suppose that you have a vivid memory as of a red car crashing into a blue car. Now suppose that as it happens there is indeed a red and a blue car involved, but as a matter of fact the blue car crashed into the red car. You have somehow confused the situation, and we can imagine various reasons as to why that would be the case. Now it seems that since you did not know when the memory was created what you now remember, you also cannot come to know it now on the basis of the memory. Suppose alternatively that it is true that the red car crashed into the blue car and you have a vivid memory of it. However, your memory was created through overhearing some other

witness, who is very unreliable, claiming that this is what happened. In that situation you had at the time your memory was created a deficient justification for the true belief. Thus you had no knowledge, and the memory cannot give you any knowledge now.

The idea is thus that memory can only preserve knowledge, but not by itself create knowledge or turn something into knowledge that was not knowledge at the time at which the memory was created. There are obvious exceptions to that as when we through examining a memory realize some detail that had escaped our attention at the time at which the memory was created. Another exception is when we combine several separate memories, and they together allow us to know by inference something else. Except for such special circumstances, if a person has at the time at which a memory created no knowledge of what the memory indicates, then the person cannot come to know on the basis of this memory alone what the memory indicates at a later point.

If this seems right then there are actually two gaps involved in knowledge through memory. First, there is the gap that characterizes the situation in which the memory is created. Second, there is the gap created by questions about whether the knowledge is accurately preserved. In a simple example, a student learning a historical fact might experience both gaps. It might be that his book is inaccurate or that when he read the book he switched lines, or something similar. It might also be that he lost the memory or that it got confused. The first situation would be one where the gap in the original situation in which the memory is created becomes a problem. The second situation would

be one where the gap in preserving the knowledge becomes a problem. Only the second gap is one that would be studied when we examine memory as a source of knowledge.

Now it does not seem that there is a physical regularity that would help rendering this second gap compatible with knowledge. We would need a tight connection between what preserves the knowledge and what renders the belief indicated by the memory true. There does not seem to be a physical regularity that links the memory capacity to what renders it true that the blue car crashed into the red car. Memory as involving the storage and preserving of knowledge is not somehow part of a physical regularity between it and cars crashing into one another. There seems to be such a physical causal regularity in the case of sensory experience. But the preserving of a memory created by a sensory experience seems to be causally removed from the cause of the sensory experience. Or so it seems at least.

A similar account can be given for knowledge by testimony. Testimony passes knowledge from one person to another person, similarly to memory passing knowledge from one time to another time. If the person giving testimony has no knowledge then the person receiving the testimony cannot gain knowledge on the basis of the testimony alone. Knowledge is not created through testimony. It is not the case that if a witness confuses the red and the blue car then the person learning about the accident through testimony can somehow gain knowledge of which car crashed into which other car. Testimony can only pass on knowledge; it cannot by itself create new knowledge.

If this is correct, there are again two gaps in knowledge through testimony. First there is the gap in the situation in which the person giving the testimony gains the knowledge. Second, there is the gap in the situation where the knowledge is passed on in testimony. Someone interviewing a witness to a car accident might experience both gaps. For example, if the accident occurred late at night in front of a bar with all the witnesses being in various stages of intoxication, then this leads to a situation where the first gap gives rise to a problem. There might also be various reasons that can render the second gap problematic. Someone might be nervous or not expressing himself very well, or someone might be hard of hearing or confuse what the witness said, or something similar. All of these would be problems due to the second gap.

Again only the second gap characterizes knowledge through testimony. And again it does not seem that there is a physical regularity ensuring communication of testimony that is linked to whether the blue car crashed into the red car. What happens in the physical world causes sensory experiences, but it does not seem that it also causes proper communication of knowledge through testimony. So it does not seem that a luck-precluding condition as understood in the dissertation can help render the second gap compatible with knowledge.

If we want to summarize this brief examination of other kinds of knowledge of the physical world, it seems that they all are characterized by a gap, however, the gap cannot

be rendered compatible with knowledge by a luck-precluding condition as understood in the dissertation. Let us now turn to knowledge of non-physical things. As it seems, if there is at all a gap in these kinds of knowledge it is very different from the gap as understood in the dissertation. It is also obvious that it is not a physical regularity and thus not a luck-precluding condition as understood in the dissertation that could solve the problem.

Let us note a number of differences between knowledge of something physical and knowledge of something non-physical. In the case of knowledge of a language or conceptual scheme, we would if in doubt probably consult a dictionary or someone else who might be knowledgeable. In the case of mathematical knowledge different mathematicians need to verify a proof for it to become accepted as a proof. In the case of philosophical knowledge, intuitions need to be shared and arguments need to be considered strong by different philosophers. Our practices of accepting or verifying something as knowledge in these areas always involve the approval of others.

As it seems, in cases of such non-physical knowledge we can easily imagine one person to be in error, however, it is very difficult to see how a whole community over some time can be in error in considering something to be knowledge. In mathematics a community can be in error as to whether a conjecture is correct, but it is very difficult and rare if a proof that is accepted by the mathematical community as correct turns out to be wrong. Something similar is true with a language community, where we can easily see how one

person can be wrong about the meaning of a word, but it is very hard to see how a whole community can be wrong about what a word means. In philosophy widespread agreement about something is very rare, and philosophy seems to be more about thinking about philosophical things than somehow convincing everybody. However, if we take for example the Gettier intuition, which is very widely shared, it is again not clear how the whole philosophical community could be wrong about its being the case that a Gettier situation is incompatible with knowledge.

It seems that this reflects something about the gap that characterizes knowledge of such non-physical things. The gap is not one between evidence and true belief but rather whether one is not confused about whether a premise is known or whether a conclusion follows from some premises. Someone might think that a premise is known or that a conclusion follows when there is some subtle error in the apparent proof. Now this kind of error can best be detected by a community of people. In contrast, in a case of knowledge of the physical world, the physics community might agree on certain evidence and that it suggests that a certain claim is true, when in fact the claim is false. Whatever errors one might make in constructing mathematical proofs or in compiling a dictionary or in developing a philosophical argument, it is extremely difficult and rare to imagine a whole community to agree on it and yet be wrong about it.

What these remarks are supposed to show is that knowledge of non-physical things is very different from knowledge of physical things, and that if there at all is a gap, it is not

the same kind of gap as in the case of perceptual knowledge. Furthermore, if there is a gap it is clear that it is not rendered compatible with knowledge by a physical regularity and thus not by a luck-precluding condition as understood in the dissertation.

Thus it seems that perceptual knowledge as discussed in the dissertation is a special kind of knowledge that is different from other kinds of knowledge of the physical world and also of non-physical things. There does not seem to be a way of applying the account of perceptual knowledge to other kinds of knowledge. If this is correct, then epistemology is more fragmented than it is usually assumed to be. We would need different accounts for different kinds of knowledge. There might be some similarities, but as far as I can see a set of necessary and sufficient conditions will have to be tailored to specific kinds of knowledge rather than somehow try to capture all kinds of knowledge.

6.2 Other unfinished work

In this section I will point to some work that is left unfinished by the dissertation. We will begin by looking at some elements of the account of perceptual knowledge that are left vague. If no account of these vague elements can be given, this would constitute an objection to the account of perceptual knowledge. We will then look at some related projects that could strengthen the account of perceptual knowledge if they are worked out in detail. Lastly, I will sketch some projects that are not directly related to the dissertation, but that might benefit from the ideas of this account of perceptual knowledge.

Let us then begin by looking at some elements of the account of perceptual knowledge that have been left vague. The luck-precluding condition is a regularity. There is considerable vagueness as to what would count as a suitable regularity. How many exceptions are permitted? How many exceptions are permitted in the close proximity of the object of which we gain perceptual knowledge? Does the answer to this question depend on the physical location of the epistemic agent, or his past locations and travels? I am not sure how to answer these questions, but would like to point out that the same problem is faced by anyone who renders a statistical regularity required for knowledge. Since most externalist accounts include such statistical regularities, the account of perceptual knowledge in the dissertation has considerable company in facing this problem.

According to the account of perceptual knowledge developed in the dissertation, the decision to accept some evidence for a belief about the physical world needs to be reasonable. Furthermore, it is not reasonable if a little bit of epistemic work would uncover a defeater. However, what exactly does a little bit of epistemic work mean? Sometimes something obvious can be hidden from an epistemic agent, so to speak, in plain sight because of his preoccupation with something else. So does a little bit of epistemic work imply that it would take only a short time and a small amount of other resources? Or that there is no psychological barrier for finding the defeater? Or that those

with whom the epistemic agent interacts have done this epistemic work? Again, it is not clear how one should make this notion more precise.

In addition, a reasonable decision to accept some evidence depends also on an objective relationship of support between this evidence and the belief in question. Much of what I say in the dissertation involves the idea that it is an objective matter whether some evidence supports a particular belief, and what would count as a defeater. Furthermore, I suppose that we are generally able to see whether some evidence indeed supports a belief and whether something counts as a defeater. I am not sure how to make this more precise, however, if no account of an objective evidence-belief relationship could be given, this would be a defeater to the account of perceptual knowledge developed in the dissertation.

Let us now turn to some related projects that could strengthen the account of perceptual knowledge. One of the most puzzling aspects of an account of perceptual knowledge is how we get from a sensory experience to the beliefs that the sensory experience naturally suggests as true. This is what I have in the dissertation called “reading” a sensory experience. It seems that we need a bundle of different abilities for this to work. We need to have sensory experiences that have a distinct character that we can recognize. Then we need language and conceptual skills. We also need to understand the sensory experience as evidence rather than just respond blindly to a stimulus. It is very puzzling to try to understand how this works. Somehow we manage to see the world through the sensory

experiences, and we manage to adopt beliefs about the physical world on the basis of the sensory experiences.

It seems to me that this is a field that should be tackled by the philosophy of mind, and be informed by epistemology and cognitive science. My inclination is to say that not any response to sensory input should count as reading a sensory experience. I am also inclined to think that the content of a sensory experience is different in kind from a proposition, where a proposition is understood as the content of a sentence. However, this would be a project of its own, one that tries to understand how this aspect of perception works.

Another puzzling aspect of an account of perceptual knowledge is how we manage to combine different and conflicting bits of evidence. We often have evidence for and against a given belief, where the evidence comes in different types and different strengths. We somehow manage to see what the total set of evidence indicates, but it is not clear what exactly allows us to do so. How does one compare different bits of evidence that vary by kind and strength? How does one, so to speak, measure what the total set of evidence indicates? This is also some work left unfinished by the dissertation.

This leads to a further bit of unfinished work. The account of evidence-based justification does not go into detail as to whether it is an essentially foundationalist or coherentist picture that underlies it. The account of evidence-based justification could be made more

precise in either of these ways. There is thus unfinished work here as to how the details of the account of evidence-based justification are supposed to go. The account developed in the dissertation is in this respect generic, and I am more concerned with delineating its broad boundaries rather than going into details of this kind.

Let us now think about some unrelated projects that might benefit from some of the ideas in the dissertation. In Kantian ethics we find the idea that the right thing to do cannot depend on any actual consequences, for those consequences are not fully known. The idea seems to be that our moral knowledge has to be more secure than our knowledge of actual consequences of our actions ever can be. It might be interesting to think about this in terms of a gap between evidence and true belief. One might think that what Kant objects to is the gap between our evidence and our true beliefs in the case of knowledge of the actual consequences of our actions. It seems that Kant would want moral knowledge to be secure from this gap. So the idea would be that Kant rejects consequentialist ideas in part because of the problem of the gap when it comes to knowledge of the actual consequences of our actions. Instead his account of moral knowledge is offered as an attempt to render it secure from this kind of a gap.

Another unrelated project would be to think about what freedom means for us as practical agents on the basis of what freedom means for us as epistemic agents. As epistemic agent one is not tempted to say that somehow we need to be able to believe independently of any reasons and causes that would prompt and support the belief. For us as epistemic

agents freedom does not consist in somehow believing whatever we want, but rather believing on the basis of good reasons. Freedom for us as epistemic agents is more concerned with freedom from manipulation and deception that would undercut our epistemic agency. We are not free epistemic agents if we are trapped in circumstances where we are unable to acquire reasons that speak against our false beliefs.

One might take these ideas that are plausible for us as epistemic agents and apply them to us as practical agents. Freedom would then not have to be understood to be freedom to act against reasons or causes. Freedom for us as practical agents might also be understood as freedom from manipulation and deception where we would be trapped in circumstances that undermine our practical agency. At least one might try to thus give a somewhat unusual argument for a compatibilist notion of freedom by drawing on the plausibility of a similar position concerning our epistemic agency.²⁹

Again, none of the ideas mentioned in this conclusion are supposed to serve as accounts of these other kinds of knowledge or of related and unrelated projects. However, what this conclusion shows is that there is still a lot of philosophical work that needs to be done. Both in epistemology, neighboring, and unrelated fields, there is a lot of philosophical work waiting to be done.

²⁹ Matthias Steup suggests that if we accept that our actions are free in a compatibilist sense, then we should accept the same for our beliefs. His claim is thus that actions and beliefs can be understood to be free in the same compatibilist sense. What I suggest here is a different approach to the same project. The idea would be to see what it means for us as epistemic agents to be free, and then try to argue that this same notion of freedom could also render us as practical agents free. Steup takes the notion of freedom for us as

REFERENCES

- O. K. **Bouwsma (1942)**: *Moore's Theory of Sense-Data*. Reprinted in: O. K. Bouwsma: *Philosophical Essays*. University of Nebraska Press, Lincoln.
- Roderick M. **Chisholm (1966)**: *Theory of Knowledge*. Prentice-Hall: Englewood Cliffs, NJ.
- Roderick M. **Chisholm (1964)**: *The Myth of the Given*. In: R. Chisholm: *Philosophy*. Englewood Cliffs, NJ: Prentice Hall, pp. 261-286
- Stewart **Cohen (1984)**: *Justification and Truth*. *Philosophical Studies*, vol. 46, no. 3, pp. 279-295
- Keith **DeRose (1995)**: *Solving the Skeptical Problem*. *The Philosophical Review* 104, 1, pp. 1-7, 17-52
- Mylan **Engel, Jr. (1992)**: *Is Epistemic Luck Compatible with Knowledge?* *The Southern Journal of Philosophy*, vol. XXX, no. 2, pp. 59-75
- Edmund **Gettier (1963)**: *Is Justified True Belief Knowledge?* *Analysis* 23.6, pp. 121-123.
- Alvin **Goldman (1976)**: *What Is Justified Belief?* In: G. S. Pappas (ed.): *Justification and Knowledge*. Dordrecht: D. Reidel, pp. 1-23.
- John **Greco (2000)**: *Putting Skeptics in Their Place. The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge University Press.
- Terry **Horgan**, John **Tienson**, and George **Graham (2005)**: *Internal-World Skepticism and the Self-Presentational Nature of Phenomenal Consciousness*. (Draft)
- Peter **Klein (1971)**: *A Proposed Definition of Propositional Knowledge*. *The Journal of Philosophy* 68, 16, pp. 471-482.
- Keith **Lehrer (2006)**: *Coherence and the Truth Connection*. *Erkenntnis*, vol. 63, no. 3, pp. 413-423
- Keith **Lehrer (2000)**: *Theory of Knowledge*. Second Edition. Westview Press.

practical agents and applies it to us as epistemic agents. I suggest taking the notion of freedom for us as epistemic agents and applying it to us as practical agents. See Steup (2006).

Keith **Lehrer** and Thomas **Paxson, Jr. (1969)**: *Knowledge: Undefeated Justified True Belief*. The Journal of Philosophy, vol. LXVI, no. 8, pp. 225-237.

Keith **Lehrer (1971)**: *Why not Skepticism?* The Philosophical Forum, vol. II, no. 3, pp. 283-298.

John **Pollock** and Joseph **Cruz (1999)**: *Contemporary Theories of Knowledge*. Second Edition. Rowman and Littlefield Publishers

Duncan **Pritchard (2005)**: *Epistemic Luck*. Oxford University Press.

Baron **Reed (2001)**: *Epistemic Agency and the Intellectual Virtues*. Southern Journal of Philosophy, vol. 39, no. 4, pp. 507-526

Ernest **Sosa (1997)**: *Reflective Knowledge in the Best Circles*. The Journal of Philosophy, vol. XCIV, no. 8, pp. 410-430

Matthias **Steup (2006)**: *Doxastic Freedom*. (Draft)

Matthias **Steup (2001)**: *Epistemic Duty, Evidence, and Internality*. In: Matthias Steup (ed.): *Knowledge, Truth, and Duty*. Oxford University Press, pp. 134-148.

Barry **Stroud (1984)**: *The Problem of the External World*. Ch. 1 of Barry Stroud: *The Significance of Philosophical Skepticism*. Oxford: Clarendon Press.

Peter **Unger (1974)**: *An Argument for Skepticism*. Philosophical Exchange 1, 4

Ted **Warfield (1998)**: *A Priori Knowledge of the World. Knowing the World by Knowing Our Minds*. Philosophical Studies, vol. 92, no. 1-2, pp. 127-147.

Timothy **Williamson (2000)**: *Knowledge and its Limits*. Oxford University Press.

Linda **Zagzebski (1994)**: *The Inescapability of Gettier Problems*. Philosophical Quarterly, vol. 44, is. 174, pp. 65-74.